Università Ca'Foscari Venezia

Corso di Dottorato di ricerca

in Economia

ciclo 34

Tesi di Ricerca

Four Essays on Social Conventions

SSD: SECS-P/01

Coordinatore del Dottorato
ch. prof. Pietro Dino Enrico Dindo

Supervisori
ch. prof. Pietro Dino Enrico Dindo
ch. prof. Ennio Bilancini

Dottorando
Roberto Rozzi
Matricola 956470

Four essays on social conventions

Rozzi Roberto

To the love of my life, Xhiselda, for your unwavering support and inspiration, both within and outside of academia.

To my mother, Nilla, for showing me in ways I did not know, what bravery really means.

To my father, Massimo, for constantly showing me the true meaning of love and care.

To the rest of my family, Stefano, Elisa and Eleonora, for all the love that they gave me which I could never fully repay.

To my friends, my second family, the Bandarlogs, your companionship is a gift that I cherish deeply.

To my PhD peers, for being an immense source of inspiration and challenge.

To Itzhak, with which we shared one of the most exciting and testing experiences of my career.

To all the people I have met during these four intense years, for all the wonderful hints.

To my supervisors, Pietro and Ennio, for their invaluable guidance, encouragement, and the many valuable lessons I have learned from them.

To my reviewers, Marco LiCalzi, Heinrich Nax, and Salvatore Modica, for the useful and thoughtful comments that have enriched all of my chapters in ways I could not conceive alone.

This thesis was only possible because of you. It is with deep humility that I acknowledge my role as a mere instrument, translating into words (and formulas) the wisdom and knowledge that I have acquired through your closeness.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

A social convention can be perceived as a mode of conduct adopted by a population or group of individuals wherein they observe the actions of others and, guided by the desire for conformity, proceed to emulate those actions (Bicchieri, 2005). Typical examples of social conventions include greeting rituals or driving on a side of the road. The latter example is particularly noteworthy as it is a behavior that entire populations adopt and sustain, even in the absence of a law. Some scientists have put forth the argument that also language can be considered as a social convention to a certain extent (Lewis, 2008). Conducts between peer groups are another example of social conventions (Borsari and Carey, 2003): friends within a social group may adopt a particular behavior due to social adaptation, whereby they conform their actions to align with their peers. Although social conventions are primarily of sociological interest, they have been widely investigated across diverse academic disciplines, including economics, underscoring their significance as a fundamental aspect of human behavior. For example, friends' activities (that are guided by social conventions) may drive the consumption behavior of that group of friends, or they may influence how those people appear in an interview (and, therefore, their possibility of getting a job). Understanding the reasons behind the formation of social conventions is also crucial from a policy-making perspective. Indeed, many economic phenomena are influenced by social conventions, and therefore, a policymaker that wants to change a given economic behavior may want to influence social conventions to change such behavior.

In this thesis, I analyze how and why social conventions emerge across distinct contexts and their impact on different economic environments. In the first and second chapters, I employ evolutionary game theory techniques together with stochastic stability to study the formation of social conventions in two different strategic situations: a coordination game and a conflict game. In the third chapter, I conducted two field experiments to assess how people react to different shares of the population engaging in a particular behavior, which is crucial in understanding the emergence of social conventions. In the fourth chapter, I use a signaling game to study the relevance of homophily in guiding the formation of social groups among employees and the consequent impact on labor market outcomes.

Applied economic theory and experiments are widely employed tools to study the formation and the consequences of social conventions.

Coordination games have a particular byte in studying social conventions since they are a simple and intuitive tool that mimics the social pressure of certain phenomena. Specifically, in a coordination game, multiple equilibria can represent different conventions, and stochastic stability has been successfully used over the years in evolutionary game theory to capture which convention is the most likely to happen in a context where more than one convention is plausible (Freidlin and Wentzell, 1998; Young, 1993a; Kandori et al., 1993; Ellison, 2000). According to stochastic stability, the convention that is the most likely to happen is the one that requires fewer errors to be reached from all the other conventions and more errors to be left (where errors are deviating behaviors of agents, not following the convention).

In the field of coordination games, the Language Game has been widely used to study the rise of conventions between two groups with different tastes (Neary, 2012; Neary and Newton, 2017). As mentioned previously, these situations can shape the consumption behaviors of agents. However, the importance of these games may extend beyond these situations; for example, they can also be used to study the formation of conventions between two ethnic groups, providing insights into the reasons beyond the integration process (Goyal et al., 2021; Tanaka et al., 2018; He and Wu, 2020; Carvalho, 2017). In this kind of context, learning the opponent's type at first glance could also be costly. For example, it may take time and energy to understand another agent's preferences or tastes. Due to this factor, the convention that emerges between two groups of agents may also depend on the cognitive cost of learning the opponent's type (Güth and Kliemt, 1994; Berger and De Silva, 2021).

On the other hand, evolutionary game theory has been studying conflict games ever since its beginning (Smith and Price, 1973). Specifically, Hawk-Dove games are used to study the formation of conventions when agents compete for a resource (Foley et al., 2018, 2021). These diverse applications of this kind of game demonstrate its flexibility and power as a framework for understanding complex social and economic systems. In particular it has been used in biology (Smith and Price, 1973), economics (Lipnowski and Maital, 1983; Herold and Kuzmics, 2020), or political science (Brams and Kilgour, 1987; Baliga and Sjöström, 2012, 2020). It is relevant to study these situations since in Hawk-Dove games, agents must choose between fighting aggressively for a resource or backing down and avoiding conflict. By analyzing the dynamics of these games, we can gain insights into the conditions under which a society may become more or less aggressive over time. Furthermore, in this kind of games, boundedly rational agents may influence the long-run equilibrium in many ways: the ability and rationality disparities among agents can play a crucial role in explaining the selection of different equilibria (Bilancini et al., 2022) and even unexpected convergence over the long-run (Arigapudi et al., 2021). Unexpected convergence may be good or bad for the welfare of society depending on factors such as the harshness of conflict. In addition, it is crucial to investigate the evolutionary basis of the coexistence of heterogeneous agents in these games, such as when agents have different theories of mind (Mohlin, 2012). This involves determining which type of agent is more likely to survive based on their fitness, and whether this coexistence affects canonical predictions.

In the experimental economics literature, the importance of descriptive social norms in guiding many economic phenomena has been documented in the literature across different fields such as charitable giving (Frey and Meier, 2004; Agerström et al., 2016), intention to vote (Gerber and Rogers, 2009), or tax evasion (Bott et al., 2020). The contribution of these experiments is important from different points of view. In the first place, different types of social conventions may form depending on how people react to different shares of the population adopting a certain behavior. Such an importance have been proven in early theoretical works such as Schelling (1971), or Granovetter (1978), and more recently in experimental works like Centola and Baronchelli (2015), Centola et al. (2018), or Andreoni et al. (2021). In these contexts, it is important to determine whether homogeneous behavior is more likely to be a long-run equilibrium or whether heterogeneity is more prevalent. In the case of homogeneous behavior, all agents will follow the same behavior in the long-run, while in the case of heterogeneous case, agents may follow different behaviors. Understanding which of these two outcomes is more likely is also important from a policymaker's perspective. Indeed, public policy decisions depend on whether a policymaker should expect long-term homogeneous or heterogeneous convergence (Nyborg et al., 2016; Efferson and Vogt, 2018).

The study of signaling games (Spence, 1973), together with the one on costly information acquisition (Grossman and Stiglitz, 1980), have been successful over the past decades in explaining inefficiencies in labor market outcomes and other kinds of markets. Lately, many scholars have been studying models that combine both approaches (Glazer and Rubinstein, 2004; Caillaud and Tirole, 2007; Gentzkow and Kamenica, 2014; Argenziano et al., 2016). Specifically, certain models study situations where the receiver has to pay to observe the signal sent by the senders; this may be the case of an employer that has to exert effort (or spend time) to screen candidates for a job position (Bilancini and Boncinelli, 2018c,b; Fosgerau et al., 2020). In these situations, the screening process is costly; therefore, the employer may decide to screen only a limited amount of information regarding the candidates, or she may decide to rely on information that does not require effort. Given that social attitudes of candidates are often more easily observable than their skills, agents' social preferences could affect how they are assigned to job tasks or, in general, labor market outcomes since the employer may rely on these easily observable social characteristics when screening candidates. This process could lead to biased or sub-optimal outcomes in the labor market (Austen-Smith and Fryer Jr, 2005; Kim and Loury, 2012; Bowles et al., 2014; Bolte et al., 2020; Okafor, 2020; Jackson, 2021). Specifically, employees may send signals through conventions (e.g. their clothes): the employer may only judge candidates based on these signals (conventions, or clothes), inducing them to not invest in their skills.

My first chapter presents a theoretical study on the formation of social conventions between two different groups in the presence of information costs. In this chapter, I use stochastic stability to assess the stability of long-run outcomes in a Language Game, where learning the opponent's type requires effort. My findings indicate that the cost of learning the opponent's type may influence the chances of agents coordinating. Specifically, there is a twofold advantage in being the majority or the group with stronger preferences. When the cost is high, nobody learns their opponent's type, and the majority (the group with stronger preferences) always plays their favorite action. When the cost is low, the majority (the group with stronger preferences) never pays for the information and always plays their favorite action. However, when the cost is high and both groups have strong preferences, a convention in which they do not coordinate when they meet may arise.

In my second chapter, I focus on another type of game, the Hawk-Dove, which is a conflict game. Precisely, I use a theoretical model to study the formation of social conventions in a Hawk-Dove game in the presence of agents with different behavioral rules. I study both the stochastic stability of strategies for all different population compositions and then the stability of each population level based on the stochastically stable distribution of strategies for that population level. I find that the less intelligent type of agent outnumbers the more intelligent one in the ultra long-run. Considering this endogenous selection mechanism, I find non-standard convergence for the Hawk-Dove. Lastly, I find that the less intelligent types behave aggressively when the conflict is harsh and defensively when it is mild. Due to this result, agents are better of than under the canonical prediction when conflict is mild (since they live in a peaceful population of doves), while they are worse off when conflict is harsh (since they live in an aggressive population of hawks).

In my third chapter, I step away from the theoretical perspective to study relevant implications for the formation of social conventions with two field experiments. Such an approach complements the theoretical one since it allows testing relevant implications from theoretical models. Indeed, my third chapter is an experimental investigation of heterogeneous responses to descriptive social norms. Using two semi-continuous randomized treatments, we show that the subjects' response to different shares of the population acting in a certain way is context-dependent. In the first experiment, we focus on face-masks, randomizing the number of people wearing the mask in the subject's immediate environment. In a second experiment, we randomize the number of people turning the camera on in the subject's immediate environment. In the first experiment, We find evidence of a quadratic relation between the different treatments, while in the second, we found evidence of linearity between the treatments. When embedded in plausible dynamical models, our estimates suggest an interior convergence in both experiments, i.e. heterogeneous behavior in the long-run. Therefore, both of our studies predict the importance of enforcing these kinds of behaviors by law and not relying on social pressure.

In my fourth chapter, I abandon dynamic games to study the implications of the social attitudes of agents on their labor market outcomes. I use a signaling model with costly information acquisition to study the effect on labor market outcomes of the formation of social conventions due to homophily. Specifically, the employer only sees candidates' social groups if she does not buy the information: social groups may be informative about candidates' abilities and skills depending on how candidates form the social groups. The receiver earns utility from matching candidates' types and skills with the right job task. There are two types of candidates. They decide their skill level and form social groups based on their level of homophily. The receiver only observes candidates' social groups unless she screens them (at a given cost). Our results show that, due to homophily, the receiver never needs to screen candidates to check their abilities: this effect leads to candidates' moral hazard. Indeed, due to homophily, candidates with similar abilities hang out together; therefore, the receiver infers their abilities from their social group (e.g. from their clothes), and she does not screen them.

Note that the first and the second chapters are single author papers: the first chapter is published on Games (Rozzi, "Competing conventions with costly information acquisition." Games 12.3 (2021): 53.). The second chapter is the last work I did during my PhD. The third chapter is joint work together with Itzhak Rasooly: we started to talk about this project after an informal chat during an online conference in June 2021, and we carried it on until we finally did our first experiment in February 2022. The fourth chapter is joint work with Ennio Bilancini and Leoardo Boncinelli: this is the second project I have worked on during my PhD. Ennio and Leonardo asked me to work on this project together: I carried on the modeling choices and the solutions to the model during my third year of the PhD.

Chapter 2

Competing conventions under costly information acquisition

Abstract

We consider an evolutionary model of social coordination in a 2×2 game where two groups of players prefer to coordinate on different actions. Players can pay a cost to learn their opponent's group: if they pay it, they can condition their actions concerning the groups. We assess the stability of outcomes in the long-run using stochastic stability analysis. We find that three elements matter for the equilibrium selection: the group size, the strength of preferences, and the information's cost. If the cost is too high, players never learn the group of their opponents in the long-run. If one group is stronger in preferences for its favorite action than the other, or its size is sufficiently large compared to the other group, every player plays that group's favorite action. If both groups are strong enough in preferences, or if none of the groups' sizes is large enough, players play their favorite actions and miscoordinate in inter-group interactions. Lower levels of the cost favor coordination. Indeed, when the cost is low, in inside-group interactions, players always coordinate on their favorite action, while in inter-group interactions, they coordinate on the favorite action of the group that is stronger in preferences or large enough.

Keywords: coordination; conventions; evolution; stochastic stability; costly information acquisition.

JEL Classification Codes: C72; C73

Notes. This chapter is a single author work published on Games. Rozzi, "Competing conventions with costly information acquisition." Games 12.3 (2021): 53. In this version of my chapter, I moved Table A1 to the Appendix, and I rewrote Theorem 4 for clarity.

2.1 Introduction

Since the seminal contribution of Kandori et al. (1993), evolutionary game theorists have used stochastic stability analysis and 2×2 coordination games to study the formation of social conventions (Lewis, 2008 and Bicchieri, 2005 are classical references on social conventions from philosophy, while for economics, see Schelling, 1980, Young, 1996, and Young, 2020). Some of these works focus on coordination games such as the battle of sexes: a class that describes situations in which two groups of people prefer to coordinate on different actions. In this framework, the long-run convention may depend on how easily people can learn each other's preferences.

Think about Bob and Andy, who want to hang out together: they can either go to a football match or to the cinema. Both Andy and Bob prefer football, but they do not know what the other prefers. In certain contexts, learning each other's preferences may require too much effort. In these cases, if Bob and Andy know that everybody usually goes to the cinema, they go to the cinema without learning each other's preferences. In other situations, learning each other's preferences may require a small effort (for instance, watching each other's Facebook walls). In this case, Bob and Andy learn that they both prefer football, so they go to a football match together.

In this work, we contribute to the literature on coordination games. We show which conventions become established between two groups of people different in preferences if people can learn each other's preferences by exerting an effort. We do so, formalizing the previous example and studying the evolution of conventions in a dynamic setting. We model the coordination problem as a repeated language game (Neary, 2012): we use evolutionary game theory solution concepts and characterize the long-run equilibrium as the stochastically stable state (see Foster and Young, 1990, Kandori et al., 1993 and Young, 1993a).

We consider a population divided into two groups, which repeatedly play a 2×2 coordination game. We assume that one group is larger than the other and that the two groups differ in preferences towards the coordination outcomes. At each period, players can learn the group of their opponent if they pay a cost. Such a cost represents the effort to exert if they want to learn their opponent's group. If they pay this cost, they can condition the action to the player they meet. If they do not pay it, they can only play the same action with every player. Given this change in the strategic set, we introduce a new possible perturbation. Players can make a mistake in the information choice and a mistake in the coordination choice. We model two situations: one where the cost is equal to zero, and players always learn their opponent's group, and one where the cost is strictly positive and players can learn their opponent's group only if they pay that cost. Players decide myopically their best reply based on the current state, which is always observable. We say that a group has a stronger preference for its favorite action than the other if it assigns higher payoffs to its favorite outcome or lower payoffs to the other outcome compared to the other group.

We find that cost level, strength in preferences, and group size are crucial drivers for the long-run stability of outcomes. Two different scenarios can happen, depending on the cost. Firstly, low cost levels favor coordination: players always coordinate on their favorite action with players of their group. If one group has a stronger preference for its favorite action or its size is sufficiently large compared to the other, every player plays the action preferred by that group in inter-group interactions. Interestingly, players from the group that is stronger in preferences never need to buy the information because they play their favorite action with everyone, while players from the other group always need to buy it.

Secondly, when the cost is high, players never learn the group of their opponents, and they play the same action with every player. Some players coordinate on one action that they do not like, even with players of their group. Indeed, we find that when one group is stronger in preferences than the other for its favorite action, or if its size is sufficiently large compared to the other, every player coordinates on that group's favorite action. Even worse, the two groups may play their favorite action and miscoordinate in inter-group interactions. We find that this outcome occurs when both groups have strong enough preferences for their favorite action or if the two groups are sufficiently close in size.

Neary (2012) considers a similar model, where each player decides one single action valid for both groups. Hence, it is as if learning an opponent's group requires too much effort, and no player ever learns it. Given this scenario, Neary's results are the same as in our analysis when the cost is high.

It is helpful to highlight our analysis with respect to the one proposed by Neary, from which we started. We firstly enlarge Neary's analysis to the case when players learn their opponent's group at zero cost. In this case, only states where all the players in one group buy the information can be stochastically stable: this result was not possible in the analysis of Neary. Overall, controlling for the cost equal to zero may be seen as a robustness exercise; nevertheless, we find that the model is more tractable under this specification than under Neary's one. Indeed, if the cost is equal to zero, we can consider inter-group dynamics separated from inside-group ones, and hence, we can consider two absorbing states at a time.

The behavioral interpretation is similar for high and low levels of the cost: either the minority adapts to the majority, or the weaker group in preferences adapt to the strongest. Indeed, when the cost is low, the weakest group always needs to buy the information, while the strongest group does not, since it plays its favorite action with everyone. Similarly, when the cost is high, everybody will play the action favored by the strongest group in preferences in the long-run. However, comparing the high-cost case with the low-cost case enriches the previous analysis. From this comparison, we can say that reducing the cost of learning the opponent's group increases the probability of inter-group coordination in the long-run. Indeed, inter-group miscoordination does not occur without incomplete information and a high cost. Unlike in Neary, strength in preferences or group size alone does not cause inter-group miscoordination.

The paper is organized as follows: In Section 2.2, we explain the model's basic features. In Section 2.3, we determine the results for the complete information case where the cost is 0. In Section 2.4, we derive the results for the case with incomplete information and costly acquisition. We distinguish between two cases: low cost and high cost. In Section 2.5, we discuss results, and in Section 2.6, we conclude. We give all proofs in the Appendix A and we give the intuition during the text.

2.2 Model

We consider N players divided into two groups A and B, $N = N_A + N_B$. We assume $N_A > N_B + 1$ and $N_B > 1$. Each period, players are randomly matched in pairs to play the 2×2 coordination game represented in Tables 2.1–2.2. Matching occurs with uniform probability, regardless of the group. Tables 2.1(a) and 2.1(b) represent inside-group interactions, while Table 2.2 represents inter-group interactions (group A row player and group B column player). We assume that $\Pi_A > \pi_A$, and thus, we name a the favorite action of group A. Equally, we assume $\Pi_B > \pi_B$, and hence, b is the favorite action of group B. We do not assume any particular order between Π_B, and Π_A. However, without loss of generality, we assume that $\Pi_A + \pi_A = \Pi_B + \pi_B$. Consider $K \in \{A, B\}$, and $K' \neq K \in \{A, B\}$. We say that group K is stronger in preferences for its favorite action than group K' if $\Pi_K > \Pi_{K'}$ or equivalently $\pi_K < \pi_{K'}$.

(a) Inside group A

	a	b
a	Π_A, Π_A	$0, 0$
b	$0, 0$	π_A, π_A

(b) Inside group B

	a	b
a	π_B, π_B	$0, 0$
b	$0, 0$	Π_B, Π_B

Table 2.1: Inside-group interactions

Each period, players choose whether to pay a cost to learn their opponent's group or not

	a	b
a	Π_A, π_B	$0, 0$
b	$0, 0$	π_A, Π_B

Table 2.2: Inter-group Interactions.

before choosing between action a and b. If they do not pay it, they do not learn the group of their opponent, and they play one single action valid for both groups. If they pay it, they can condition the action on the two groups. We call information choice the first and coordination choice the second.

Consider player $i \in K$. τ_i is the information choice of player i: if $\tau_i = 0$, player i does not learn the group of her/his opponent. If $\tau_i = 1$, player i pays a cost c and learns the group. We assume that $c \geq 0$. $x_{0i} \in \{a, b\}$ is the coordination choice when $\tau_i = 0$. If $\tau_i = 1$, $x_{1i}^K \in \{a, b\}$ is the coordination choice when player i meets group K, while $x_{1i}^{K'} \in \{a, b\}$ is the coordination choice when player i meets group K'.

A pure strategy of a player consists of her/his information choice, τ_i, and of her/his coordination choices conditioned on the information choice, i.e.,

$$s_i = \left(\tau_i, x_{0i}, x_{1i}^K, x_{1i}^{K'}\right) \in \mathcal{S} = \{0, 1\} \times \{a, b\}^3.$$

Each player has sixteen strategies. However, we can safely neglect some strategies because they are both payoff-equivalent (a player earns the same payoff disregarding which strategy s/he chooses) and behaviorally equivalent (a player earns the same payoff independently from which strategy the other players play against her/him).

We consider a model of noisy best-response learning in discrete time (see Kandori et al., 1993, Young, 1993a).

Each period $t = 0, 1, 2, \ldots$, independently from previous events, there is a positive probability $p \in (0, 1)$ that a player is given the opportunity to revise her/his strategy. When such an event occurs, each player who is given the revision opportunity chooses with positive probability a strategy that maximizes her/his payoff at period t. $s_i(t)$ is the strategy played by player i at period t. $U_s^i(s', s_{-i})$ is the payoff of player i that chooses strategy s' against the strategy profile s_{-i} played by all the other players except i. Such a payoff depends on the random matching assumption and the payoffs of the underlying 2×2 game. At period $t + 1$, player i chooses

$$s_i(t+1) \in \arg\max_{s' \in \mathcal{S}} U_s^i(s', s_{-i}(t)).$$

If there is more than one strategy that maximizes the payoff, player i assigns the same probability to each of those strategies. The above dynamics delineates a Markov process that is ergodic thanks to the noisy best response property.

We group the sixteen strategies into six analogous classes that we call behaviors. We name behavior a (b) as the set of strategies when player $i \in K$ chooses $\tau_i = 0$, and $x_{0i} = a$ (b). We name behavior ab as the set of strategies when player i chooses $\tau_i = 1$, $x_{1i}^K = a$, and $x_{1i}^{K'} = b$, and so on. Z is the set of possible behaviors: $Z = (a, b, ab, ba, aa, bb)$. $z_i(t)$ is the behavior played by player i at period t as implied from $s_i(t)$. $z_{-i}(t)$ is the behavior profile played by all the other players except i at period t as implied from $s_{-i}(t)$. Note that behaviors catch all the relevant information as defined when players are myopic best repliers. $U_z^i(z', z_{-i}(t))$ is the payoff for player i that chooses behavior z' against the behavior profile $z_{-i}(t)$. Such a payoff depends on the random matching assumption and the payoffs of the underlying 2×2 game. The dynamics of behaviors as implied by strategies coincide with the dynamics of behaviors, assuming that players myopically best reply to a behavior profile. We formalize the result in the following lemma.

Lemma 1. *Given the dynamics of $z_i(t+1)$ as implied by $s_i(t+1)$, it holds that $z_i(t+1) \in \underset{z' \in Z}{\text{argmax}}\, U_z^i(z', z_{-i}(t))$.*

Consider a player $i \in A$ such that the best thing to do for her/him is to play a with every player s/he meets regardless of the group. In this case, both $(0, a, a, b)$ and $(0, a, b, b)$ maximize her/his payoff. In contrast, $(0, b, a, b)$ does not maximize her/his payoff since in this case, s/he plays b with every player s/he meets. Moreover, the payoff of player i is equal whether $s_{-i} = (0, a, a, b)^{N-1}$ or $s_{-i} = (0, a, b, b)^{N-1}$ but different if $s_{-i} = (0, b, a, b)^{N-1}$. Therefore, all the strategies that belong to the same behavior are payoff equivalent and behaviorally equivalent.

A further reduction is possible because aa (bb) is behaviorally equivalent to a (b) for each player. The last observation and the fact that we are interested in the number of players playing a with each group lead us to introduce the following state variable. We denote with n^{AA} (n^{BB}) the number of players of group A (B) playing action a with group A (B), and n^{AB} (n^{BA}) the number of players of group A (B) playing action a with group B (A). We define states as vectors of four components: $\omega = \left\{n^{AA}, n^{AB}, n^{BA}, n^{BB}\right\}$, with Ω being the state space and $\omega_t = \left\{n_t^{AA}, n_t^{AB}, n_t^{BA}, n_t^{BB}\right\}$ the state at period t. At each t, all the players know all the components of ω_t. Consider player i playing behavior $z_i(t)$ at period t. $U^i_{z_i(t)}(z', \omega_t)$ is the payoff of i if s/he chooses behavior z' at period $t+1$ against the state ω_t. All that matters for a decision-maker is ω_t and $z_i(t)$. We formalize the result in the following lemma.

Lemma 2. *Given the dynamics of ω_{t+1} generated by $z_i(t+1)$, it holds that $U^i_z(z', z_{-i}(t)) = U^i_{z_i(t)}(z', \omega_t)$. Moreover, $U^i_a(z', \omega_t) = U^i_{aa}(z', \omega_t) = U^i_{ab}(z', \omega_t)$, and $U^i_b(z', \omega_t) = U^i_{bb}(z', \omega_t) = U^i_{ba}(z', \omega_t)$.*

If players are randomly matched, it is as if each player plays against the entire population. Therefore, each player of group K myopically best responds to the current period by looking at how many players of each group play action a with group K. Moreover, a player that is given the revision opportunity subtracts her/himself from the component of ω_t where s/he belongs. If $i \in K$ is playing behavior a, aa or ab at period t, s/he knows that $n^{KK}_t - 1$ players of group K are playing action a with group K at period t.

Define with θ_{t+1} the set of players that are given the revision opportunity at period t. Given Lemma 2, it holds that ω_{t+1} depends on ω_t and on θ_{t+1}. That is, we can define a map $F(\cdot)$ such that $\omega_{t+1} = F(\omega_t, \theta_{t+1})$. The set θ_{t+1} reveals whether the players who are given the revision opportunity are playing a behavior between a, aa, and ab, or a behavior between b, bb, and ba. In the first case we should look at U^i_a, while in the second at U^i_b.

From now on, we will refer to behaviors and states following the simplifications described above.

We illustrate here the general scheme of our presentation. We divide the analysis into two cases: complete information and incomplete information. For each case, we consider unperturbed dynamics (players choose the best reply behavior with probability 1) and perturbed dynamics (players choose a random behavior with a small probability). First, we help the reader understand how each player evaluates her/his best reply behavior and which states are absorbing. Second, we highlight the general structure of the dynamics with perturbation and then determine the stochastically stable states. In the next section, we analyze the case with complete information, hence, when the cost is zero.

2.3 Complete information with free acquisition

In this section, we assume that each player can freely learn the group of her/his opponent when randomly matched with her/him. Without loss of generality, we assume that players always learn the group of their opponent in this case. We refer to this condition as free information acquisition. Each player has four possible behaviors as defined in the previous section. $Z = \{aa, ab, ba, bb\}$, with $a = aa$, and $b = bb$ in this case.

Define $\pi_a^K = \begin{cases} \Pi_A & \text{if } K = A \\ \pi_B & \text{if } K = B \end{cases}$ and $\pi_b^K = \begin{cases} \pi_A & \text{if } K = A \\ \Pi_B & \text{if } K = B \end{cases}$.

Equations (2.1)–(2.4) are the payoffs for a player $i \in K$ playing aa or ab at period t.

$$U_a^i(aa, \omega_t) = \frac{n_t^{KK} - 1}{N-1}\pi_a^K + \frac{n_t^{K'K}}{N-1}\pi_a^K, \tag{2.1}$$

$$U_a^i(ab, \omega_t) = \frac{n_t^{KK} - 1}{N-1}\pi_a^K + \frac{N_{K'} - n_t^{K'K}}{N-1}\pi_b^K, \tag{2.2}$$

$$U_a^i(ba, \omega_t) = \frac{N_K - n_t^{KK}}{N-1}\pi_b^K + \frac{n_t^{K'K}}{N-1}\pi_a^K, \tag{2.3}$$

$$U_a^i(bb, \omega_t) = \frac{N_K - n_t^{KK}}{N-1}\pi_b^K + \frac{N_{K'} - n_t^{K'K}}{N-1}\pi_b^K. \tag{2.4}$$

2.3.1 Unperturbed dynamics

We begin the analysis for complete information by studying the dynamics of the system when players play their best reply behavior with probability one.

We can separate the dynamics of the system into three different dynamics. The two regarding inside-group interactions, i.e., n_t^{AA} and n_t^{BB}, and the one regarding inter-group interaction, i.e., n_t^{AB} and n_t^{BA}. We call this subset of states $n_t^I = \left(n_t^{AB}, n_t^{BA}\right)$. Both n_t^{AA} and n_t^{BB} are one-dimensional; n_t^I instead is two-dimensional.

Lemma 3. *Under free information acquisition, $n_{t+1}^{AA} = F_1(n_t^{AA}, \theta_{t+1})$, $n_{t+1}^{BB} = F_4(n_t^{BB}, \theta_{t+1})$ and $(n_{t+1}^{AB}, n_{t+1}^{BA}) = F_{2,3}(n_t^{AB}, n_t^{BA}, \theta_{t+1})$.*

The intuition behind the result is as follows. If players always learn their opponent's group, the inter-group dynamics does not interfere with the inside-group and vice-versa. If player $i \in K$ is given the revision opportunity, s/he chooses x_{1i}^K only based on n_t^{KK}.

Consider a subset of eight states: $\omega^R = \{(N_A, N_A, N_B, N_B)$, $(0, N_A, N_B, N_B)$, $(N_A, N_A, N_B, 0)$, $(N_A, 0, 0, N_B)$, $(0, N_A, N_B, 0)$, $(N_A, 0, 0, 0)$, $(0, 0, 0, N_B)$ and $(0, 0, 0, 0)\}$.

Lemma 4. *Under free information acquisition, the states in ω^R are the unique absorbing states of the system.*

We call (N_A, N_A, N_B, N_B) and $(0, 0, 0, 0)$ Monomorphic States (MS from now on). Specifically, we refer to the first one as MS_a and to the second as MS_b. We label the remaining six as Polymorphic States (PS from now on). We call $(N_A, N_A, N_B, 0)$ PS_a and $(N_A, 0, 0, 0)$ PS_b. In MS, every player plays the same action with any other player; in PS, at least one group is conditioning the action. In MS_a, every player plays aa; in MS_b, every player plays bb. In PS_a, group A plays aa and group B plays ba. In PS_b, group A plays ab while group B plays bb. In both PS_a and PS_b, all players coordinate on their favorite action with their similar.

In the model of Neary, only three absorbing states were possible: the two MS and a Type Monomorphic State where group A plays aa and group B plays bb. The PS were not present in the previous analysis. We observe these absorbing states in our analysis, thanks to the possibility of conditioning the action on the group.

We can break the absorbing states in ω^R into the three dynamics in which we are interested. This simplification helps in understanding why only these states are absorbing. For instance, in inter-group interactions, there are just two possible absorbing states, namely (N_A, N_B) and $(0, 0)$. For what concerns inside-group interactions, N_A and 0 matter for n_t^{AA}, and N_B and 0 for n_t^{BB}. For each dynamic, the states where every player plays a or where every player plays b with one group are absorbing. In this simplification, we can see the importance of Lemma 3. As a matter of fact, in all the dynamics we are studying, there are just two candidates to be stochastically stable. This result simplifies the stochastic stability analysis.

2.3.2 Perturbed dynamics

We now introduce perturbations in the model presented in the previous section; that is, players can make mistakes while choosing their behaviors: there is a small probability that a player does not choose her/his best response behavior when s/he is given the revision opportunity. We use tools and concepts developed by Freidlin and Wentzell (1998) and refined by Ellison (2000).

Given perturbations, ω_{t+1} depends on ω_t, θ_{t+1} and on which players make a mistake among those who are given the revision opportunity. We define with ψ_{t+1} the set of players who do not choose their best reply behavior among those who are given the revision opportunity. Formally, $\omega_{t+1} = F(\omega_t, \theta_{t+1}, \psi_{t+1})$.

We use uniform mistakes: the probability of making a mistake is equal for every player and every state. At each period, if a player is given the revision opportunity, s/he makes a mistake with probability ε. In this section, we assume that players make mistakes only in the coordination choice: assuming $c = 0$, adding mistakes also in the information choice would not influence the analysis. Note that Lemma 3 is still valid under this specification.

If we consider a sequence of transition matrices $\{P^\varepsilon\}_{\varepsilon>0}$, with associated stationary distributions $\{\mu^\varepsilon\}_{\varepsilon>0}$, by continuity, the accumulation point of $\{\mu^\varepsilon\}_{\varepsilon>0}$ that we call $\mu^\star$, is a stationary distribution of $P := \lim_{\varepsilon\to 0} P^\varepsilon$. Mistakes guarantee the ergodicity of the Markov process and the uniqueness of the invariant distribution. We are interested in states which have positive probability in $\mu^\star$.

Definition 1. *A state $\bar{\omega}$ is stochastically stable if $\mu^\star(\bar{\omega}) > 0$ and it is uniquely stochastically stable if $\mu^\star(\bar{\omega}) = 1$.*

We define some useful concepts from Ellison (2000). Let $\bar{\omega}$ be an absorbing state of the unperturbed process. $D(\bar{\omega})$ is the basin of attraction of $\bar{\omega}$: the set of initial states from which the unperturbed Markov process converges to $\bar{\omega}$ with probability one. The radius of $\bar{\omega}$ is the number of mistakes needed to leave $D(\bar{\omega})$ when the system starts in $\bar{\omega}$. Define a path from state $\bar{\omega}$ to state ω' as a sequence of distinct states $(\omega_1, \omega_2, \dots, \omega_T)$, with $\omega_1 = \bar{\omega}$ and $\omega_T = \omega'$. $\Upsilon(\bar{\omega}, \omega')$ is the set of all paths from $\bar{\omega}$ to ω'. Define $r(\omega_1, \omega_2, \dots, \omega_T)$ as the resistance of the path $(\omega_1, \omega_2, \dots, \omega_T)$, namely the number of mistakes that occurs to pass from state $\bar{\omega}$ to state ω'. The radius of $\bar{\omega}$ is then

$$R(\bar{\omega}) = \min_{(\omega_1,\omega_2,\dots,\omega_T)\in\Upsilon(\bar{\omega},\Omega-D(\bar{\omega}))} r(\omega_1, \omega_2, \dots, \omega_T).$$

Now define the Coradius of $\bar{\omega}$ as

$$CR(\bar{\omega}) = \max_{\omega\notin D(\bar{\omega})} \min_{(\omega_1,\omega_2,\dots,\omega_T)\in\Upsilon(\omega,D(\bar{\omega}))} r(\omega_1, \omega_2, \dots, \omega_T)$$

Thanks to Theorem 1 in Ellison (2000), we know that if $R(\bar{\omega}) > CR(\bar{\omega})$, then $\bar{\omega}$ is uniquely stochastically stable.

We are ready to calculate the stochastically stable states under complete information.

Theorem 1. *Under free information acquisition, for N large enough, if $\frac{\pi_B}{\pi_A} < \frac{N_B}{N_A}$, then PS_b is uniquely stochastically stable. If $\frac{\pi_B}{\pi_A} > \frac{N_B}{N_A}$, then PS_a is uniquely stochastically stable.*

When the cost is null, players can freely learn the group of their opponent. Therefore, in the long-run, they succeed in coordinating on their favorite action with those who are similar in preference. Hence, n_t^{AA} always converges to N_A, and n_t^{BB} always converges to 0. This result rules out Monomorphic States and the other four Polymorphic States: only PS_a and PS_b are left. Which of the two is selected depends on strength in preferences and group size. Two effects determine the results in the long-run. Firstly, if $\pi_A = \pi_B$, PS_a is uniquely stochastically stable. The majority prevails in inter-group interactions if the two groups are equally strong in preferences.

Secondly, if $\pi_A \neq \pi_B$, there is a trade-off between strength in preferences and group size. If $\frac{\pi_B}{\pi_A} > \frac{N_B}{N_A}$, either group A is stronger in preferences than group B, or group A is sufficiently larger than group B. In both of the two situations, the number of mistakes necessary to leave PS_a is bigger than the one to leave PS_b: in a sense, more mistakes are needed to make b best reply for A players than to make a best reply for B players. Therefore, every player will play action a in inter-group interactions. Similar reasoning applies if $\frac{\pi_B}{\pi_A} < \frac{N_B}{N_A}$.

Interestingly, in both cases, only players of one group need to learn their opponent's group: the players from the group that is weaker in preferences or sufficiently smaller than the other.

Unlike in the analysis of Neary, if learning the opponent's group is costless, the Monomorphic States are never stochastically stable. This result is a consequence of the possibility to condition the action on the group. Indeed, if players can freely learn the opponent's group, they will always play their favorite action inside the group.

We provide two numerical examples to explain how the model works in Figures 2.1 and 2.2. We represent just n_t^I, hence, a two-dimensional dynamics. Red states represent the basin of attraction of $(0,0)$, while green states the one of (N_A, N_B). From gray states, there are paths of zero resistance both to $(0,0)$ and to (N_A, N_B). Any path that involves more players playing a within red states has a positive resistance. Every path that involves fewer people playing a within green states has a positive resistance. The radius of $(0,0)$ is equal to the coradius of (N_A, N_B), and it is the minimum resistance path from $(0,0)$ to gray states. The coradius of $(0,0)$ is equal to the radius of (N_A, N_B), and it is the minimum resistance path from (N_A, N_B) to gray states.

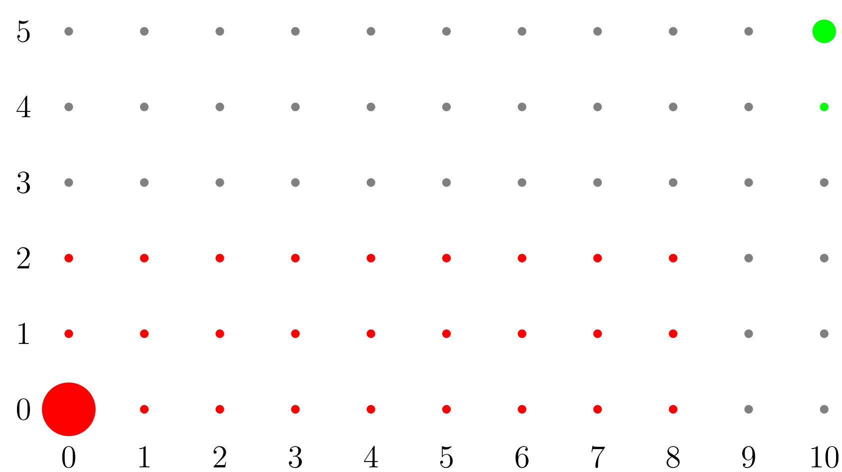

Figure 2.1: $PS_b = (0,0)$ is uniquely stochastically stable: $\frac{\pi_B}{\pi_A} < \frac{N_B}{N_A}$.

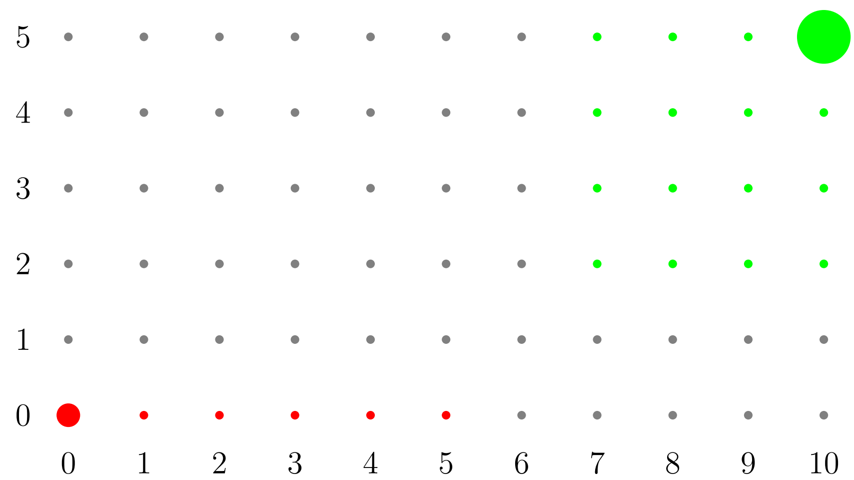

Figure 2.2: $PS_a = (10, 5)$ is uniquely stochastically stable: $\frac{\pi_B}{\pi_A} > \frac{N_B}{N_A}$.

Firstly, consider the example in Figure 2.1. $N_A = 10$, $N_B = 5$, $\pi_A = 8$, $\Pi_A = 10$, $\pi_B = 3$, $\Pi_B = 15$. Clearly, $\frac{\pi_B}{\pi_A} = \frac{3}{8} < \frac{5}{10} = \frac{N_B}{N_A}$. In this case $R(10, 5) = CR(0, 0) = 1$, while $R(0, 0) = CR(10, 5) = 3$. Hence, $(0, 0)$ is the uniquely stochastically stable state. We give here a short intuitive explanation. Starting from $(0, 0)$, the minimum-resistance path to gray states is the one that reaches $(0, 3)$. The minimum resistance path from $(10, 5)$ to gray states is the one that reaches $(9, 5)$. Hence, fewer mistakes are needed to exit from the green states than to exit from the red states, and $PS_b = (10, 0, 0, 0)$ is uniquely stochastically stable.

Secondly, consider the example in Figure 2.2. $N_A = 10$, $N_B = 5$, $\pi_A = 3$, $\Pi_A = 15$, $\pi_B = 8$, $\Pi_B = 10$. Note that $\frac{\pi_B}{\pi_A} = \frac{8}{3} > \frac{5}{10} = \frac{N_B}{N_A}$. In this case, $R(10, 5) = CR(0, 0) = 4$, $CR(10, 5) = R(0, 0) = 1$. Hence, $PS_a = (10, 10, 5, 0)$ is uniquely stochastically stable. In this case, the minimum resistance path from $(10, 5)$ to gray states is the one that reaches $(6, 5)$ or $(10, 1)$. The one from $(0, 0)$ to gray states is the one that reaches $(0, 1)$.

2.4 Incomplete information with costly acquisition

In this section, we assume that each player can not freely learn the group of her/his opponent. Each player can buy this information at cost $c > 0$. We refer to this condition as costly information acquisition. It is trivial to notice that Lemma 3 is not valid anymore. Indeed, since players learn the group of their opponent conditional on paying a cost, not every player pays it, and the dynamics are no longer separable.

This time, $Z = \{a, b, ab, ba, aa, bb\}$. It is trivial to show that there are four strictly dominant behaviors; indeed, $U^i_{z_i(t)}(aa, \omega_t) = U^i_{z_i(t)}(a, \omega_t) - c$ and $U^i_{z_i(t)}(bb, \omega_t) = U^i_{z_i(t)}(b, \omega_t) - c$. Hence,

$U^i_{z_i(t)}(aa, \omega_t) < U^i_{z_i(t)}(a, \omega_t)$ and $U^i_{z_i(t)}(bb, \omega_t) < U^i_{z_i(t)}(b, \omega_t)$, $\forall i \in N$ and $\forall \omega_t \in \Omega$. We define strictly dominant behaviors as $Z^o = \{a, b, ab, ba\}$, with z^o_i being a strictly dominant behavior of player i.

Equations (2.5)–(2.8) are the payoffs at period t, for a player $i \in K$ currently playing a or ab.

$$U^i_a(a, \omega_t) = \frac{n^{KK}_t + n^{K'K}_t - 1}{N - 1} \pi^K_a, \tag{2.5}$$

$$U^i_a(b, \omega_t) = \frac{N - n^{KK}_t - n^{K'K}_t}{N - 1} \pi^K_b, \tag{2.6}$$

$$U^i_a(ab, \omega_t) = \frac{n^{KK}_t - 1}{N - 1} \pi^K_a + \frac{N_{K'} - n^{K'K}_t}{N - 1} \pi^K_b - c, \tag{2.7}$$

$$U^i_a(ba, \omega_t) = \frac{N_K - n^{KK}_t}{N - 1} \pi^K_b + \frac{n^{K'K}_t}{N - 1} \pi^K_a - c. \tag{2.8}$$

Note that if $c = 0$, then $aa = a$ and $bb = b$. We begin the analysis with the unperturbed dynamics.

2.4.1 Unperturbed dynamics

So far, there are no more random elements with respect to Section 2.3. Therefore, $\omega_{t+1} = F(\omega_t, \theta_{t+1})$. Nine states can be absorbing under this specification.

Lemma 5. *Under costly information acquisition, there are nine possible absorbing states:* $\omega^R \cup (N_A, N_A, 0, 0)$.

We summarize all the relevant information in Table A1. The reader can note two differences with respect to Section 2.3: firstly, some states are absorbing if and only if some conditions hold, and secondly, there is one more possible absorbing state, that is, $(N_A, N_A, 0, 0)$. Such an absorbing state was also possible in Neary under the same conditions on payoffs and group size.

Where we write "none", we mean that a state is always absorbing for every value of group size, payoffs, and/or the cost. We name $(N_A, N_A, 0, 0)$ the Type Monomorphic State (TS from now on): each group is playing its favorite action in this state, causing miscoordination in inter-group interactions. In both MS and TS, no player is buying the information, while in PS, at least one group is buying the information.

Monomorphic States are absorbing states for every value of group size, payoffs, and cost. Indeed, when each player is playing one action with any other player, players do not need to learn their opponent's group (the information cost does not matter): they best reply to these states by playing the same action.

Polymorphic States are absorbing if and only if the cost is low enough: if the cost is too high, buying the information is too expensive, and players best reply to Polymorphic States by playing a or b. The Type Monomorphic State is absorbing if group B is either sufficiently close in size to group A or strong enough in preferences for its favorite action and if the cost is high enough. The intuition is the following. On the one hand, if the cost is high and if group B is weak in preferences or small enough, every player of group B best replies to TS by playing a. On the other hand, if the cost is low enough, every player best replies to this state by buying the information and conditioning the action.

2.4.2 Perturbed dynamics

We now introduce perturbed dynamics. In this case, we assume that players can make two types of mistakes: they can make a mistake in the information choice and in the coordination choice. Choosing the wrong behavior, in this case, can mean both. We say that with probability η, a player who is given the revision opportunity at period t chooses to buy the information when it is not optimal. With probability ε, s/he makes a mistake in the coordination choice. We could have chosen to set only one probability of making a mistake with a different behavior or strategy.

The logic behind our assumption is to capture behaviorally relevant mistakes. We assume a double punishment mechanism for players choosing by mistake the information level and the coordination action. Specifically, our mistake counting is not influenced by our definition of behaviors. We could have made the same assumption starting from the standard definition of strategies assuming that players can make different mistakes in choosing the two actions that constitute the strategy. Our assumption is in line with works such as Jackson and Watts (2002) and Bhaskar and Vega-Redondo (2004), which assume mistakes in the coordination choice and the link choice.

Formally, $\omega_{t+1} = F(\omega_t, \theta_{t+1}, \psi^c_{t+1})$, where $\psi^c_{t+1} = \{\psi^\varepsilon_{t+1}, \psi^\eta_{t+1}\}$ is the set of players who make a mistake at period t among those who are given the revision opportunity, ψ^ε_{t+1} is the set of players who make a mistake in the coordination choice, and ψ^η_{t+1} the set of players that make a mistake in the information choice.

Since we assume two types of mistakes, the concept of resistance changes. We then need to consider three types of resistances. We call $r_\varepsilon(\omega_t, \dots, \omega_s)$ the path from state ω_t to state ω_s with ε mistakes (players make a mistake in the coordination choice). We call $r_\eta(\omega_t, \dots, \omega_s)$

the path with η mistakes (players make a mistake in the information choice). Finally, we call $r_{\varepsilon\eta}(\omega_t, \ldots, \omega_s)$ the path with mistakes both in the coordination choice and the information choice. Since we do not make further assumptions on ε and η (probability of making mistakes uniformly distributed), we can assume $\eta \propto \varepsilon$.

We count each mistake in the path of both ε and η mistakes as 1; however, $r_{\varepsilon\eta}(\omega_t, \ldots, \omega_s)$ is always double since it implies a double mistake. Indeed, we can see this kind of mistake as the sum of two components, one in η and the other in ε, namely $r_{\varepsilon\eta}(\omega_t, \ldots, \omega_s) = r_{\varepsilon\eta|_\varepsilon}(\omega_t, \ldots, \omega_s) + r_{\varepsilon\eta|_\eta}(\omega_t, \ldots, \omega_s)$.

For example, think about $\omega_t = MS_a$, and that one player from B is given the revision opportunity at period t. Consider the case where s/he makes a mistake both in the information choice and in the coordination choice. For example, s/he learns the group and s/he plays a with A and b with B. This mistake delineates a path from MS_a to the state $(N_A, N_A, N_B, N_B - 1)$ of resistance $r_{\varepsilon\eta}(MS_a, \ldots, (N_A, N_A, N_B, N_B - 1)) = 2$. Next, think about $\omega_t = TS$: the transition from TS to $(N_A, N_A - 1, 0, 0)$ happens with one η mistake. One player from A should make a mistake in the information choice and optimally choosing ab. In this case, $r_\eta(TS, \ldots, (N_A, N_A - 1, 0, 0)) = 1$. With a similar reasoning, $r_\varepsilon(MS_a, \ldots, (N_A - 1, N_A - 1, N_B, N_B)) = 1$: a player of group A makes a mistake in the coordination choice and chooses b.

Before providing the results, we explain why using behaviors instead of strategies does not influence the stochastic stability analysis. Let us consider all the sixteen strategies as presented in Section 2.2, and just one kind of mistake in the choice of the strategy. Let us take two strategies $s', s'' \in z'$ and a third strategy $s''' \in z''$. Now consider the state $\bar{\omega}$, where $s_i = s'$, $\forall i \in N$, and the state ω', where $s_i = s'$, $\forall i \in \{0, \ldots, N - m - 1\}$ and $s_j = s''$, $\forall j \in \{N - m, \ldots, N\}$. Since s' and s'' are both payoff-equivalent and behaviorally equivalent, s' and s'' are the best reply strategies $\forall i \in N$ in both states $\bar{\omega}$ and ω'. Therefore at each period, every player who is given the revision opportunity in state $\bar{\omega}$ or ω' chooses s' and s'' with equal probability. Now let us consider the state $\bar{\omega}'$ where $s_i = s'''$, $\forall i \in N$. When considering the transition between $\bar{\omega}$ and $\bar{\omega}'$, the number of mistakes necessary for this transition is the same whether the path passes through ω' or not because the best reply strategy is the same in both ω' and $\bar{\omega}$. Therefore, when computing the stochastically stable state, we can neglect s'' and ω'.

We divide this part of the analysis into two cases, the first one where the cost is low and the second one when the cost is high.

Low cost

This section discusses the case when c is as low as possible but greater than 0.

Corollary 1. *Under costly information acquisition, if* $0 < c < \frac{1}{N-1}\min\{\pi_A, \pi_B\}$*,* MS *and* PS *are absorbing states, while* TS *is not an absorbing state.*

The proof is straightforward from Table A1. In this case, there are eight candidates to be stochastically stable equilibria.

Theorem 2. *Under costly information acquisition, for large enough* N*, take* $0 < c < \frac{1}{N-1}$ $\min\{\pi_A, \pi_B\}$*. If* $\frac{\pi_B}{\pi_A} < \frac{N_B}{N_A}$*, then* PS_b *is uniquely stochastically stable. If* $\frac{\pi_B}{\pi_A} > \frac{N_B}{N_A}$*, then* PS_a *is uniquely stochastically stable.*

The conditions are the same as in Theorem 1. When the cost is low enough, whenever a player can buy the information, s/he does it. Consequently, the basins of attraction of both Monomorphic States and Polymorphic States have the dimension they had under free information acquisition. Due to these two effects, the results are the same as under free information acquisition. This result is not surprising per se but serves as a robustness check of the results of Section 2.3.2.

High cost

In this part of the analysis, we focus on a case when only MS and TS are absorbing states.

Define the following set of values:

$$\Xi_{PS} = \{N_B\pi_A, N_A\pi_B, (N_B - 1)\,\Pi_B, (N_A - 1)\,\pi_A, N_B\Pi_A, (N_B - 1)\,\pi_B\}\,.$$

Corollary 2. *Under costly information acquisition, if* $c > \frac{1}{N-1}\max\{\Xi_{PS}\}$ *and* $\frac{\pi_B}{\Pi_B} < \frac{N_B-1}{N_A}$*, then only* MS *and* TS *are absorbing states. If* $\frac{\pi_B}{\Pi_B} \geq \frac{N_B-1}{N_A}$*, then only* MS *are absorbing states.*

The proof is straightforward from Table A1, and therefore, we omit it. We previously gave the intuition behind this result. Let us firstly consider the case in which TS is not an absorbing state, hence, the case when $\frac{\pi_B}{\Pi_B} \geq \frac{N_B-1}{N_A}$.

Theorem 3. *Under costly information acquisition, for large enough* N*, take* $\frac{\pi_B}{\Pi_B} \geq \frac{N_B-1}{N_A}$ *and* $c > \frac{1}{N-1}\max\{\Xi_{PS}\}$*. If* $N_A > \frac{2N\pi_A+\Pi_A-\pi_A}{\Pi_A+\pi_A}$*, then* MS_a *is uniquely stochastically stable. If* $N_A < \frac{2N\pi_A+\Pi_A-\pi_A}{\Pi_A+\pi_A}$*, then* MS_b *is uniquely stochastically stable.*

If group A is sufficiently large or strong enough in preferences, the minimum number of mistakes to exit from the basin of attraction of MS_a is higher than the minimum number of mistakes to exit from the one of MS_b. Therefore, MS_a is uniquely stochastically stable: every player plays behavior a in the long-run.

Now we analyze the case when also TS is a strict equilibrium.

Theorem 4. *Under costly information acquisition, for large enough N, take $\frac{\pi_B}{\Pi_B} < \frac{N_B-1}{N_A}$ and $c > \frac{1}{N-1}\max\{\Xi_{PS}\}$.*

- *If $\pi_B - \pi_A$ is high enough, MS_a is uniquely stochastically stable.*
- *If $\pi_A - \pi_B$ is high enough, MS_B is uniquely stochastically stable.*
- *If $\pi_A + \pi_B$ is small enough, TS is uniquely stochastically stable.*

The first condition expresses a situation where group A is stronger in preferences than group B or group A is sufficiently larger than group B. In this case, there is an asymmetry in the two costs for exiting the two basins of attraction of MS_a and MS_b. Exit from the first requires more mistakes than exit from the second. Moreover, reaching MS_a from TS requires fewer mistakes than reaching MS_b from TS. For this reason, $R(MS_a) > CR(MS_a)$ and MS_a is uniquely stochastically stable in this case. A similar reasoning applies to the second condition.

The third condition expresses a case where both groups are strong enough in preferences or have sufficiently similar sizes. Many mistakes are required to exit from TS, compared to how many mistakes are required to reach TS from the two MS. Indeed, TS is the state where both groups are playing their favorite action. Since they are both strong in preferences or large enough, in this case, all the players play their favorite action in the long-run, but they miscoordinate in inter-group interactions.

The results of Theorems 3 and 4 reach the same conclusions as Neary. However, our analysis allows us to affirm that only with a high cost, the MS or the TS is stochastically stable. This result enriches the previous analysis.

As a further contribution, comparing these results with those in Section 2.4.2, we can give the two conditions for inter-group miscoordination to happen in the long-run. First, the cost to pay to learn the opponent's group should be so high that players never learn their opponent's group. Second, both groups should be strong enough in preferences or sufficiently close in size. The following lemma states what happens when the cost takes medium values.

Lemma 6. *If $\frac{1}{N-1}\max\{\pi_A, \pi_B\} < c < \frac{1}{N-1}\min\{\Xi_{PS}\}$, then the stochastically stable states must be in the set $M = \{PS_a, PS_b, MS_a, MS_b\}$.*

When the cost lowers a tiny quantity from the level of Section 2.4.2, TS is not absorbing anymore. Therefore, only PS and MS can be stochastically stable when the cost is in the interval above. However, not all the PS can be stochastically stable, only the two where all the players play their favorite action in inside-group interactions. The intuition of this result is simple: if players condition their action on the groups in the long-run, they play their favorite action with those with similar preferences.

We do not study when MS are stochastically stable or when PS are: we leave this question for future analysis. Nevertheless, given the results of Sections 2.4.2 and 2.4.2, we expect that for higher levels of cost, MS is stochastically stable, and for lower levels, PS is stochastically stable.

2.5 Discussion

The results of our model involve three fields of the literature. Firstly, we contribute to the literature on social conventions. Secondly, we contribute to the literature on stochastic stability analysis, and lastly, we contribute to the literature on costly information acquisition.

For what concerns social conventions, many works in this field study the existence in the long-run of heterogeneous strategy profiles. We started from the original model of Neary (2012), which considers players heterogeneous in preferences, but with a smaller strategic set than ours (Heterogeneity has been discussed in previous works such as Smith and Price, 1973, Friedman, 1998, Cressman et al., 2001, Cressman et al., 2003 or Quilter et al., 2007). Neary's model gives conditions for the stochastic stability of a heterogeneous strategy profile that causes miscoordination in inter-group interactions in a random matching case. Neary and Newton (2017) expands the previous idea to investigate the role of different classes of graphs on the long-run result. It finds conditions on graphs such that a heterogeneous strategy profile is stochastically stable. It also considers the choice of a social planner that wants to induce heterogeneous or homogeneous behavior in a population.

Carvalho (2017) considers a similar model, where players choose their actions from a set of culturally constrained possibilities and the heterogeneous strategy profile is labeled as miscoordination. It finds that cultural constraints drive miscoordination in the long-run. Michaeli and Spiro (2017) studies a game between players with heterogeneous preferences and who feel pressure from behaving differently. Such a study characterizes the circumstances under which a biased norm can prevail on a non-biased norm. Tanaka et al. (2018) studies how local dialects survive in a society with an official language. Naidu et al. (2017) studies the evolution of egalitarian and inegalitarian conventions in a framework with asymmetry similar to the language game. Likewise, Belloc and Bowles (2013) examines the evolution and the persistence of inferior cultural conventions.

We introduce the assumption that players can condition the action on the group if they pay a cost. This assumption helps to understand the conditions for the stability of the Type Monomorphic State, where players miscoordinate in inter-group interactions. We show that a low cost favors inter-group coordination: incomplete information, high cost, strength in preferences, and group size are key drivers for inter-group miscoordination. Like many works in this literature, we show the importance of strength in preferences and group size in the equilibrium selection.

Concerning network formation literature, Goyal et al. (2021) conducts an experiment on the language game, testing whether players segregate or conform to the majority. van Gerwen and Buskens (2018) suggests a variant of the language game similar to our version but in a model with networks to study the influence of partner-specific behavior on coordination. Concerning auctions theory, He (2019) studies a framework where each individual of a population divided into two types has to choose between two skills: a "majority" and a "minority" one. It finds that minorities are advantaged in competition contexts rather than in coordination ones. He and Wu (2020) tests the role of compromise in the battle of sexes with an experiment.

Like these works, we show that group size and strength in preferences matter for the long-run equilibrium selection. The states where the action preferred by the minority is played in most of the interactions (MS_b or PS_b) are stochastically stable provided that the minority is strong enough in preferences or sufficiently large.

A parallel field is the one of bilingual games such as the one proposed by Goyal and Janssen (1997) or Galesloot and Goyal (1997): these models consider situations in which players are homogeneous in preferences towards two coordination outcomes, but they can coordinate on a third action at a given cost.

Concerning the technical literature on stochastic stability, we contribute by applying standard stochastic stability techniques to an atypical context, such as costly information acquisition. Specifically, we show that with low cost levels, Polymorphic States where all players in one group condition their action on the group are stochastically stable. Interestingly, only one group of players needs to learn their opponent's group. With high cost levels, Monomorphic States where no player conditions her/his action on the group are stochastically stable. Since the seminal works by Bergin and Lipman (1996) and Blume (2003), many studies have focused on testing the role of different mistake models in the equilibrium selection. We use uniform mistakes, and introducing different models could be an interesting exercise for future studies.

Among the many models that can be used, there are three relevant variants: payoff/cost-dependent mistakes (Sandholm, 2010; Dokumacı and Sandholm, 2011; Klaus and Newton, 2016; Blume, 1993; Myatt and Wallace, 2003a), intentional mistakes (Naidu et al., 2010; Hwang et al., 2016), and condition-dependent mistakes (Bilancini and Boncinelli, 2020). Important experimental works in this literature have been done by Lim and Neary (2016), Hwang et al. (2018), Mäs and Nax (2016), and Bilancini et al. (2021).

Other works contribute to the literature on stochastic stability from the theoretical perspective (see Newton, 2018 for an exhaustive review of the field). Recently, Newton (2021) has expanded the domain of behavioral rules regarding the results of stochastic stability. Sawa and Wu (2018a) shows that with loss aversion individuals, the stochastic stability of Risk-Dominant equilibria is no longer guaranteed. Sawa and Wu (2018b) introduces reference-dependent preferences and analyzes the stochastic stability of best response dynamics. Staudigl (2012) examines stochastic stability in an asymmetric binary choice coordination game.

For what concerns the literature on costly information acquisition, many works interpret the information's cost as costly effort (see the seminal contributions by Simon, 1955 or Grossman and Stiglitz, 1980). Our paper is one of those. Many studies place this framework in a sender-receiver game. This is the case of Dewatripont and Tirole (2005), which builds a model of costly communication in a sender-receiver setup.

More recent contributions in this literature are Dewatripont (2006), Caillaud and Tirole (2007), Tirole (2009), and Butler et al. (2013). Bilancini and Boncinelli (2021) applies this model to persuasion games with labeling. Both Bilancini and Boncinelli (2018a) and Bilancini and Boncinelli (2018c) consider coarse thinker receivers, combining costly information acquisition with the theory of Jehiel (2005). Rational inattention is a recent field where the information cost is endogenous (see Maćkowiak et al., 2023 for an exhaustive review). We assume that the cost is exogenous and homogeneous for each player.

Güth and Kliemt (1994) firstly uses costly information acquisition in evolutionary game theory in a game of trust. It finds conditions such that developing a conscience can be evolutionarily stable. More recently, Berger and De Silva (2021) uses a similar concept in a deterrence game where agents can buy costly information on past behaviors of their opponents.

Many works use similar concepts of cost in the evolutionary game theory literature on coordination games. For example, Staudigl and Weidenholzer (2014) considers a model where players can pay a cost to form links. The main finding is that when agents are constraint in the possible number of interactions, the payoff-dominant convention emerges in the long-run.

The work by Bilancini and Boncinelli (2018d) extends Staudigl and Weidenholzer (2014). The model introduces the fact that interacting with a different group might be costly for a player. It finds that when this cost is low, the Payoff-Dominant strategy is the stochastically stable one. When the cost is high, the two groups in the population coordinate on two different strategies: one on the risk-dominant and the other on the payoff-dominant. Similarly, Bilancini et al. (2018) studies the role of cultural intolerance and assortativity in a coordination context. In that model, there is a population divided into two cultural groups, and each group sustains a cost from interacting with the other group. It finds interesting conditions under which cooperation can emerge even with cultural intolerance.

2.6 Conclusions

We can summarize our results as follows. When players learn the group of their opponent at a low cost, they always coordinate: they play their favorite action with their similar, while in inter-group interactions, they play the favorite action of the group that is stronger in preferences or with large enough size. If the cost is high, players never learn the group of their opponent. All the players play the same action with every player, or they play their favorite action.

By comparing Sections 2.4.2 and 2.4.2, we can see the impact of varying the cost levels on the long-run results. Surely a low cost favors inter-group coordination. However, a change in the cost level produces two effects that perhaps need further investigation. The first effect concerns the change in the payoff from the interactions between players. The second concerns the change in the purchase of the information.

Consider a starting situation where the cost is low. Players always coordinate on their favorite action in inside-group interactions. If the cost increases, players stop learning their opponent's group (hence, they stop paying the cost), and they begin to play the same action as any other player. If this happens, either Monomorphic States are established in the long-run, or the Type Monomorphic State emerges. In the first case, a group of players coordinates on its second best option, even in inside-group interactions. For this group, there could be a certain loss in terms of welfare. In the second case, players miscoordinate in inter-group interactions, and hence, all of them could have a certain loss in welfare.

Nevertheless, when the cost is low, there is a "free-riding" behavior that vanishes if the cost increases. In fact, with low cost levels, only one group needs to pay the cost, and the other never needs to pay it. In one case, players of group A play their favorite action both in inside-group and inter-group interactions; hence, they never need to pay the cost, while group B always needs to afford it. In the other case, the opposite happens. Therefore, when the cost increases, one of the two groups will benefit from not paying for the information anymore. Future studies could address the implications of this trade-off between successful coordination and the possibility of not paying the cost.

We conclude with a short comparison of our result with the one of Neary (2012). Indeed, it is worthwhile to mention a contrast that is a consequence of the possibility of conditioning the action on the group of the player. In the model of Neary, a change in the strength of preferences or the group size of one group does not affect the behavior of the other group. We can find this effect even in our model when the cost is high. For example, when MS_a is stochastically stable and group B becomes strong enough in preferences or sufficiently large, the new stochastically stable state becomes TS. Therefore, group A does not change its behavior. However, when the cost is sufficiently low, the change in payoffs or group size of one group influences the other group's behavior in inter-group interactions. For instance, when PS_a is stochastically stable, if group B becomes strong enough in preferences or sufficiently large, PS_b becomes stochastically stable. Hence, both groups change the way they behave in inter-group interactions.

Nevertheless, we can interpret similarly the passing from MS_a to TS and the one from PS_a to PS_b. In both cases, both groups keep playing their favorite action in inside-group interactions, and what happens in inter-group interactions depends on strength in preferences and group size. Therefore, in this respect, the behavioral interpretation of our results is similar to Neary's.

Chapter 3

The rise of conformism in the Hawk-Dove game

Abstract

I consider a double selection mechanism to study the evolution of conformists (the less intelligent type) and myopic best repliers (the more intelligent type) in the Hawk-Dove game. Firstly, I derive the stochastically stable distribution of strategies for each population level. Secondly, I determine the fitness of each type in the stochastically stable distribution of strategies for each population level. Myopic best repliers pay a cognitive cost for being the more intelligent type. The process goes on until both systems reach stability. My results highlight three main findings. First, conformists outnumber myopic best repliers in the stable state. Second, the fraction of hawks is always higher (lower) than the one predicted by standard models, depending on the harshness of the conflict. Third, conformists play hawk when the conflict is harsh and dove when it is mild.

Keywords: hawk-dove, learning, evolutionary stability, stochastic stability, behavioral rules.

JEL Classification Codes: C73, D74, D83

Notes. This chapter is a single author work.

3.1 Introduction

"Alla riscossa stupidi che i fiumi sono in piena
Potete stare a galla"

To the rescue stupids that rivers are in flood
You can stay afloat.

– F. Battiato (1980), "Up Patriots to Arms"

Intelligence may seem like a prerequisite for survival in strategic interactions. Specifically, it sounds reasonable to think that conflict games should be particularly selective against players who do not play intelligently. However, there is often a trade-off between being intelligent and not. Firstly, being intelligent causes cognitive costs. Secondly, intelligent players may be forced to adapt the way they play to unintelligent ones in a way that favors the latter.

In this paper, I study the evolution of two types of agents in the Hawk-Dove game. I am interested in mainly three research questions. Firstly, does the less intelligent type survive in such a game? Secondly, is the stable distribution of strategies the canonical mixed Nash Equilibrium? Lastly, which strategy does the less intelligent type play in the stable state?

I consider a population of agents repeatedly randomly matched to play a Hawk-Dove game in discrete time. Each agent is either one of two types: conformists or myopic best repliers, the former being less intelligent than the latter since they do not give importance to payoff while making decisions. Specifically, myopic best repliers always calculate the best response to the average payoff at the current period, while conformists choose the strategy played by the majority in the current period. I study the evolution of strategies with stochastic stability for all possible population levels. Then, I compute the evolution of types, calculating the fitness of each type at each stochastically stable distribution of strategies. The stable state is such that each type has the same fitness, and the distribution of strategies is stochastically stable for that population level.[1] Each agent chooses the best reply to the current period following their own behavioral rule (Kandori et al., 1993; Young, 1993a). Types evolve following a replicator equation (Hamilton, 1967; Smith, 1982). Myopic best repliers play more intelligently than conformists; therefore, their fitness is discounted by a cognitive cost. This mechanism generates a trade-off between being more or less intelligent.

With my model, I can answer the three research questions proposed above. First, not only do conformists survive, they outnumber myopic best repliers in the stable state. This result

[1] I am implicitly assuming that the stochastic stability process evolves at a faster rate than the evolution of types, see Section 3.2.4.

is due to the behavior of myopic best repliers that keep conformists alive. Indeed, as long as conformists do not outnumber myopic best repliers, the latter compensates for what the former plays and pushes the system towards a distribution of strategies where each type earns the same payoff. In this distribution of strategies, superior intelligence becomes unnecessary since more and less intelligent types achieve the same fitness. Second, the stochastically stable distribution of strategies in the stable state is not the one predicted by canonical models. It is a type monomorphic distribution of strategies where each type plays a different strategy (which one is selected depends on the harshness of conflict as I should explain below). In this distribution of strategies, the fraction of agents playing hawk is different from the one predicted by standard Hawk-Dove games. Specifically, it is always greater than the one predicted by canonical model when the conflict is mild and it is always lower when the conflict is harsh. Third, in the stable state, the most intelligent type plays hawk when the conflict is mild and dove when it is harsh.

The intuition behind my results is as follows. Considering the evolution of strategies, when the fraction of conformists is a minority, the unique absorbing distribution of strategies is the mixed one such that part of the agents plays hawk, and part play dove. Consequently, also the stochastically stable distribution of strategies is the mixed one. Note that, in this distribution of strategies, all conformists play hawk or dove depending on the harshness of conflict, while myopic best repliers play a mixed strategy. When conformists are the majority, type monomorphic distributions of strategies become absorbing. In such distributions, conformists play hawk (dove), and myopic best repliers play dove (hawk). However, such distributions of strategies are never stochastically stable as long as the level of conformists is lower than a critical threshold. When conformists are the vast majority, the mixed distributions of strategies are no longer absorbing since the fraction of conformists is so high that myopic best repliers can only play a pure strategy as a response to conformists behavior. Therefore, only the type monomorphic ones can be stochastically stable. In these situations, the harshness of conflict guides the results. When conflict is mild, conformists play dove, and myopic best repliers play hawk. When conflict is harsh, conformists play hawk, and myopic best repliers play dove.

In the mixed distribution of strategies, conformists and myopic best repliers earn the same payoff. However, myopic best repliers pay the cognitive cost for being the most intelligent type. Hence, for all population levels such that the mixed distribution of strategies is stochastically stable, the fraction of conformists keeps rising. This effect continues to the point where conformists become the vast majority in the population, and type monomorphic distributions of strategies become stochastically stable. In these distributions of strategies, myopic best repliers earn a higher payoff on average, but they still pay the cognitive cost for being the most intelligent type. Therefore, conformists still grow up to the point where the benefits of being myopic best repliers meet the costs of being such types.

I present some welfare considerations about the stable state in the ultra long-run. When compared to the mixed distribution of strategies, the type monomorphic distribution of strategies is desirable when conflict is mild, but not when conflict is harsh.

In addition to these results, I validate the robustness of my findings in different setups. I verify the robustness of my findings to the relaxation of some crucial assumptions of my model. I find that the first finding (conformists being the majority in the long-run) is the most robust among the three, while the second and the third depend on the cost for being intelligent, but they are still keep a level of robustness. Further, I consider a game where agents play at each step a mild conflict Hawk-Dove with frequency p and a harsh conflict Hawk-Dove with frequency $1-p$. In this setup, agents do not know which game they are playing, but they play the best reply to the linear combination of the two games. The linear combination between two Hawk-Doves games is itself an Hawk-Dove game. Hence, not surprisingly, my results are robust to this setup. Afterwards, I consider a setup where at each time a mild conflict Hawk-Dove is realized with probability p and a harsh conflict Hawk-Dove is realized with probability $1-p$. In this setup, agents are doubly myopic since they assume that the game in a given time will be the same in the following time. The results from this extension show that uncertainty (between steps) decreases the chances of survival of the least intelligent behavioral rule. However, a small level for the cost for being intelligent is sufficient to make conformists prosper. Lastly, I extend the results in my model to the Hawk-Dove-Bourgeois model (Smith, 1982). I find that in the ultra long-run only conformists survive, and they all play the bourgeois strategy (in line with canonical models). With this extension, I show that conformists have a further advantage if an element of coordination is introduced in the game, and therefore, myopic best repliers become extinct.

My paper is mostly related to three works in the literature: Mohlin (2012), Arigapudi et al. (2021), and Bilancini et al. (2022).

Mohlin (2012) considers a model to study the evolution of theories of mind (namely, the level-k model and the hierarchy model) in different classes of games. He firstly shows the mechanism such that more sophisticated theories of mind can coexist with less sophisticated ones even in the Hawk-Dove (through the mixed Nash Equilibrium). With respect to this paper, I study the robustness of the analysis to the introduction of different levels of intelligence and a cognitive cost for being more intelligent. In fact, in Mohlin (2012), the stable fraction of the population is a set of states where the stable distribution of strategies is the mixed one. In my model, the stable fraction of the population is a precise value (such that conformists are always the majority), and the stable distribution of strategies is type monomorphic. Importantly, in Mohlin (2012), in the stable distribution of strategies, all agents (not acting randomly) play hawk in the same fraction predicted by the mixed Nash Equilibrium. In my model, not only the distribution of strategies is type monomorphic, but the fraction of agents playing hawk is never the one predicted by the mixed Nash Equilibrium.

Arigapudi et al. (2021) proposes a Hawk-Dove game with two populations where agents have heterogeneous sampling dynamics. They show that if agents have a limited sampling ability, the two-populations Hawk-Dove game can converge to a mixed distribution of strategies (instead of a type monomorphic one). I show that a type monomorphic distribution of strategies (instead of a mixed one) arises in the Hawk-Dove with one population due to heterogeneity in the intelligence of agents. In contrast with the previous work, such heterogeneity arises endogenously in my model. In a few words, my result can be seen as a complement to Arigapudi et al. (2021). They show non-standard results for a two-populations Hawk-Dove due to heterogeneity in agents' abilities; I show non-standard results for a one-population Hawk-Dove due to heterogeneity in agents' intelligence.

Bilancini et al. (2022) presents a Hawk-Dove game with two populations where one has a longer memory than the other. They show that when conflict is harsh, the population with the longest memory plays hawk, while they play dove when it is mild. Therefore, in their model, the group with the better cognitive ability plays hawk when the conflict is harsh and dove when it is mild. My model shows that the opposite happens in a one-population game. Although the results are different and hardly comparable, in both models, greater cognitive ability (longer memory in their model, and more intelligence in mine) pushes towards dove when conflict is mild and hawk when conflict is harsh. Indeed, in my model, conformists play dove (hawk) when conflict is mild (harsh) because they occasionally play intelligently. However, given that the stable composition of the population is such that conformists are the majority in the ultra long-run, myopic best repliers react to conformists' behavior. This mechanism generates the opposite results than Bilancini et al. (2022).

The remaining of the paper is organized as follows. In Section 3.1.1, I review the existing literature. In Section 3.2, I present the model. In Section 3.3, I provide the results: in 3.3.1, I present the results of the stochastic stability analysis, in 3.3.2, the analysis on the evolution of types, and in 3.3.3, the evolutionary stable state. In Section 3.3.4, I show some welfare considerations about the stable state. In Section 3.3.5, I study some robustness tests by relaxing some crucial assumptions to my model. In Section 3.4, I provide some extensions: in 3.4.1 and 3.4.2, I extend my result to uncertain environments (uncertainty over games within and between steps), and in 3.4.3, I extend the results to the Hawk-Dove-Bourgeois game. I offer some concluding remarks in Section 3.5. All the proof are provided in Appendix B.1 and additional details on the simulations are given in Appendix B.2.

3.1.1 Literature review

The importance of the Hawk-Dove game in studying conflict situations has been documented in various disciplines such as biology (Hamilton, 1964; Smith and Price, 1973), economics (Lipnowski and Maital, 1983; Kimbrough et al., 2020), political science (Brams and Kilgour, 1987, 2001; Dixit et al., 2019), and in other interdisciplinary fields (Huntingford and Turner, 1987; Archer, 1988; Baliga and Sjöström, 2012, 2020). The experimental literature has produced some interesting contributions to the understanding of the Hawk-Dove game. Oprea et al. (2011) conducted an experiment testing the precision of the convergence to both one and two populations games, finding positive answers.[2] Benndorf et al. (2016) and Benndorf and Martinez-Martinez (2017) studied both theoretically and in the lab an hybrid version of the Hawk-Dove game where agents play a combination between the one and two populations game. Such an hybrid version could explain empirical puzzles in the laboratory. Recent developments in the theoretical literature on Hawk-Dove games have focused on understanding when the standard predictions could fail. Standard models predict a mixed Nash Equilibrium in the one-population game and a type monomorphic state in the two-populations game (Friedman, 1991; Weibull, 1997). Novel contributions in the two-populations game have shown potential mechanisms in driving the selection towards one type monomorphic state or the other (Herold and Kuzmics, 2020; Bilancini et al., 2022), or in driving the selection towards a mixed Nash Equilibrium (Arigapudi et al., 2021). Foley et al. (2018, 2021) are noteworthy works too in the social networks literature. They study the emergence of conventions as a way to solve conflicts. They find that competitive ability together with the partner choice guide the convergence to possibly cooperative equilibria.

Concerning this literature, I show an endogenous mechanism that creates the conditions for the coexistence between more and less intelligent types in the Hawk-Dove game. This mechanism also leads the system to a type monomorphic distribution of strategies, where the fraction of agents that plays hawk is never the one predicted by standard evolutionary models. I also show the role of the harshness of conflict in driving the results, showing that the less intelligent type plays hawk when conflict is harsh and dove when conflict is mild.

[2] See Van Huyck et al. (1995) and Friedman (1996) for notable previous attempt to reproduce such convergence in the lab.

Regarding the behavioral rules literature in evolutionary game theory (Alós-Ferrer and Buckenmaier, 2020), many relevant works focus on the level-k model of intelligence. A consistent number of models have tried to generalize Young (1993b) bargaining model, introducing the level-k model (Sáez-Martı and Weibull, 1999; Matros, 2003; Khan and Peeters, 2014). Among the most important contributions, Nax and Newton (2022) proves that in coordination games, when generalized risk dominance exists, the long-run result is indifferent to the introduction of any kind of level-k agents.[3] Mohlin (2012) studies the evolution of theories of mind in different classes of games, showing that the coexistence between different level-ks is possible in different kinds of games. Heller (2015) shows that the coexistence between agents with different foresight abilities is possible in a prisoner's dilemma with an uncertain horizon. LiCalzi and Mühlenbernd (2019) studies different learning rules in a model where agents categorize games differently. They show the potential benefit of the imitation rule regarding cooperation levels. Related to the industrial organization literature but still important for the behavioral rules one, Alós-Ferrer and Ritschel (2021) studies the coexistence between imitators and myopic best repliers in Cournot oligopolies (see also Goerg and Selten, 2009 or Schipper, 2009).

Concerning this literature, I show how conformism can survive in the Hawk-Dove game (a game where conforming is often the opposite of playing the best reply strategy). By introducing a cognitive cost, I show the robustness of the states that predicts, in the ultra long-run, a state where conformists outnumber myopic best repliers. Such an introduction allows me to study the distribution of strategies in the stable state, with implications on welfare, and in understanding the behavior of less intelligent types.

More generally, my study relates to other fields in the literature, studying the evolution of learning models in different classes of games. This approach has been particularly prosperous in the heuristic switching models literature (Brock and Hommes, 2001, 1998; Chiarella and He, 2002, 2003; Anufriev and Hommes, 2012a,b). Such literature studies traders that adapt their heuristics based on the relative performance of these heuristics. These works have successfully explained many stylized facts in the financial markets. Relevant theoretical works in this literature also include Dindo and Tuinstra (2011); Tuinstra et al. (2014); Schmitt et al. (2017). All of the above cited works share one common characteristic with my model: suggesting an evolutionary argument to explain the coexistence between more and less intelligent agents. Specifically, I show how a less intelligent behavioral rule can coexist with a more intelligent one in the Hawk-Dove game. Also related to my model is the literature on

[3] Such a study generalizes some previous results by Blume (2003), Myatt and Wallace (2003b), and Peski (2010).

the "Learning to Forecast" experiments (Hommes et al., 2005, 2007; Heemeijer et al., 2009). The focus of this literature is the different forecasting heuristics that subjects use. With these papers, I share the intent to explain which choice different subjects adopt. Specifically, I study the role of the harshness of conflict in driving the result: the less intelligent type plays aggressively in the stable state if the conflict is harsh and defensively if the conflict is mild.

Lastly, it is worth mentioning two possibly related fields of the literature: the one on animal behavior and the one on "zero-intelligence" players. In the animal behavior literature, one old question concerns whether, for animals, it is better to be scroungers or producers, namely, to follow other animals' behaviors or to search on their own for their source of food. Many works have been conducted in this literature, and their scope is far beyond the one of my paper (Laland, 2004; Rendell et al., 2010; Dridi and Lehmann, 2015). However, my work and these works share some relevant features. First, in both cases, one type of behavior implies a cost (in my case, a cognitive cost, in the other case, a cost from try and error). Secondly, the other type of behavior implies social learning. The main difference is that I consider conformists that follow one strategy (playing hawk or dove), while in the other literature, social learning often implies imitative behavior, which is a less extreme form of conformism. In the "zero-intelligence" literature, it is shown that even in the presence of traders with zero intelligence of the market, the equilibrium can be the one predicted by models predicated on utility maximization (Gode and Sunder, 1993; Ladley, 2012). With this literature, I share the result that players with zero intelligence of the game are kept alive by players with full intelligence. However, my model shows that this coexistence brings unexpected results compared to canonical models.

3.2 Model

I consider a continuum of agents, each agent is labelled $i \in [0, 1]$. Each agent is one of two types $\theta_i \in \{C, M\}$. As I will define better in Section 3.2.1, C types are conformists and M types are myopic best repliers. I name α the fraction of C types, and $1 - \alpha$ the fraction of M types.

I consider a dynamical process divided in two steps. At the largest time scale there are generation of types. Each generation is labelled t. The evolution of types between generations is moved by the fitnesses of each type at the equilibrium play in each generation. Within each generation I model how this equilibrium comes about using stochastic stability. The system reaches a stable state whenever the population is stable, and agents reaches an equilibrium play given the stable population. In the next section I describe the evolution of strategies within each generation t.

3.2.1 Evolution of strategies within a generation

In this part of the process, the fraction of conformists α is fixed. Time is discrete, indexed by $\tau = 0, 1, 2...$, at each step, agents are repeatedly randomly matched to play a Hawk-Dove game as represented in Table 3.1.

	H	D
H	$\frac{v-c}{2}, \frac{v-c}{2}$	$v, 0$
D	$0, v$	$\frac{v}{2}, \frac{v}{2}$

Table 3.1: The Hawk-Dove game

Trivially, each agent has two strategies: H and D, i.e. $s_i = \{H, D\}$. To make it a Hawk-Dove game, I assume $v < c$. In the standard two players version of this model, there exists three Nash Equilibria. One where one player plays H, and the other plays D, the opposite, and a mixed Nash Equilibrium where both players play H with probability $\frac{v}{c}$.

I call $n_H(\tau, t)$ the fraction of agents playing H at time τ within generation t; hence, $1-n_H(\tau, t)$ is the fraction of agents playing D at time τ within generation t. For simplicity, I will refer to $n_H(\tau, t)$ as $n_H(\tau)$. $n_H^C(\tau)$ ($n_H^M(\tau)$) is the fraction of type C (M) agents playing H at time τ. Trivially, $n_H(\tau) = n_H^C(\tau) + n_H^M(\tau)$. I define $\boldsymbol{\sigma(\tau, t)} \in [0, 1]^2$ as the distribution of strategies at time τ within generation t. For simplicity, I will refer to $\boldsymbol{\sigma(\tau, t)}$ as $\boldsymbol{\sigma(\tau)}$. Note that $\boldsymbol{\sigma(\tau)} = \big(n_H^C(\tau), n_H^M(\tau)\big)$. I define $\pi(s, s')$ as the payoff for playing strategy s against an opponent playing strategy s', and $\bar{\pi}(s, n_H(\tau)) = n_H(\tau)\pi(s, H) + (1 - n_H(\tau))\pi(s, D)$ as the average payoff for playing strategy s at time τ.

Each agent earns utility $U_i(s, n_H(\tau))$ by playing strategy s at time τ, given the distribution of strategies at time τ. The parameter λ_i measures the relative weight that agent i gives to the average payoff in the decision making process. Trivially, $1 - \lambda_i$ is the importance given to the fraction of agents playing hawk (dove) in the decision making process, without taking into account payoffs. I take λ_i as a proxy fro intelligence, as players with a higher λ_i give more importance to a relevant dimension such as the payoffs.

$$U_i(s, n_H(\tau)) = \lambda_i \left[\bar{\pi}\left(s, n_H(\tau)\right)\right] + (1-\lambda_i)\mathbb{1}(s, n_H(\tau)), \tag{3.1}$$

where $\mathbb{1}(s, n_H(\tau)) = \begin{cases} n_H(\tau) & \text{if } s = H \\ 1 - n_H(\tau) & \text{if } s = D \end{cases}$.

In line with the stochastic stability literature, I assume *inertia*, i.e. only a small fraction of agents has the opportunity to revise their strategy at time τ. I name $s_i(\tau+1)$ the strategy choice of an agent i, who is given the revision opportunity at time τ, such an agent choose $s_i(\tau+1)$ as shown below

$$s_i(\tau+1) \in \arg\max_{s \in S} U_i(s, n_H(\tau)), \forall i. \tag{3.2}$$

I only consider pure strategies unless agents are indifferent between the two strategies; in that case, they randomize.[4] From (3.2), it emerges that all agent *myopically best reply* to the distribution of strategies at time τ; however, for each $\theta_i = C$, $\lambda_i = \lambda_C \approx 0$, and for each $\theta_i = M$, $\lambda_i = \lambda_M = 1$.[5] Therefore, only M types play intelligently since, to make their decision, they give weight 1 to the average payoff at time τ. C types take their decision giving a weight close to 0 to the payoffs: they make it giving weight close to 1 to the fraction of agents playing strategy H. Indeed, C types give a weight close to 1 to $\mathbb{1}(s, n_H(\tau))$. Therefore, they play H, approximately, if and only if $n_H(\tau) > \frac{1}{2}$. With this formulation I give a simple microfundation to the conformists' behavior of C types, results are likely to be not sensible to other formulation of conformists' behavior, as long as they have a grain of intelligence. This definition of conformists and myopic best repliers also allows me to call the former as unintelligent types and the latter as intelligent.

The choice of behavioral rules is not random. Firstly, conformism and myopic best reply are among the most frequently used in evolutionary models (Newton, 2018). Secondly, the contrast between conformism and myopic best reply can be a proxy of the contrast between social and asocial learning rules, which have particular importance in the animal behavior literature (Rendell et al., 2010). Indeed, conformism is an extreme version of social learning since a conformist copies (i.e. learns) whatever the majority plays. Moreover,

[4] A mixed distribution of strategies can be achieved by agents of the same type playing different pure strategies since I am considering a population game.

[5] In particular λ_C should be low enough such that C types choose H if and only if $n_H(\tau) > \frac{1}{2}$. $\lambda_C \approx 0$ serves a selection mechanism for certain levels of the population. Results in Section 3.3 would otherwise have a coexistence between two distributions of strategies in the stable state. Importantly, this assumption only breaks a tie where α is high enough, but does not modify the other results.

myopic best reply is an extreme form of asocial learning since a myopic best replier computes (i.e. experiences/tries) all the payoffs on their own.

I assume that M types pay a higher cognitive cost than C types: I normalize to zero the cognitive cost of C types. As a matter of fact, both types calculate the distribution of strategies at each step. However, M types give importance close to 1 to the average payoff. On the other hand, C types give importance close to 0 to the average strategy with a frequency close to 1 and to the average payoff with a frequency close to 0. Given that myopic best repliers give importance to a relevant dimension, they are for sure more intelligent than conformists. Considering that higher level of intelligence also correlates with higher cognitive loads/costs, I consider myopic best repliers to pay a higher cognitive sot than conformists. Formally, I assume a cost $\kappa_i(\lambda_i) = \kappa\lambda_i$, with $\kappa > 0$ for each agent. This means that $\kappa_i = \kappa$, for all M types and $\kappa_i \approx 0$ for all C types. I assume that only M types pay a cost $\kappa > 0$.[6]

Within each generation of agents, they converge to a stochastically stable distribution of strategies (see Section 3.3.1 for technical details) that I call $\boldsymbol{\sigma}^*(t)$.[7] Such a distribution of strategies is the distribution of strategies where the generation t stays the most. Therefore, types evolving between generations based on their fitness in $\boldsymbol{\sigma}^*(t)$. That is, the fraction of a type in generation $t+1$ positively depends on their fitness in the stochastically stable distribution of strategies of generation t. In the next section I describe the process behind the evolution of types.

3.2.2 Evolution of types between generations

This process represents the evolution of types between generations, each generation is labelled $t = 0, 1, 2, ...$

I name $\alpha(t)$ the fraction of type C in generation t. I evaluate the fitness of each type in generation t at $\boldsymbol{\sigma^*(t)}$, using a replicator equation. To define such an equation, I need to define the average payoff for each type in $\boldsymbol{\sigma^*(t)}$. Firstly, I call s^*_C the strategy played by all C types in $\boldsymbol{\sigma^*(t)}$, and s^*_M the strategy played by all M types in $\boldsymbol{\sigma^*(t)}$. Secondly, I name $\overline{\Pi}_C(\boldsymbol{\sigma^*(t)})$ the payoff earned by C types in $\boldsymbol{\sigma^*(t)}$, and $\overline{\Pi}_M(\boldsymbol{\sigma^*(t)})$ the payoff earned by M types in $\boldsymbol{\sigma^*(t)}$.

[6] It is like agents maximize a given problem that involves deciding how much to be intelligent and which strategy to play given the level of intelligence they choose. Results are likely to be robust in a model where agents' strategies consist of choosing their behavioral rule and the strategy implied by such a rule (see Mohlin, 2012 for a similar model with level-k).

[7] Note that the stochastically stable distribution of strategies is labelled with t, since it denotes the convergence result of a generation t.

I compute the average payoffs in $\boldsymbol{\sigma^*(t)}$ consequently:

$$\overline{\Pi}_C(\boldsymbol{\sigma^*(t)}) = \bar{\pi}(s_C^*, \boldsymbol{\sigma^*(t)}),$$

$$\overline{\Pi}_M(\boldsymbol{\sigma^*(t)}) \;=\; \bar{\pi}(s_M^*, \boldsymbol{\sigma^*(t)}) - \kappa.$$

Lastly, I define $\overline{\Pi}(\boldsymbol{\sigma^*(t)}) = \alpha\overline{\Pi}_C(\boldsymbol{\sigma^*(t)}) + (1-\alpha)\overline{\Pi}_M(\boldsymbol{\sigma^*(t)})$ as the average payoff of the population in $\boldsymbol{\sigma^*(t)}$.

Following standard evolutionary models (Hamilton, 1967; Smith and Price, 1973; Smith, 1982), I consider the following replicator equation:

$$\Delta\alpha(t+1) = \alpha(t)\left[\overline{\Pi}_C(\boldsymbol{\sigma^*(t)}) - \overline{\Pi}(\boldsymbol{\sigma^*(t)})\right],$$

where $\Delta\alpha(t+1) = \alpha(t+1) - \alpha(t)$.

Since there are only two types in the population, the above equation can be specified as

$$\Delta\alpha(t+1) = \alpha(t)(1-\alpha(t))\left[\overline{\Pi}_C(\boldsymbol{\sigma^*(t)}) - \overline{\Pi}_M(\boldsymbol{\sigma^*(t)})\right]. \tag{3.3}$$

To recap the two processes implied by my model, I illustrate them in Figure 3.1.

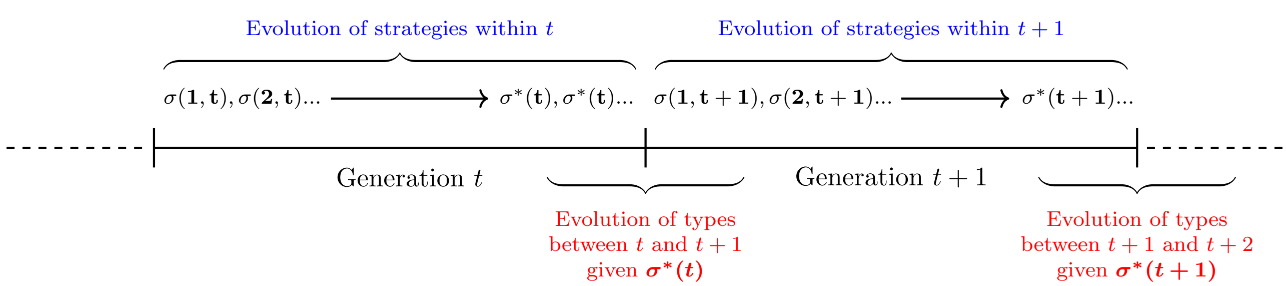

Figure 3.1: A graphical illustration of the evolution of strategies and types

3.2.3 Evolutionary stable state.

To conduct my analysis I am interested in knowing the evolutionary stable composition of the population and the stochastically stable distribution of strategies for this composition. Therefore, I call $\boldsymbol{\omega(t,\tau)} = (\alpha(t), \boldsymbol{\sigma(\tau,t)})$ the state of the world at time τ in generation t. I am interested in $\overline{\boldsymbol{\omega}} = (\overline{\alpha}, \overline{\boldsymbol{\sigma}})$: the evolutionary stable state such that $\overline{\alpha}$, i.e. the composition of the population is stable and $\overline{\boldsymbol{\sigma}}$, i.e. the distribution of strategies is stochastically stable given $\overline{\alpha}$.

3.2.4 Technical assumptions behind the process

The dynamical process explained in the previous sections is reasonable under some assumptions. To track them down, I define formally the process implied by Equation (3.2).

$$\Delta\boldsymbol{\sigma}(\boldsymbol{\tau}+\mathbf{1},\boldsymbol{t}) = F\left(\boldsymbol{\sigma}(\boldsymbol{\tau},\boldsymbol{t}), \epsilon, \boldsymbol{\beta}(\boldsymbol{\tau},\boldsymbol{t}), \alpha(t)\right). \tag{3.4}$$

Where $\Delta\boldsymbol{\sigma}(\boldsymbol{\tau}+\mathbf{1},\boldsymbol{t}) = \boldsymbol{\sigma}(\boldsymbol{\tau}+\mathbf{1},\boldsymbol{t}) - \boldsymbol{\sigma}(\boldsymbol{\tau},\boldsymbol{t})$. ϵ is the probability with which agents make mistakes, i.e. the probability with which they choose a random strategy and not the one as in Equation (3.2): this is a standard assumption in stochastic stability models. The vector $\boldsymbol{\beta}(\boldsymbol{\tau},\boldsymbol{t})$ is the vector of agents who are given the revision opportunity at time τ (as implied by the *inertia* assumption).

I define a parameter γ that represents the difference in the speed of evolution of the two processes. I define the dynamical system driving the evolution within generations and between generations as follows

$$\begin{cases} \gamma\Delta\boldsymbol{\sigma}(\boldsymbol{\tau}+\mathbf{1},\boldsymbol{t}) = F\left(\boldsymbol{\sigma}(\boldsymbol{\tau},\boldsymbol{t}), \epsilon, \boldsymbol{\beta}(\boldsymbol{\tau},\boldsymbol{t}), \alpha(t)\right) \\ \Delta\alpha(t+1) = \alpha(t)(1-\alpha(t))\left[\overline{\Pi}_C(\boldsymbol{\sigma}(\boldsymbol{\tau},\boldsymbol{t})) - \overline{\Pi}_M(\boldsymbol{\sigma}(\boldsymbol{\tau},\boldsymbol{t}))\right] \end{cases}. \tag{3.5}$$

To solve the system, I study the solutions of the system when $\gamma = 0$ and when τ goes to infinity and ϵ is small enough. Both the conditions on τ and ϵ are required by the stochastic stability process. In such a way, strategies converge, according to stochastic stability, much faster than types. Referring to the literature on fast-slow dynamical systems, putting $\gamma = 0$ is equivalent to calculating the solution of the system in the critical manifold, a proxy used in such dynamical systems to calculate the solution (see Cortez and Ellner, 2010). Using the critical manifold together with stochastic stability is reasonable under some assumptions that ensure that a stochastically stable distribution of strategies is reached within each generation t.

Firstly, as assumed in standard stochastic stability models, it must be that agents interact repeatedly within each generation, in order to infer the distribution of strategies at each τ (see Young, 1993a; Kandori et al., 1993). Moreover, Equation (3.3) is reasonable under the assumption that within each t, agents interact with all the other agents, in order to experience all the payoffs. Another usual assumption in standard stochastic stability is to set ϵ close to zero. Indeed, the convergence to a stochastically stable distribution of strategies is guaranteed by τ going to infinity and ϵ being very small (Young, 1993a). This means that I am assuming that τ goes to infinity for each generation and that ϵ is small enough at each generation. I divide the following assumption in two parts, with different implications on the convergence within a generation.

The first part concerns strategies converging before a new generation is born. I am implicitly assuming that within a generation agents are incapable of evolving their intelligence, while they are capable of converging to the stochastically stable distribution of strategies having fixed the fraction of intelligent agents in the population. In other words, I am assuming that strategies evolve faster than behavioral rules.

Assumption 1.

Within each t, the system implied by (3.4) converge to a stable distribution of strategies.

The second assumption states that not only strategies converge to a stable distribution of strategies, this distribution is also the stochastically stable one. Such an assumption relies on τ going to infinity and ϵ being small enough within each generation t.

Assumption 2.

Within each t, τ goes to infinity and ϵ is small enough to select the stochastically stable distribution of strategies.

Note that Assumption 1 differs from Assumption 2 since if the former is relaxed, it means that within a generation, agents may not converge to a stable distribution of strategies (they may still be converging to one of those when the new generation begins). On the other hand, if the latter is relaxed, agents still converge to at least one stable distribution of strategies within a generation, but they maybe stuck between two of them (since the error mechanism fails to select one of them as in stochastic stability). I test the robustness of my result to the relaxation of such assumptions in Section 3.3.5.

As it is discussed in seminal works in the stochastic stability literature such as Kandori et al. (1993) or Young (1993a), stochastic stability is a process that could require a considerable amount of time to be completed. Therefore, Assumption 1 needs to be discussed deeply. In my model, I consider the stochastic stability process of one generation to be limited by the fraction of conformists during that generation. Therefore, I may refer to the length of each generation as the long-run, while the length of the entire process as the ultra long-run.[8] Such an assumption can be justified for three main reasons. First, the stochastically stable distribution of strategies is the one such that agents spend the most time in it. Hence, in

[8] In this sense, I am referring to the ultra long-run in the way of Binmore and Samuelson (1999), which considers the long-run to be limited by initial conditions, while the ultra long-run to be such that structural situations can change and the convergence is no longer dependent on initial conditions.

the passing from one generation to another, I am giving weight 1, to the distribution of strategies which is the most likely to be observed for the previous generation. Second, the evolution of types can be viewed as the evolution of preferences or the evolution of learning rules which are usually assumed to evolve slower than strategies; third, it is arguable that the convergence of strategies can be quick in the case of study of this paper.[9] The rationale behind why the evolution of preferences or learning rules is usually assumed to be slower than the evolution of strategies is as follows. It is reasonable to suppose that for an agent, thinking about how intelligently they take a decision (i.e. their behavioral rule) implies an additional level of cognitive effort than simply thinking about which decision they should take (i.e. their strategy). Therefore, it is plausible to expect strategies to converge faster than behavioral rules. A similar reasoning applies to preferences. Concerning the speed of convergence of strategies in my model, it is reasonable to assume a fast convergence for mainly two factors. First, when $\alpha < \frac{1}{2}$, there is only one stable distribution of strategies, and second, when $\alpha > \frac{1}{2}$, the majority of agents (conformists) take decisions very quickly since conformism is a relatively easy behavioral rule to follow (do whatever the majority does).

In the next section I compute the results generated by the two above mentioned mechanisms.

Given the double selection mechanism and the many names, I summarize the important variables in the model in Table 3.2. For simplicity, from now on, I will refer to the term "distribution/s of strategies" as DS.

Variable	Referred to as	Meaning
α	Fraction of conformists	The fraction of conformists
$n_H(\tau)$	Fraction of hawks at time τ	The fraction playing H at time τ
$\boldsymbol{\sigma}(\boldsymbol{\tau})/\ \boldsymbol{\sigma}^*(\boldsymbol{t})$	Distribution of strategies (DS)	... at time τ/ stochastically stable in generation t
$\boldsymbol{\omega}(\boldsymbol{t},\boldsymbol{\tau})/\overline{\boldsymbol{\omega}}$	State	Both $\alpha(t)/\overline{\alpha}$ and $\boldsymbol{\sigma}(\boldsymbol{\tau})/\overline{\boldsymbol{\sigma}}$.

Table 3.2: Recap on the important notation.

3.3 Results

I divide this section into three subsections to make comprehension easier for the reader. Firstly (3.3.1), I show the stochastically stable DS for each composition of the population. Secondly (3.3.2), I compute the fitnesses of the types at each stochastically stable DS. Thirdly

[9] Concerning the evolution of preferences or learning rules, see Dekel et al. (2007), Kuran and Sandholm (2008) or Dridi and Lehmann (2015). For some considerations on the speed of evolution of strategies, see Oprea et al. (2011), Kreindler and Young (2013), Ellison et al. (2016) or Arieli et al. (2020).

(3.3.3), I show the stable composition of the population implied by the stochastic stability analysis.

Before I move on, it is useful to define the concepts of mild and harsh conflict that will be frequently used in the remaining of the paper.

Definition 2. *The conflict is harsh if* $\frac{v}{c} < \frac{1}{2}$.

Definition 3. *The conflict is mild if* $\frac{v}{c} > \frac{1}{2}$.

Trivially, the closer v gets to c the milder $\pi(H, H) = \frac{v-c}{2}$, i.e. the *conflict* between two hawks.

I also define some important DS for my results. I name $(1, \frac{v-\alpha c}{(1-\alpha)c})$ HM, $(0, \frac{v}{(1-\alpha)c})$ DM, $(0, 1)$ DH, and $(1, 0)$ HD. Throughout the rest of the paper, I call HM or DM as mixed DS since M types play a mixed strategy in these cases. Instead, I refer to HD and DH as type monomorphic DS since each type plays a different strategy in these situations. Importantly, in the mixed DS, the fraction of hawks is $\frac{v}{c}$, since in these DS, conformists play a pure strategy, while myopic best repliers play a mixed strategy such that the total fraction of hawks is $\frac{v}{c}$. In the type monomorphic ones, the fraction of hawks depends on α since in these DS, each type plays a different pure strategy.

3.3.1 Stochastic stability

In this section, I compute the stochastically stable DS for each level of α. The stochastic stability analysis is usually divided into two parts: unperturbed and perturbed dynamics. In the first part, absorbing DS are computed assuming that agents do not make mistakes. In the second part, first, it is assumed that agents make mistakes with a small probability, and then, the stochastically stable DS is determined considering that the probability with which agents make mistakes approaches the 0 and that time (τ) approaches infinity. The DS which requires the minimum fraction of errors to be reached starting from all the other DS is the stochastically stable DS. For details see Freidlin and Wentzell (1998), Foster and Young (1990), Kandori et al. (1993), Young (1993a), or Ellison (2000) (see also Wallace and Young, 2015 for practical intuition). I provide the details about the unperturbed dynamics in Appendix B.1 (see Lemma 12).

The next lemma states results of the stochastic stability analysis.

Lemma 7.

If conflict is mild, $\forall \alpha(t) \in [0, \frac{v}{c})$, $\boldsymbol{\sigma^*(t)} = HM$, $\forall \alpha(t) \in \left[\frac{v}{c}, 1\right]$, $\boldsymbol{\sigma^*(t)} = DH$.

If conflict is harsh, $\forall \alpha(t) \in [0, \frac{c-v}{c})$, $\boldsymbol{\sigma^{**}(t)} = DM$, $\forall \alpha(t) \in \left[\frac{c-v}{c}, 1\right]$, $\boldsymbol{\sigma^{**}(t)} = HD$.

I prove the results in Appendix B.1, and give the intuition here. Results are summarized in Figure 3.2.

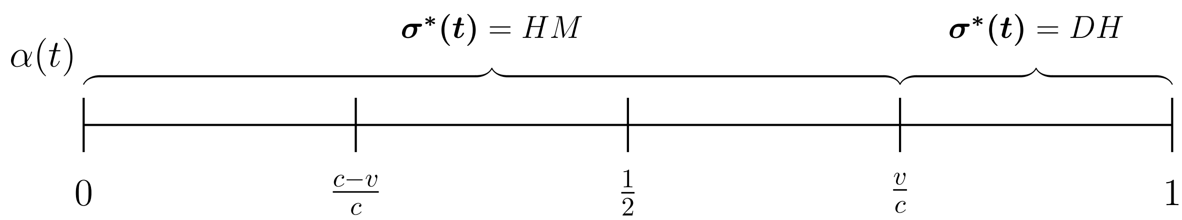

(a) Mild conflict

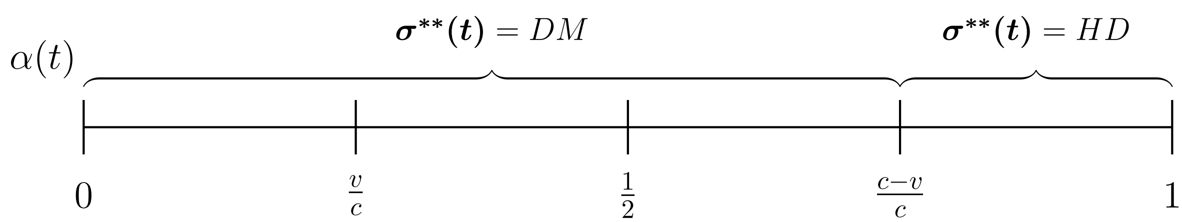

(b) Harsh conflict

Figure 3.2: Stochastically stable distributions of strategies for different levels of α.

Figure 3.2 highlights the importance of $\frac{v}{c}$ and $\frac{c-v}{c}$. From now on, I will refer to them as the critical thresholds.

Lemma 7 implies that unless conformists are a vast majority, the mixed DS is the stochastically stable one. This result is perhaps not surprising but deserves a better explanation.

When α is lower than the critical thresholds, the system always converge to a mixed DS since myopic best repliers are always in a fraction such that they influence the direction of the system. Therefore, the system is driven towards the DS which are stochastically stable with pure myopic best repliers, i.e. the mixed ones.

When α is higher than the critical thresholds, the long-run convergence is guided by conformists. Their grain of sand of intelligence pushes them to play H when the conflict is harsh and D when the conflict is mild. Indeed, if $\lambda_C = 0$, it takes the same fraction of mistakes to go from DH to HD than vice-versa. However, given that $\lambda_C \approx 0$, the only thing that drives the result is intelligence. For a myopic best replier, it takes fewer mutations to change from H to D than to change from D to H when conflict is mild and vice-versa when conflict is harsh. Hence, when conflict is mild, DH is the stochastically stable DS, and when conflict is harsh, HD is the stochastically stable one. This result is the opposite of Bilancini et al. (2022), although the result is driven by the same principle (greater cognitive ability pushes towards D when conflict is mild and towards H when conflict is harsh).

3.3.2 Evolution of types

In this section I analyze the evolution of types, computing Equation (3.3) at the stochastically stable DS from Lemma 7.

Lemma 8.

If conflict is mild,

- $\forall \alpha(t) \in [0, \frac{v}{c})$, $\forall \kappa > 0$, $\Delta\alpha(t+1) > 0$.
- $\forall \alpha(t) \in [\frac{v}{c}, 1]$, $\forall \kappa > 0$, $\exists \overline{\alpha} \geq \frac{v}{c}$ *s.t.* $\forall \alpha(t) < \overline{\alpha}$, $\Delta\alpha(t+1) > 0$, *and* $\forall \alpha(t) \geq \overline{\alpha}$, $\Delta\alpha(t+1) \leq 0$.

If conflict is harsh,

- $\forall \alpha(t) \in [0, \frac{c-v}{c})$, $\forall \kappa > 0$, $\Delta\alpha(t+1) > 0$.
- $\forall \alpha(t) \in [\frac{c-v}{c}, 1]$, $\forall \kappa > 0$, $\exists \overline{\alpha} \geq \frac{c-v}{c}$ *s.t.* $\forall \alpha(t) < \overline{\alpha}$, $\Delta\alpha(t+1) > 0$, *and* $\forall \alpha(t) \geq \overline{\alpha}$, $\Delta\alpha(t+1) \leq 0$.

Corollary 3.

If conflict is mild, and if $\kappa \leq \frac{2v-c}{2}$, $\overline{\alpha} = \frac{v}{c}$. *If* $\kappa > \frac{2v-c}{2}$, $\overline{\alpha} > \frac{v}{c}$.
If conflict is harsh, and if $\kappa \leq \frac{c-2v}{2}$, $\overline{\alpha} = \frac{c-v}{c}$. *If* $\kappa > \frac{c-2v}{2}$, $\overline{\alpha} > \frac{c-v}{c}$.

Lemma 8 and Corollary 3 set the last conditions for the main results of the paper. I prove the results in Appendix B.1 and give the intuition here. I summarize the results in Figure 3.3.

When α is smaller than the critical thresholds, conformists always grow in the stochastically stable DS. In other words, under mild and harsh conflict, conformists can always invade a population of myopic best repliers and evolve up to a point where the former outnumbers the latter. The intuition is as follows. As long as myopic best repliers can "influence" stochastic stability in the way previously discussed, the stochastically stable DS is the mixed one, where by definition, all the agents earn the same payoff. Therefore, since M types pay a cognitive cost, they have a lower fitness than conformists. On top of that, M types drive the result in a way that keeps "alive" conformists since, as it is shown in Lemma 8, they push the system towards the mixed DS. However, by doing so, myopic best repliers lead the system in a DS that is disadvantageous for them.

When α reaches the critical thresholds, the stochastically stable DS are the type monomoprhic ones. In these DS, there is a trade-off between being more or less intelligent. In fact,

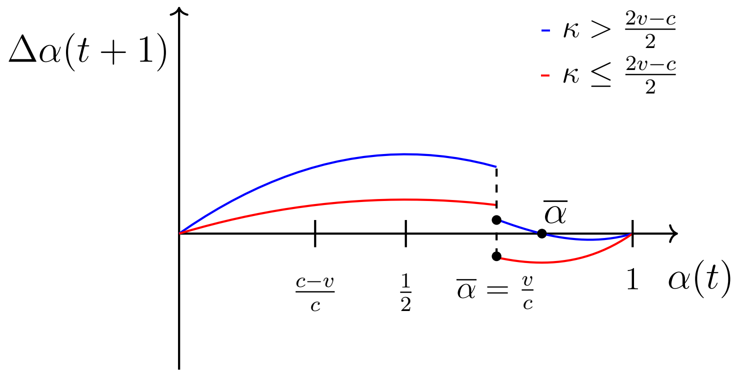

(a) Mild conflict

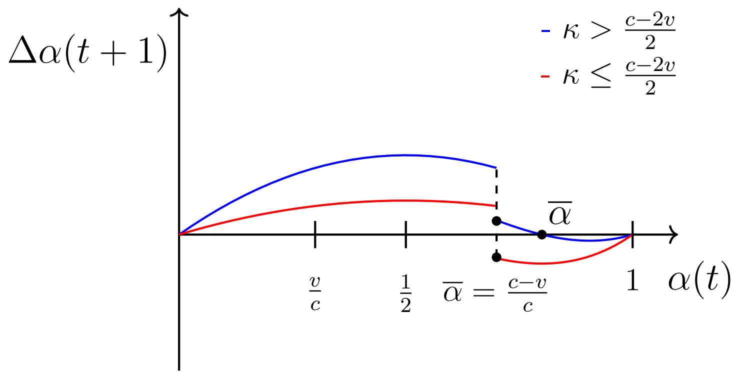

(b) Harsh conflict

Figure 3.3: The sign of the derivative of α with respect to time for different values of α.

M types have a benefit on C types since they play the correct best response to the DS. However, M types also suffer from the cost κ. If α is higher than the critical thresholds but still lower than $\overline{\alpha}$, the cost for being more intelligent is higher than the benefit for being so, and hence, the level of conformists continues to grow. The cost meets the benefit in $\overline{\alpha}$, so, such a composition of the population is stable. Importantly, whenever $\alpha > \overline{\alpha}$, the benefit for being myopic best replier is higher than the cost for being so; therefore, the level of conformists decreases. This makes $\overline{\alpha}$ globally stable (see Figure 3.3 for a graphical illustration). Note that $\overline{\alpha}$ depends on the level of κ: the higher κ, the higher the level of $\overline{\alpha}$. The discontinuity in Figure 3.3 comes from the fact that the stochastically stable DS changes when $\alpha > \frac{v}{c}$ $(\frac{c-v}{c})$.

3.3.3 Evolutionary stable state

I now state the main result of the paper.

Theorem 5.

Consider $a^* \in [0, \frac{c-v}{c}]$*, and* $a^{**} \in [0, \frac{v}{c}]$*. If conflict is mild,* $\overline{\boldsymbol{\omega}} = (\frac{v}{c} + a^*, DH)$*. If conflict is* *harsh,* $\overline{\boldsymbol{\omega}} = (\frac{c-v}{c} + a^{**}, HD)$*.*

Corollary 4.

If conflict is mild, and if $\kappa \leq \frac{2v-c}{2}$*,* $\overline{\boldsymbol{\omega}} = (\frac{v}{c}, DH)$*. If conflict is harsh, and if* $\kappa \leq \frac{c-2v}{2}$*,* $\overline{\boldsymbol{\omega}} = (\frac{c-v}{c}, HD)$*.*

The proof of both statements is straightforward and, therefore, is omitted. Indeed, the previous lemmas constitute the proof of Theorem 5, and the proof of Corollary 4 is a consequence of Corollary 3 (and Theorem 5). Results are summarized in Figure 3.4.

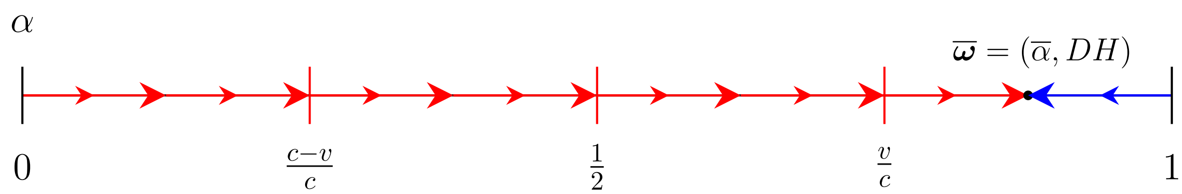

(a) Mild conflict for $\kappa > \frac{2v-c}{2}$

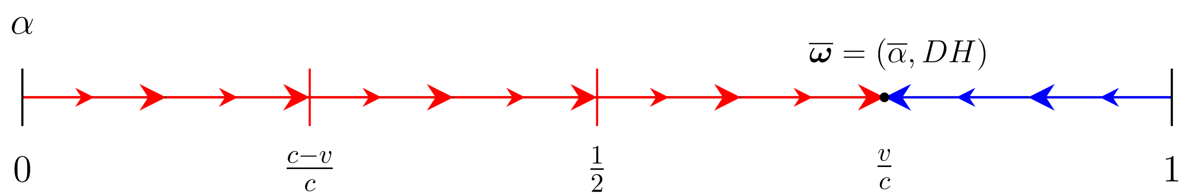

(b) Mild conflict for $\kappa \leq \frac{2v-c}{2}$

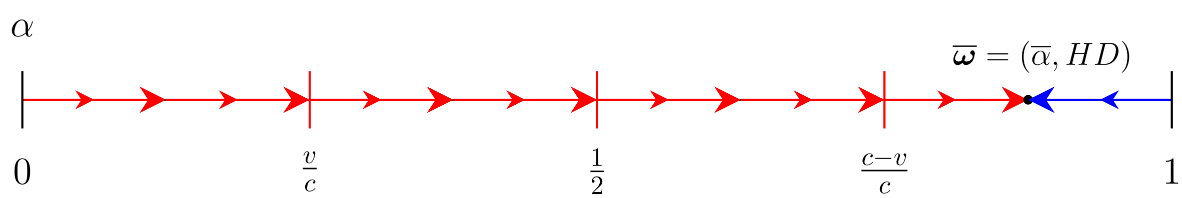

(c) Harsh conflict for $\kappa > \frac{c-2v}{2}$

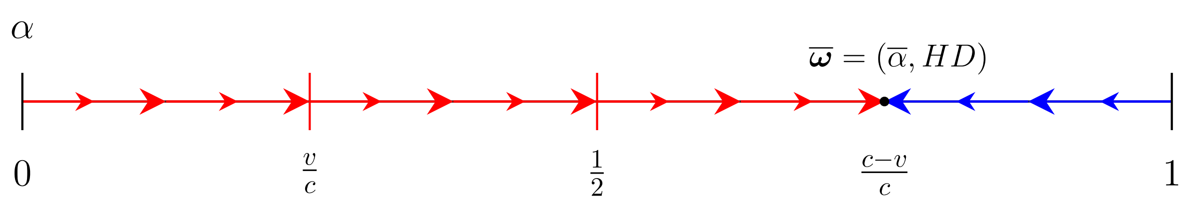

(d) Harsh conflict for $\kappa \leq \frac{c-2v}{2}$

Figure 3.4: The stable state in terms of the evolution of α.

First of all, the stable state is such that the type C is the predominant one in the ultra long-run. We know from Lemma 8 that as long as α is lower than the critical thresholds, α continues to grow. Since the critical thresholds are always greater than $\frac{1}{2}$, α is always going to be greater than $\frac{1}{2}$ in the ultra long-run. In my paper, I evaluate the robustness of Mohlin (2012) for level-k, introducing different levels of intelligence and a cognitive cost. Thanks to the introduction of κ, I can say that the level of conformists is always higher than the one of myopic best repliers in the stable state. Perhaps, this result is even more surprising than the one in Mohlin (2012) since conformists play a strategy that is far from approximating the best reply strategy in a Hawk-Dove game.

Secondly, in the stable state in the ultra long-run, the behavior of agents is not the one predicted by canonical models. Indeed, in a single population game, all canonical models (Smith and Price, 1973; Smith, 1982) agree on the fact that the only stable DS is the one where $\frac{v}{c}$ agents play H and $\frac{c-v}{c}$ agents play D. From Theorem 5, it emerges that the fraction of agents playing H is always lower or equal than $\frac{c-v}{c}$ in the case of mild conflict, while it is always greater or equal than $\frac{c-v}{c}$ in the case of harsh conflict. I formalize the result in the following corollary and summarise it in Figure 3.5.

Corollary 5.

If conflict is mild, $n_H^* < \frac{c-v}{c}$ *for* $\kappa > \frac{2v-c}{2}$*, and* $n_H^* = \frac{c-v}{c}$ *for* $\kappa \leq \frac{2v-c}{2}$.
If conflict is harsh, $n_H^* > \frac{c-v}{c}$ *for* $\kappa > \frac{c-2v}{2}$*, and* $n_H^* = \frac{c-v}{c}$ *for* $\kappa \leq \frac{c-2v}{2}$.

As discussed in Section 3.1, this result is complementary to Arigapudi et al. (2021). They show that due to heterogeneity and limited cognitive ability of agents, a two-populations Hawk-Dove may converge to a mixed DS instead of to a type monomorphic one. In my model, I show that due to a different kind of heterogeneity and limited intelligence of agents, a one-population Hawk-Dove game converges to a type monomorphic DS. Importantly, in my model, the heterogeneity arises endogenously since the level of conformists is determined by their fitness in the stochastically stable DS.

Thirdly, from Theorem 5, it emerges how more intelligent types behave in the stable state depending on the harshness of conflict. When conflict is mild, the more intelligent type (M type) plays H, while when conflict is harsh, they play D. My result differs from the one of Bilancini et al. (2022), although both results are driven by the same mechanism. In Bilancini et al. (2022), the agents with the greater cognitive ability (longer memory in their case) play D when the conflict is mild and H when the conflict is harsh. Even though in my model, the opposite happens, such a result is driven by the grain of sand of intelligence of conformists (the less cognitively able agents). The non-trivial part of my result comes from the fact that the stable level of unintelligent agents is above $\frac{1}{2}$ in the ultra long-run: due to this endogenous mechanism, the less intelligent agents behave the way they do.

In summary, when conflict is mild (harsh), conformists play D (H) in the stochastically

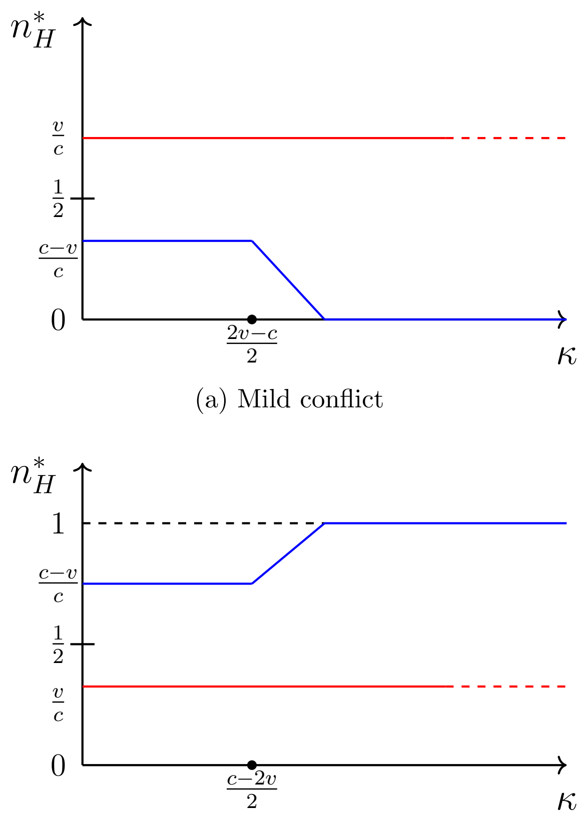

(a) Mild conflict

(b) Harsh conflict

Figure 3.5: The fraction of hawks in the stable state as a function of κ.

stable DS in the stable state: the best reply to this DS is H (D). Therefore, myopic best repliers play H (D) when the conflict is mild (harsh). Despite this result being interesting per se, its relevance comes from the fact that the high share in the population of conformists is endogenously determined.

One interesting question that may arise from these results is whether the spread of conformists is good or bad in terms of welfare. I will answer this question in the next section.

3.3.4 Welfare implications

In this section, I compare the welfare from Theorem 5 to the case when a mixed DS is the stable DS, in the style of Arigapudi et al. (2021), with similarities and differences between our works being discussed in previous sections.

To compute such an analysis, I first need to calculate the average payoffs in the stable state both in the case of mild conflict and in the case of harsh one. In the former case, the stable DS is DH; therefore, in this state, α agents play D, and $1-\alpha$ play H. I name $\overline{\Pi}(DH)$ the average payoff in the population in such a DS.

$$\overline{\Pi}(DH) = \alpha \left(\alpha \frac{v}{2}\right) + (1-\alpha) \left(\alpha v + (1-\alpha)\frac{v-c}{2}\right) - (1-\alpha)\kappa.$$

Similarly, I name $\overline{\Pi}(HD)$ the average payoff in the population in HD.

$$\overline{\Pi}(HD) = \alpha \left(\alpha \frac{v-c}{2} + (1-\alpha)v\right) + (1-\alpha) \left((1-\alpha)\frac{v}{2}\right) - (1-\alpha)\kappa.$$

Finally, I name $\overline{\Pi}(mixed)$ the average payoff in the population in HM or DM. Such a payoff is independent by which mixed DS is realized since $\frac{v}{c}$ is always the fraction of agents playing H.

$$\overline{\Pi}(mixed) = \frac{v}{c} \left(\frac{v}{c}\frac{v-c}{2} + \frac{c-v}{c}v\right) + \frac{c-v}{c} \left(\frac{c-v}{c}\frac{v}{2}\right) - (1-\alpha)\kappa.$$

Note that $(1-\alpha)\kappa$ is present in all of the three equations defined above; therefore, the cognitive cost will not be a disclaimer in determining the final welfare.

Corollary 6. *If conflict is mild* $\overline{\Pi}(DH) > \overline{\Pi}(mixed), \forall c, v > 0$. *If conflict is harsh,* $\overline{\Pi}(HD) < \overline{\Pi}(mixed), \forall \alpha > \frac{v}{c}$.

The proof is straightforward and therefore is omitted. From Corollary 6, it emerges that the welfare of agents depends on the harshness of conflict. Indeed, when conflict is mild, the welfare in DH is higher than the mixed DS (HM). The reason is simple: in DH, the majority of the population plays D, while in HM, the majority of the population plays H.[10] A similar reasoning applies to harsh conflict. Indeed, in this case, in the stable state $\boldsymbol{\omega}^*$, the stochastically stable DS is HD, where the vast majority of the population (i.e. conformists) plays H. In this case, the mixed DS (DM) is always better as long as $\alpha > \frac{v}{c}$, which is always the case given that $\alpha > \frac{c-v}{c} > \frac{v}{c}$ in the stable state. The reason for this inequality is straightforward: given that $\alpha > \frac{v}{c}$, and that all $i \in [0, \alpha]$ play H, there is a population of mostly hawks in a world where conflict is harsh.

[10] Note that the result does not depend on the fraction of agents playing H in the mixed DS, so it does not matter if we compare DH to HM or DM, but given that the mixed DS under mild conflict is HM I chose to compare DH with it.

3.3.5 Robustness checks

In this section, I relax some of the assumptions depicted in Section 3.2.4 to validate the robustness of my findings to possibly different settings.

In a first test, I study the robustness of the statement in Theorem 5 to non using stochastic stability to compute the prediction within a generation t. In other words, I relax Assumption 2. To understand the implication of this exercise, it means that whenever there is more than one possible distributions of strategies, I put equal weight on all the possible absorbing distribution of strategies. I summarize the results releasing such an assumption in the following corollary.

Corollary 7.

Consider a system where Assumption 2 is relaxed, $\overline{\alpha} \geq \frac{1}{2}$ $\forall \kappa > 0$. When the conflict is mild, $\overline{\alpha} > \frac{v}{c}$ if and only if $\kappa > \frac{2v-c}{4}$; when the conflict is harsh, $\overline{\alpha} > \frac{c-v}{c}$ if and only if $\kappa > \frac{c-2v}{4}$.

I provide the proof in Appendix B.1 and a simple intuition here for the case of mild conflict (the argument stands also for the case of harsh conflict).

Given that whenever $\alpha(t) \in [0, \frac{1}{2})$, there is only one absorbing distribution of strategies, that is a mixed one, conformists always grow from 0 to $\frac{1}{2}$. However, given the coexistence between a mixed distribution of strategies and a type monomorphic one for $\alpha(t)$ between $\frac{1}{2}$ and $\frac{v}{c}$, myopic best repliers have an advantage half of the times in this range. Therefore, conformists grow more than $\frac{1}{2}$ only if κ is big enough. In other words, the κ such that $\overline{\alpha} > \frac{1}{2}$ is higher without Assumption 2 than with it (see Corollary 4). Notably, the κ such that $\overline{\alpha} > \frac{v}{c}$ is lower without Assumption 2 than with it (see again Corollary 4).

In a second test, I relax Assumption 1 and 2 not allowing τ to go to infinite within each generation. In this way, the system may not converge to a stable DS within each generation. To implement this test, I simulate the results in NetLogo.[11] I summarized the parameters used in the model in Table D1; I condensed all results in Table D2 to D9 and in Figure D1 to D4. For completeness I test two different settings, in the first, each generation lasts 50 periods, while in the second each generation lasts 20 (I choose such kind of length because I wanted to prevent convergence within a generation). In order to assess the fitness, I took the average payoff of the last 4 (30) periods in the case of 50 periods generation, and the average payoff of the last 4 (12) periods in the case of 20 periods generation.

The first result which is evident from the first column of the tables is that still conformists prevail in the long-run even if Assumption 1 and 2 are relaxed. In line with results in Theo-

[11] The code is available at https://github.com/rrozzi/hawk-dove-uncertainty.

rem 5, when conflict is harsh (mild), the higher (lower) v, the lower the level of conformists in the ultra long-run; the higher κ, the higher the level of conformists in the ultra long-run. Secondly, it emerges that the behavioral prediction is reversed. According to Table D2 to D9, on average, conformists play H when conflict is mild and D when conflict is harsh. However, the level of conformists playing H (D) when conflict is mild (harsh) decreases with increasing value of κ. Once I set a $\kappa > \frac{2v-c}{2}$ ($\frac{c-2v}{2}$), the equilibrium where conformists play D (H) when conflict is mild (harsh) becomes more likely in line with Theorem 5. This result can be seen both from the average of n_H^C that decreases and its standard deviation that increases, indicating the alternation between different behaviors of conformists. However the other equilibria (HD in case of mild conflict and DH in case of harsh one) are still more likely. This result is due to the small length of each generation. Due to this phenomenon, the ergodicity property does no longer hold, and a sort of behavioral inertia drives the system towards the equilibrium of the previous generation. Such an equilibrium is often HD in case of mild conflict and DH in case of harsh one: as in Lemma 7, HD (DH) is stochastically stable for all levels of $\alpha \leq \frac{v}{c}$ ($\alpha \leq \frac{c-v}{c}$), in case of mild (harsh) conflict. The third result which we can observe by looking at the tables is that again the fraction of hawks in the stable state is different from the one predicted by canonical models. Indeed, the fraction of hawks is consistently different from $\frac{v}{c}$ across all specifications: this result is due to the fact that conformists do not play following the payoffs of the game, and they are the type with the biggest share in the population.

Concluding, from this second test, I shown that my results are robust to a certain degree to the relaxation of Assumption 1 and 2.

3.4 Extensions

In this section, I provide some extensions to the model presented in Section 3.2.

3.4.1 Uncertain environments within steps

In this section, I consider a model where agents play at each step, within generations one of two kinds of games. I introduce the following notation: $G = \{G_1, G_2\}$ is the set of games played by agents at each τ. Particularly, at each τ agents play G_1 with frequency p and G_2 with frequency $1 - p$. G_1 and G_2 are characterized as follows. The rationale behind such an extension is as follows: each day, agents face different types of situations, and their behavior might be different depending on how many situations of each type they face. For the moment, I am not assuming that agents condition their behavior on the type of situation (game) they face.

G_1

	H	D
H	$\frac{v_1-c}{2}, \frac{v_1-c}{2}$	$v_1, 0$
D	$0, v_1$	$\frac{v_1}{2}, \frac{v_1}{2}$

p

G_2

	H	D
H	$\frac{v_2-c}{2}, \frac{v_2-c}{2}$	$v_2, 0$
D	$0, v_2$	$\frac{v_2}{2}, \frac{v_2}{2}$

$1-p$

Table 3.3: G_1 and G_2.

Both G_1 and G_2 are Hawk-Dove games. However $\frac{v_1}{c} > \frac{1}{2}$, while $\frac{v_2}{c} < \frac{1}{2}$. Therefore, G_1 represents every situation of mild conflict, while, G_2 represents every situation of harsh conflict. Agent do not distinguish between G_1 and G_2, and therefore, they still play one strategy for both games, $s_i = \{H, D\}$. Consequently, the distribution of strategies is still $\boldsymbol{\sigma(\tau)}$, while the state is still $\boldsymbol{\omega(t,\tau)} = (\alpha(t), \boldsymbol{\sigma(\tau)})$. I still am interested in $\overline{\boldsymbol{\omega}} = (\overline{\alpha}, \overline{\boldsymbol{\sigma}})$. Equation (3.1), (3.2) and (3.3) still rule agents' decisions and the states evolution.

To improve the interpretation of the reader I define harsh and mild environments in the following way.

Definition 4. *The environment is harsh if* $\frac{pv_1+(1-p)v_2}{c} < \frac{1}{2}$.

Definition 5. *The environment is mild if* $\frac{pv_1+(1-p)v_2}{c} > \frac{1}{2}$.

For this section, I only present the main theorem, I provide the proof in Appendix B.1.

Theorem 6.

Consider $a^* \in [0, \frac{c-(pv_1+(1-p)v_2)}{c}]$, *and* $a^{**} \in [0, \frac{pv_1+(1-p)v_2}{c}]$. *If the environment is mild,* $\overline{\boldsymbol{\omega}} = \left(\frac{pv_1+(1-p)v_2}{c} + a^*, DH\right)$. *If the environment is harsh,* $\overline{\boldsymbol{\omega}} = \left(\frac{c-(pv_1+(1-p)v_2)}{c} + a^{**}, HD\right)$.

The intuition behind the result is straightforward. The convex combination between G_1 and G_2 is a Hawk-Dove game itself. In such a game, there is mild conflict if $\frac{(pv_1+(1-p)v_2)}{c} > \frac{1}{2}$ (if the environment is mild), and harsh conflict if $\frac{(pv_1+(1-p)v_2)}{c} < \frac{1}{2}$ (if the environment is harsh). From this observation, it is easy to determine the results in Theorem 6 from Theorem 5.

The above extension is rather simple: things could get more interesting when considering agents conditioning their actions on the game they play. I leave such a case for future extensions. In the following extension, I study these robustness of the results to a new definition of uncertainty, allowing myopic best repliers to condition their behavior (up to a certain degree).

3.4.2 Uncertain environments between steps

In this section, I add a further layer of uncertainty to the environment. I still consider G_1 and G_2 as defined in Table 3.3. However, I consider the game to be randomly determined at each step within generations. Specifically, I assume that the probability that G_1 is played at each τ is p, while the probability that G_2 is played is $1-p$. Agents know that they are playing G_1 (or G_2) at time τ, but they do not know which game they will play at time $\tau+1$: they believe that they will still be playing the same game they are playing at time τ. The interpretation of such an extension can be as follows: due to external shocks or institutional changes, the environments may change from time to time, but from one time to another, agents believe that the game they play will be the same.

The distribution of strategies is still $\boldsymbol{\sigma(\tau)}$, the state is still $\boldsymbol{\omega(t,\tau)} = (\alpha(t), \boldsymbol{\sigma(\tau)})$, and I am still interested in $\overline{\boldsymbol{\omega}} = (\overline{\alpha}, \overline{\boldsymbol{\sigma}})$. However, Equation (3.1) needs to be updated. Name $G_{k,\tau}$ the game played at time τ, with $k = \{1,2\}$. $U_i(s, n_H(\tau)|G_{k,\tau})$ is the utility earned by agent i for playing strategy s at time τ given that G_k is played, and similarly, $\bar{\pi}\left(s, n_H(\tau)|G_{k,\tau}\right)$ is the payoff of agent i from playing strategy s at time τ given that G_k is played.

$$U_i(s, n_H(\tau)|G_{k,\tau}) = \lambda_i \left[\bar{\pi}\left(s, n_H(\tau)|G_{k,\tau}\right)\right] + (1-\lambda_i)\mathbb{1}(s, n_H(\tau)). \tag{3.6}$$

Equation (3.2) and (3.3) are still unchanged. From Equation (3.6), combined with Equation (3.2), it emerges that agents myopically maximize their utility from time τ to time $\tau+1$, assuming that the game played at τ will also be the game in $\tau+1$. Concerning this assumption, agents are doubly myopic: they believe that the distribution of strategies and the game will be the same from τ to $\tau+1$. Importantly, in this application, agents know less about the uncertainty of the game than in the one in 3.4.1. Indeed, believing that the game will be the same from τ to $\tau+1$ is a good proxy of the true realization only for very high or very low levels of p (i.e. when the uncertainty is low).

Since the model does not have a closed form solution, I simulate the results for such a model using NetLogo.[12] The parameters used in the simulations are summarized in Table D10. I computed the results for four different couples of v_1 and v_2. Specifically, I chose v_1 and v_2 such that v_1 and $\frac{c-v_2}{c}$ are close or far from $\frac{1}{2}$. In two cases both v_1 and $\frac{c-v_2}{c}$ are either far or

[12] See again https://github.com/rrozzi/hawk-dove-uncertainty.

close to $\frac{1}{2}$, while in the other two cases v_1 is far from $\frac{1}{2}$ and $\frac{c-v_2}{c}$ is close to $\frac{1}{2}$ and vice-versa.[13] I displayed the results from these simulations from Table D11 to D22 and Figure D5 to D8.

From a first look at the tables, the most apparent result is that allowing for this kind of uncertainty favors the most intelligent behavioral rule. More precisely, when $\kappa = 0.01$, myopic best repliers always become the majority across all specifications. This result is due to the fact that even though myopic best repliers are doubly myopic, they react to the changes in the environment in a better way than conformists since the latter only react to changes in the share of agents playing hawk, while the former react to changes in the game's payoffs. A second look at the results for $\kappa = 0.01$ reveals a "U shape" relationship between p and $\overline{\alpha}$. Such a shape is coherent with the fact that extreme levels of p denote situations close to certainty in the game played at each step where we know (from Theorem 5) that conformists prosper. On the other hand, when there is more uncertainty (e.g. p between 0.3 and 0.7), the myopic best repliers adapt better to uncertainty, and therefore, with a low κ, they prosper more. The "U shape" relationship is asymmetric with respect to $p = 0.5$ when $v_1 \neq \frac{c-v_2}{c}$. This result is imputable to the fact that, even where there is no uncertainty in the game played, $\overline{\alpha}$ is smaller when v_1 or $\frac{c-v_2}{c}$ are closer to $\frac{1}{2}$ (see Corollary 4).

Interestingly, a small increase in κ (to 0.05) is sufficient to increase the level of $\overline{\alpha}$ up to a point where conformists outnumber myopic best repliers for almost all levels of p (except 0.05) across all specifications. This result shows that even for low levels for the cost of being more intelligent, the less intelligent behavioral rule could prosper: the explanation may be similar to the one of Theorem 5 since even though myopic best repliers react to the game switches in a better way, they still drive the system towards a distribution of strategies such that all players have the same fitness. Not surprisingly, if κ is greater than the thresholds of Corollary 4, conformists always outnumber myopic best repliers in the ultra long-run across all specifications. The last result suggests us that the first result of my paper is robust for high values of κ.

Another result from these simulations discloses the stable fraction of hawks under uncertain environments. Not strikingly, as p (the probability of playing the mild conflict game) increases, the fraction of hawks increases. This result is due to two different features. Firstly, if conformists are a few, the myopic best repliers push the system towards an equilibrium where the fraction of hawks exceeds $\frac{1}{2}$. Secondly, even when they are the majority, as shown in Section 3.3.5, as the ergodicity property no longer holds (inevitable in simulations), a behavioral inertia pushes the fraction of hawks towards $\frac{v_1}{c}$ > in G_1 (more likely for high

[13] Such a choice is coherent with the results in Theorem 5 (or Figure 3.4) such that the higher v when conflict is mild, the higher $\overline{\alpha}$, and the lower v when conflict is mild, the higher $\overline{\alpha}$.

levels of p) and $\frac{v_2}{c} > \frac{1}{2}$ in G_2 (more likely for low levels of p). Consistently with results in Section 3.3.5, higher levels of κ lead to higher standard deviation of the fractions of hawks across the simulations, meaning that when κ is big enough to kill all myopic best repliers, conformists may also end-up playing D when conflict is mild and vice-versa when conflict is harsh.

3.4.3 Hawk-Dove-Bourgeois

In this section, I extend the analysis to another version of the Hawk-Dove game called the Hawk-Dove-Bourgeois. The peculiarity of this game is that there is one more strategy called "Bourgeois" such that agents play H with probability $\frac{1}{2}$ and D with probability $\frac{1}{2}$. The interpretation of such a strategy is that an agent might be a hawk in certain kinds of situations (e.g. when they are the owner of the house and someone invades them), and a dove in other kinds of situations (e.g. when they invade someone else house).[14] Payoffs are the ones in Table 3.4. Note that the introduction of B makes the game non trivial. In fact, there is an additional pure NE that is (B, B), i.e. a "coordination" NE such that both players play B.

	H	D	B
H	$\frac{v-c}{2}, \frac{v-c}{2}$	$v, 0$	$\frac{v}{2} + \frac{v-c}{4}, \frac{v-c}{4}$
D	$0, v$	$\frac{v}{2}, \frac{v}{2}$	$\frac{v}{4}, \frac{3}{4}v$
B	$\frac{v-c}{4}, \frac{v}{2} + \frac{v-c}{4}$	$\frac{3}{4}v, \frac{v}{4}$	$\frac{v}{2}, \frac{v}{2}$

Table 3.4: The Hawk-Dove-Bourgeois game

I have to introduce some new notation to deal with this model. Starting from the strategies of each agent that are three: $s_i = \{H, D, B\}$. Moreover, in order to build the distribution of strategies we need one additional piece of information. I name $n_D(\tau)$ as the fraction of doves at time τ and maintain $n_H(\tau)$ as the fraction of doves; consequently, the fraction of bourgeois is $1 - n_H(\tau) - n_D(\tau)$. Equation (3.1) turns into

$$U_i(s, n_H(\tau), n_D(\tau)) = \lambda_i \left[\bar{\pi}\left(s, n_H(\tau), n_D(\tau)\right)\right] + (1 - \lambda_i)\mathbb{1}(s, n_H(\tau), n_D(\tau)), \tag{3.7}$$

where $\mathbb{1}(s, n_H(\tau), n_D(\tau)) = \begin{cases} n_H(\tau) & \text{if } s = H \\ n_D(\tau) & \text{if } s = D \\ 1 - n_H(\tau) - n_D(\tau) & \text{if } s = B \end{cases}$.

[14] See Foley et al. (2018) and Foley et al. (2021) for recent interesting work on this concept.

Moreover, I define $n_D^C(\tau)$ as the fraction of C types playing D at time τ and $n_H^C(\tau)$ as the fraction of C types playing H at time τ. Trivially, $\alpha - n_D^C(\tau) - n_H^C(\tau)$ is the fraction of C types playing B at time τ. Similarly, $n_D^M(\tau)$ is the fraction of M types playing D at time τ, $n_H^M(\tau)$ is the fraction of M types playing H at time τ, and $(1-\alpha) - n_D^M(\tau) - n_M^C(\tau)$ is the fraction of M types playing B at time τ. I define $\boldsymbol{\sigma(\tau)} = \big(n_H^C(\tau), n_D^C(\tau), n_H^M(\tau), n_D^M(\tau)\big)$. The state is still $\boldsymbol{\omega(t,\tau)} = (\alpha(t), \boldsymbol{\sigma(\tau)})$ and the main target of the analysis is still to individuate $\overline{\boldsymbol{\omega}} = (\overline{\alpha}, \overline{\boldsymbol{\sigma}})$. To distinguish this state from the one in the previous sections, I call the stable state for the Hawk-Dove-Bourgeois game $\overline{\boldsymbol{\omega}}_{\boldsymbol{hdb}}$.

Importantly, the underlying assumptions driving Equation (3.2) and (3.3) are still the same.

Lastly, I define $(0,0,0,0)$ as BB: in this DS, all agents play B. All the part of the analysis concerning stochastic stability and the evolution of types is in Appendix B.1. In this section I only state the main result in the following theorem.

Theorem 7.

$\forall \kappa > 0$, both under mild and harsh conflict, $\overline{\boldsymbol{\omega}}_{\boldsymbol{hdb}} = (1, BB)$.

The intuition behind this result is quite simple. Introducing B gives an element of coordination to the original game. Moreover, B is evolutionary stable in all standard models; that is, a population of only myopic best repliers converges to play B. This fact makes B the main attractor in all compositions of the population. For this reason, BB is the stochastically stable DS for all $\alpha(t) \in [0,1]$. If this is the case, $\Delta\alpha(t+1) > 0$, for every level of $\alpha(t) \in (0,1)$, that is, conformists always grow more than myopic best repliers for every value of $\alpha(t)$. This result is due to the fact that in BB, both types earn the same payoff, given that they coordinate on playing the same strategy. However, given that M types pay a cognitive cost $\kappa > 0$, they always have a lower fitness than C types. For this reason, $\alpha = 1$ is the only possible value in the stable state. Intuitively, the Hawk-Dove-Bourgeois introduces coordination in the Hawk-Dove game such that even intelligent players push towards coordination. Given that, by nature, conformists push towards coordination, both types of players will push towards playing B. Hence, there will be no trade-off between being more or less intelligent, and conformists will be the unique type in the ultra long-run.

This result is not surprising, but it generalizes the model in previous sections; it shows that in a game with an element of coordination, conformists prevail even more in the ultra long-run due to the nature of the interaction.

3.5 Conclusions

In this paper, I studied a double selection mechanism to assess the evolution of conformism and myopic best reply in the Hawk-Dove game. Myopic best repliers correctly play the strategy that best replies to the current distribution of strategies; conformists play the one chosen by the majority at the current distribution of strategies. Myopic best repliers pay a cognitive cost for being the most intelligent type. I consider the evolution of the distributions of strategies through stochastic stability and the evolution of behavioral rules through replicator dynamics. I find that conformists always outnumber myopic best repliers in the ultra long-run. Moreover, in the stable distribution of strategies, the fraction of agents playing hawk is never the one predicted by canonical models. In this distribution, conformists play dove if the conflict is mild and hawk if it is harsh, while myopic best repliers behave oppositely. Such a result has important implications for welfare. A population with a majority of conformists is always better off than a population with a majority of myopic best repliers if the conflict is mild, while it is always worse off if it is harsh.

I provide some relevant robustness checks to my result. I show that relaxing some critical assumptions does not affect most of the main predictions. Specifically, I show that the prediction concerning conformists being the majority in the ultra long-run is stable to many robustness checks that include the relaxation of the assumptions that make stochastic stability within each generation realistic. The prediction of the distribution of strategies depends on the cost of playing more intelligently. However, across all robustness checks, the fraction of hawks observed in the ultra long-run is consistently different from $\frac{v}{c}$ (the canonical prediction).

I also provide some extensions to my model. Firstly, I prove the robustness of my results to an environment where at each step, agents play with a certain frequency, a mild conflict game, and with a certain frequency, a harsh conflict game, but they do not condition their strategy on the game they are playing. Secondly, I study a case where the uncertainty is between steps, namely, a mild (harsh) conflict game is realized with a certain probability at each step. I show that in this environment, the survival of the least intelligent behavioral rule is harder (comes at a higher cognitive cost) than when the uncertainty is within steps. However, a small level of the cost of being more intelligent is sufficient to make conformists outnumber myopic best repliers in the ultra long-run. I also provide results for the Hawk-Dove-Bourgeois game, and I show that when an element of coordination is put in the game, there are no longer any benefits for being more intelligent. This mechanism brings to the extinction of myopic best repliers in the ultra long-run.

Chapter 4

Masks, cameras and social pressure

Abstract

In contrast to classical social norm experiments, we conduct experiments that *semi-continuously* randomise the share of individuals who are taking a particular action in a given environment. Using our experimental results, we are able to estimate the distributions of individual tipping points across our settings. We find that tipping points are very heterogenous, and that a substantial share choose to do the action (or not) regardless of what others are doing. We also show that, once embedded in dynamic models, our estimates predict that individuals will end up doing very different things despite engaging in copying-like behaviour.

Keywords: social norms, field experiment, dynamic models

JEL Classification Codes: D90, C93, C73

Notes. This chapter is a joint work with Itzhak Rasooly.

4.1 Introduction

There is a large literature demonstrating the power of peer effects and descriptive social norms across a range of domains. For instance, studies have found that we look to others when deciding whether to evade our taxes (Bott et al., 2020), donate to charity (Agerström et al., 2016) and even whether to vote (Gerber and Rogers, 2009). Of course, these examples are somewhat arbitrary: it is hard to think of even one activity that is not somehow shaped by our expectations about the behaviour of others.

Despite the obvious importance of social norms, however, current studies only provide limited evidence regarding the exact relationship between our beliefs about the share of people who do an activity and our own inclination to do that activity. To take a fairly typical example, consider Frey and Meier (2004)'s study of the impact of informing individuals that 64% as opposed to 46% of their peers donate to charity. While their experiment reveals that higher beliefs about the prevalence of charitable giving can lead to higher donation rates, it reveals little about the exact shape of this relationship; and more generally how this function looks over the full possible range of beliefs (0% to 100%).

There are at least three reasons why one might care about how actions depend on exact beliefs about prevalence (we call this relationship the f function). First, there is a clear policy motivation: the shape of this function reveals the returns to altering perceptions about prevalence (e.g. by disclosing information). Second, estimating the shape of this function allows us to test economic theories since certain economic models, e.g. those in evolutionary game theory, make distinctive predictions about the observed functional form. Third, the shape of this function turns out to be absolutely crucial for understanding long-run equilibria in dynamic models.

In this paper, we begin by elaborating on this third point by demonstrating theoretically how long-run equilibria in a plausible dynamic model depend on the shape of this function. We show that curvature is critical. If the function is first convex and then concave, then our dynamic system converges to an extreme equilibrium (as in Kreindler and Young 2013). On the other hand, if the function is first concave and then convex, we obtain convergence to an interior equilibrium; meaning that individuals end up doing very different things even though each of them is engaging in copying-like behaviour. As a result, understanding the curvature of this function is crucial for understanding long-run outcomes. In addition, we show that the intercepts of this function play a key role in pinning down its fixed points, thereby also shaping long-run equilibria.

This motivates our first experiment which aims to estimate the shape of this function in a particular setting, namely face mask usage. The basic idea of the experiment was straightforward. Subjects entered a room (one at a time) thinking that they were there solely to answer a decision problem involving lotteries. Unbeknownst to them, the number of the four experimenters in the room wearing a face mask had been randomised (leading to treatments in which 0/4, 1/4, 2/4, 3/4, or 4/4 experimenters were wearing a mask). We then observed whether each subject themselves chose to wear a face mask.

The experiment took place in Oxford over the course of nine days in February/March 2022. In total, we conducted fourteen three-hour sessions across twelve different colleges; and repeated our experimental protocol 646 times (each time with a different subject). Importantly, the experiment took place at a time in which face masks were no longer required by law or university rules, but still remained not abnormal. As a result, this was an ideal setting for capturing the implications of social pressure.

Our first experiment yielded four main results. First, according to our point estimates, the function is strictly increasing. That is, the greater the number of experimenters who were wearing a mask, the more likely were subjects to wear a mask. Reassuringly, this increasing relationship is evident across all of the specifications we estimate, including those that include and omit demographic controls and college fixed effects.

Second, we observe that many individuals defy social pressure. For example, 20% of the subjects chose to wear a mask even when none of the experimenters were wearing one (the 0/4 treatment); and 51% of the subjects did not wear the mask when all the other experimenters were wearing it (the 4/4 treatment). Similar results can be obtained by looking at changes, i.e. whether individuals chose to put on or take off a face mask during the experiment. For instance, out of the 106 subjects in the 4/4 treatment who entered the room without wearing a mask, only 39 chose to put on a mask during the experiment — a fact that illustrates the limits of social pressure in our setting.

Third, according to our point estimates, the largest jump in mask wearing arises between the 3/4 and 4/4 treatments. For instance, while increasing the number of mask wearers in the room from 1 to 2 experimenters raises the probability that a subject will wear a mask by around 4 percentage points, increasing the number of mask wearers from 3 to 4 raises the probability that a subject will wear a mask by a full 12 percentage points. This finding is consistent with an 'everybody effect' where social pressure becomes especially acute if everybody in the relevant environment chooses to do a particular activity.

Fourth, and perhaps most importantly, our estimated function has an interior fixed point, which is close to 23%. When embedded in our dynamic models, our results therefore suggest that, in the long-run, around 23% will choose to wear a face mask. As a result, calibrating our models using our estimates predicts convergence to an interior equilibrium, despite the existence of copying-like behaviour.

In order to assess the robustness of our findings, we conducted an analogous experiment in a very different context: camera use in online calls. The idea of this experiment was also straightforward. Subjects joined a Zoom call (one at a time) knowing only that they were attending in order to participate in an economics experiment. Unbeknownst to them, the number of the four experimenters on the call with their laptop camera on had been randomised (leading again to five treatments, corresponding to 0/4, 1/4, 2/4, 3/4, and 4/4 experimenters with their camera on). We then observed whether each subject themselves chose to use their video camera. In total, we repeated this process 1,114 times, leading to a sample size that was almost twice as large as that obtained in our first experiment.

Conducting this experiment led to similar, although not identical, results. We again find evidence of an everywhere increasing f function, i.e. that the share who use their camera is everywhere increasing in the number of experimenters who use their camera. We also again find high levels of non-compliance, with many participants choosing to use their cameras (or not) regardless of how many others are doing the same. Most importantly, once we use our estimates to calibrate our dynamic models, we again obtain convergence to an interior equilibrium, now with around 37% using a camera. Despite these similarities, the estimated f function in this context is not precisely the same as that estimated in the mask setting; and appears to be substantially more linear.

Finally, we discuss which models could give rise to our experimental findings; and could explain both the commonalities and differences between them. We observe that, assuming that all individuals have tipping point preferences, our f function can be interpreted as the cumulative distribution of individual tipping points. Viewed in this way, our experiments can be interpreted as an attempt to estimate the distribution of individual tipping points using randomisation. In both experiments, we find that individual tipping points are very heterogeneous, in contrast to canonical models in evolutionary game theory (e.g. Young 1993a). We also provide a simple model to explain where these tipping points come from. In our model, tipping points are the result of the interaction of intrinsic preferences to take the action along with (potentially non-linear) social pressure effects.

Our study contributes to a number of literatures across economics and related disciplines. First, our study contributes to the broader literature on the importance of peer effects and descriptive social norms. The current literature consists of a series of generally binary experiments across a variety of domains (see Cialdini 2007, Mascagni 2018 and Farrow et al. 2017 for reviews).[1] In contrast, our study is the first to semi-continuously randomise the share taking an action in subjects' immediate environment; and the first to do so in any setting (not just the settings of face masks and video calls).[2]

Second, our study contributes to the literature on tipping points and long-run dynamics. Especially relevant references include papers like Young (1993a), Kandori et al. (1993), Jackson and Yariv (2007), Young (2009) and Kreindler and Young (2013) in the economics literature; as well as the sociology literature following Schelling (1971) and Granovetter (1978) (see Dodds and Watts 2011 for an overview). Our study can be viewed as a first attempt to use randomisation in order to estimate the shape of the 'aggregate best response function' (or equivalently, tipping point distribution) that is crucial for driving the results of such models. The works by Damon Centola in the sociology literature are particularly relevant to this extent. In Centola and Baronchelli (2015) and Centola et al. (2018), the authors experimentally estimate the aggregate best response function in a different way from our experiment (see also Andreoni et al. (2021)). In the first experiment, they show how a group of subjects can converge to a convention even without knowing the share of the other subject doing one action nor the size of the population, while in the second experiment, they show in the same setting, how many subjects going against the status quo are necessary to break the convention.

Although these experiments have many advantages in studying dynamical systems compared to our study, we differentiate our aims and results because we show the effect of directly showing subjects the share of other people doing one activity. This is possible since, in our study, subjects observe what the people in the immediate environment ¡ are doing, while in the previously mentioned papers, they do not.

Other works using and/or trying to estimate the aggregate best response function is the one on harmful norms such as FGM, child marriage, or domestic violence, among others.

[1] Studies which provide some information about f functions in various contexts include: Cialdini et al. (1990); Cason and Mui (1998); Ichino and Maggi (2000); Borsari and Carey (2003); Heldt (2005); Fortin et al. (2007); Goldstein et al. (2008); Martin and Randal (2008); Krupka and Weber (2009); Gerber and Rogers (2009); Allcott (2011); Ferraro and Price (2013); Ayres et al. (2013); Costa and Kahn (2013); Bursztyn et al. (2014); Damm and Dustmann (2014); Smith et al. (2015); Thöni and Gächter (2015); Efferson et al. (2015); Lefebvre et al. (2015); Allcott and Kessler (2019); Linek and Traxler (2021).

[2] Our study also connects with the conformity literature following Asch (1951). In contrast to this literature, our study concerns individuals' actions (e.g. whether to wear a mask) as opposed to their cognitive judgements. Perhaps more importantly, our study also uses semi-continuous randomisation, in contrast to experiments in the Asch paradigm (see Bond and Smith 1996 for an overview).

Efforts in these studies have focused on finding a way to reverse these detrimental norms.[3] Many works have exploited evolutionary game theory works to try to empirically explain the persistence or reversing of these norms (among the others, see Howard and Gibson, 2017; Efferson et al., 2020; Novak, 2020; Gulesci et al., 2021). Although our studies share similar goals with these works, the fundamental difference between our work and the above mentioned ones is that we can directly estimate the effect of varying the share of people doing one activity on the subjects. For obvious reasons, it is harder to manipulate such a share in harmful norms; therefore, it is hard to study the direct effect of descriptive social norms in those settings. Although our results can hardly be compared to those mentioned above, we believe that the discussion in Section 4.5 applies more generally to those works.

Third, and more narrowly, we contribute to the literature on the social determinants of face mask wearing. The existing papers in this literature rely either on vignette-based experiments and surveys (Bokemper et al., 2021; Barceló and Sheen, 2020; Rudert and Janke, 2022; Goldberg et al., 2020; Barile et al., 2021) or instead observational data (Freidin et al., 2022; Woodcock and Schultz, 2021). We contribute to this literature by conducting the first ever randomised field experiment on the social determinants of face mask use.[4]

Fourth, we contribute to the literature on the social determinants of video camera use. Existing papers in this literature are again based on surveys: see, for example, Castelli and Sarvary (2021), Gherheș et al. (2021), Sederevičiūtė-Pačiauskienė et al. (2022) and Bedenlier et al. (2021). Our study is the first to examine this topic through use of a randomised field experiment.

The remainder of this article is structured as follows. Section 4.2 motivates our experiments with a theoretical discussion of the long-run implications of various f functions. Section 4.3 outlines the design of our face mask experiment and the associated results. Section 4.4 presents the design and results for our experiment on video cameras. Section 4.5 uses our results to calculate the distribution of individual tipping points across our contexts and discusses what could give rise to these distributions. Finally, Section 4.6 concludes with a

[3] Among the others, see Shell-Duncan and Herniund (2006); Shell-Duncan et al. (2011); Lee-Rife et al. (2012); Bicchieri et al. (2014); Bellemare et al. (2015); Efferson et al. (2015); Nyborg et al. (2016); Howard and Gibson (2017); Vogt et al. (2016); Efferson and Vogt (2018).

[4] Our use of a randomised field experiment allows us to side-step some of the issues that afflict previous studies of the social determinants of face mask wearing. For example, attempts to study this problem using hypothetical questions (as in Bokemper et al. 2021) suffer from the issue that individuals may not know what they would do in a hypothetical situation — an especially pressing concern since imitative behaviour may well rest on unconscious cognition. Meanwhile, attempts to study this problem using observational data (as in Woodcock and Schultz 2021) can suffer from both omitted variable bias and reverse causality issues (see Manski 1993 for an influential exposition of this latter point). Our randomised experiment avoids both of these issues.

discussion of future research suggested by our experiments.

4.2 Dynamics

To motivate our experiments, we begin by discussing how the relationship between beliefs about prevalence and the actual prevalence of an activity pin down long-run equilibria in a plausible dynamic model. Time is discrete, indexed by $t = 0, 1, 2, \ldots$. Let $\hat{s}_t \in [0,1]$ denote the belief (assumed to be common) about the share doing an activity at time $t \in \mathbb{N}$. The belief $\hat{s}_t \in [0,1]$ generates the actual share $s_t \in [0,1]$ via the function $f\colon [0,1] \to [0,1]$. That is, $s_t = f(\hat{s}_t)$ for all t. If we assume that $\hat{s}_t = s_{t-1}$, then we obtain the relation $s_t = f(s_{t-1})$; a dynamic process whose outcomes we can study. We will write f^t to denote the tth iterate of f. For example, $f^3(s_0) = f(f(f(s_0)))$.

In canonical evolutionary game theory models (Young, 1993a; Kandori et al., 1993; Young, 2009), there exists some 'tipping point' at which all individuals will switch from not taking the action to taking it. This gives rise to a convex and then concave shaped f function (see Figure 4.1(a) for a smooth version). While this is a plausible model in many settings, it is not the only possible model. For example, one might instead think that individuals are quantitatively quite insensitive, so treat shares like 40% and 45% as 'the same'. In contrast, however, one might think that there is an important qualitative difference between nobody and a minority taking an action, which leads to an f function which is steep near zero; and one might similarly assume an f function that is steep near 1. The resulting f function — which is reminiscent of the probability weighting function proposed in Kahneman and Tversky (1979) — is concave and then convex, and is plotted in Figure 4.1(b).

Proposition 1. Suppose that f is continuous, strictly increasing, and has three fixed points at $s = 0$, $s = \hat{s} \in (0,1)$ and $s = 1$. Then

- If f is convex on $[0, \hat{s}]$ and concave on $[\hat{s}, 1]$, then $\lim_{t\to\infty} s_t \in \{0,1\}$ provided that $s_0 \neq \hat{s}$.
- If f is concave on $[0, \hat{s}]$ and convex on $[\hat{s}, 1]$, then $\lim_{t\to\infty} s_t = \hat{s}$ provided that $s_0 \notin \{0,1\}$.

Proposition 1 shows how long-run equilibria crucially depend on the shape of the f function. If one assumes a convex and then concave f function, as in Kreindler and Young (2013), then one generically obtains convergence to an extreme equilibrium in which either nobody or everybody does the action.[5] We do not outline the dynamics in any detail since they are already familiar, but they are displayed graphically in Figure 4.1(a). On the other hand, Proposition 1 also states that if the f function is concave and then convex, then one obtains convergence to an *interior* equilibrium as displayed in Figure 4.1(b). This illustrates how differently shaped f functions can generate very different equilibria.

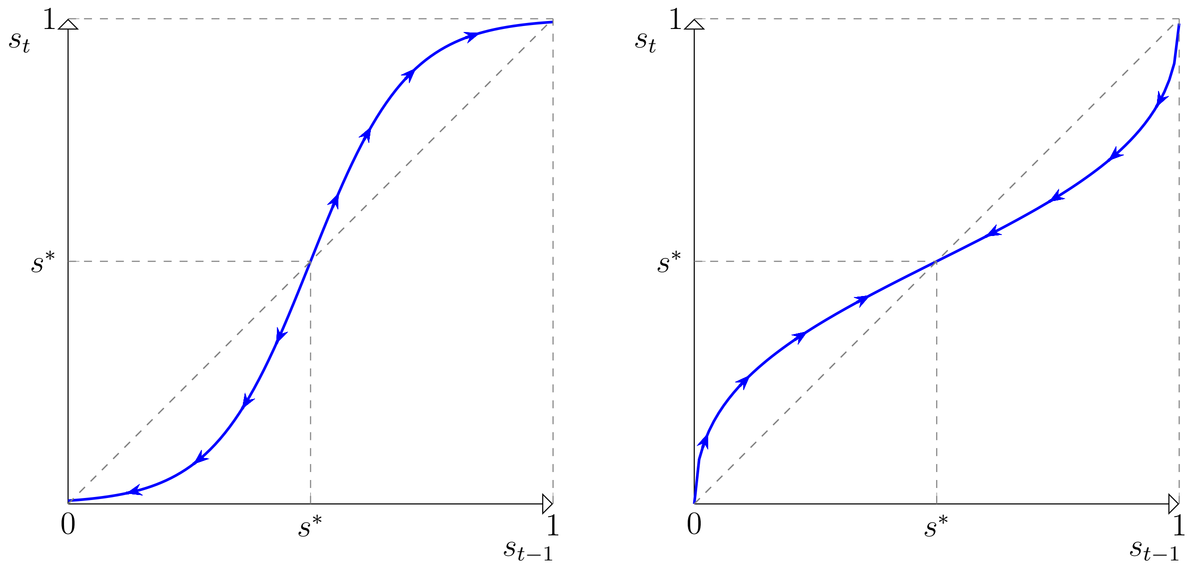

(a) Convex then concave (b) Concave then convex

Notes. This figure describes the evolution of s_t given two different f functions. When $f(s_{t-1}) > s_{t-1}$, the share doing the activity rises. When $f(s_{t-1}) < s_{t-1}$, the share doing the activity falls.

Figure 4.1: Two possible f functions

While Proposition 1 assumes that $f(0) = 0$ and $f(1) = 1$, this need not be the case. Instead, one might think that some individuals always do the action (leading to $f(0) > 0$) and that some others never do the action (leading to $f(1) < 1$). We now show that understanding the intercepts of the f function is also crucial for understanding long-run equilibria.

Proposition 2. Suppose that f is increasing. Then if s^* is the limit of $f^t(s_0)$ as $t \to \infty$, $s^* \in [f(0), f(1)]$.

Proposition 2 states that, assuming the f function is increasing, then the long-run share of individuals who do the activity is bounded by $f(0)$ and $f(1)$. Intuitively, this is because the fixed points of the f function must be bounded in this way; and any limit of the sequence $\{s_t\}_{t=0}^{\infty}$ must be a fixed point of f. As a result, estimating $f(0)$ and $f(1)$ can provide valuable information about long-run equilibria.

[5] One can then use stochastic stability arguments to identify which of these equilibria is more likely to emerge: see, for example, Young (1993a); Kandori et al. (1993).

In our experimental settings, behaviour is largely pinned down by beliefs about the share doing the action in the individual's immediate environment. Our results apply immediately to such cases, barring some discreteness issues, if one defines this environment as the relevant population. Alternatively, one can consider a network model in which individuals interact locally in small but interconnected communities (see Appendix C.2). As one might expect, this model yields very similar results.

In this section, our primary goal is not to insist on a particular dynamic model. Indeed, we believe that a large number of reasonable models are possible; and that one can make substantial variations on the assumptions made above. Rather, the main goal is to emphasise how, in any reasonable model, the shape of the f function is going to be a crucial determinant of long-run outcomes. This motivates our experimental investigation of f functions in the next two sections.

4.3 Masks

4.3.1 Experimental design

We now describe our first experiment aimed at estimating the shape of the f function in a particular context. The basic idea of the experiment was straightforward. Subjects entered a room thinking that they were there solely to answer a decision problem involving lotteries. Unbeknownst to them, the number of experimenters in the room wearing a face mask had been randomized. We then observed whether each subject themselves chose to wear a face mask (and how this varied with the number of experimenters wearing a mask in their immediate environment).[6]

This first experiment took place in Oxford in late February and early March of 2022. At this time, masks were not required by either law or university rules – however, they were also not unusual. This gave us an ideal setting in which to study the effects of social pressure. In total, we conducted 14 three-hour sessions in 12 different colleges over 7 days (with the help of 16 research assistants, some of whom participated in multiple sessions). On average, around 46 participants attended each session; which led to a total sample size of 646 experimental

[6] The experiment received approval from the University of Oxford's Departmental Research Ethics Committee (ECONCIA21-22-50). In line with the recommendations of the committee, we told subjects in advance that taking part in the experiment might involve interacting with unmasked individuals (which was common at the University of Oxford at the time). We also took reasonable social distancing precautions, including making sure that the experimental settings were well ventilated. We should also emphasise that, although we did not reveal the main purpose of our experiment to participants (as is not unusual in social science experiments), we did not explicitly deceive participants at any stage.

subjects (see Table C3 for the distribution of subjects across treatment groups).[7]

The structure of the experiment was as follows:

1. Subjects were asked to arrive at a room within a particular time slot.
2. Before each subject entered the room, the number of the four experimenters in the room who were wearing a mask (and the allocation of masks to experimenters) had been randomised. Thus, there were five treatment groups, corresponding to: 0/4 masks, 1/4 masks, 2/4 masks, 3/4 masks, 4/4 masks. We denote these treatments by T0, T1, etc.
3. Once a subject entered, they were asked to sit at a table in a way that gave them a clear view of the four experimenters. On the table were a box of masks as well as a bottle of hand sanitiser (such a set-up was common within the University of Oxford at the time). As a result, any subject who wished to wear a mask was able to do so.
4. Once the subject had sat down, each of the four experimenters introduced themselves by stating their name and subject of study. The purpose of this was to further ensure that each subject fully processed the number of experimenters who were wearing a mask.
5. The subject was asked their name, age, college and subject of study; and then given a decision problem involving lotteries.
6. We then asked the subject to leave the room, and repeated the process for the next subject (see Appendix C.4 for a more detailed description of the experimental protocol which includes the decision problem).

We recorded whether each subject was wearing a mask when they entered the room (this variable is labelled 'pre' in the tables). Naturally, we also recorded whether they chose to wear a mask after interacting with the experimenters. Finally, we recorded their choice in the lottery problem; as well as whether they asked if they ought to wear a mask (in such cases, each was told 'it's up to you' by the data recorder).

Based on post-experimental conversations, it seemed that most subjects believed that our goal was to measure risk aversion. Importantly, none of the subjects appeared to suspect that the experiment had anything to do with face masks; and there was nothing in the

[7] The experiment was pre-registered here: `https://www.socialscienceregistry.org/trials/9013`. The pre-registration contains the details about the design, the sample size we were expecting to obtain, and the approval from the Oxford ethical committee. Given that social norms have already been studied across different fields, we did not have the urgency to pre-recording our expected results. Indeed, most of the findings in the dedicated section are in line with other findings in the descriptive norms literature.

experimental design that could have revealed this.[8] This is reassuring since subjects might have acted in unnatural and unrepresentative ways if they had known that they were taking part in a face mask experiment.

Once all experimental sessions had been completed, we debriefed all subjects on the underlying purpose of the experiment. During the debriefing, subjects were given the opportunity to take part in an online survey. In the survey, subjects were asked to imagine that they entered a room and saw 4 people sitting around a table. They were then asked if they would wear a mask if none of the 4 people were wearing a face mask, if 1 of the 4 people were wearing a face mask, and so forth. Finally, they were asked to give an explanation for their answers, as well as whether they had contracted COVID-19 at any point during the pandemic. The purpose of the follow-up survey was to obtain some suggestive evidence on mechanisms, as well as some data on individual level f functions (see Section 4.5 for discussion).

4.3.2 Results

We now turn to our main results, beginning with a brief description of our sample. As shown by Table C4, our average participant was around 21 years old; and approximately half of our sample was male. Participants were fairly evenly distributed across subject divisions, with social sciences students being most represented (33% of the sample). Turning to Table 4.1, we see that genders, subjects and ages were reasonably balanced across our five treatment groups. However, we do observe some imbalance in the share of participants who entered the room wearing a mask: for example, the share is 27% in treatment T2 but only 14% in T0. Given that this variable turns out to be highly predictive for our outcome (whether participants continued to wear a mask), we control for it in our main specification.

Our regressions take the form

$$y_i = \beta_0 + \sum_{i=1}^{4} \beta_i T_i + \gamma x_i + u_i, \tag{4.1}$$

where y_i denotes whether an individual chose to wear a mask, the T_i are dummy variables indicating treatment assignment, and x_i is a vector of covariates (including whether they entered the room wearing a mask). In our main specification, we control for participant age, gender, and whether they entered the room wearing a mask (the 'pre' variable). However, we also report uncontrolled regressions, as well as regressions that use the full set of controls

[8] We also required all research assistants to sign an agreement specifying that they would keep the main purpose of the experiment confidential throughout its duration.

Variable	T0	T1	T2	T3	T4	p-value
Age	21.0	21.3	20.1	20.6	20.8	.143
	[.361]	[.539]	[.165]	[.219]	[.268]	
Pre	.142	.157	.266	.242	.203	.060
	[.031]	[.032]	[.039]	[.039]	[.035]	
Male	.535	.522	.461	.548	.421	.189
	[.044]	[.043]	[.044]	[.045]	[.043]	
Humanities	.323	.246	.250	.347	.256	.237
	[.042]	[.037]	[.038]	[.043]	[.038]	
Social	.268	.403	.336	.298	.353	.177
	[.039]	[.043]	[.042]	[.041]	[.042]	
MPLS	.213	.209	.305	.242	.233	.380
	[.036]	[.035]	[.041]	[.039]	[.037]	
Medical	.181	.104	.102	.105	.143	.235
	[.034]	[.027]	[.027]	[.028]	[.030]	

Notes. This table shows the average value of various variables across the five treatments. The variables are age, whether the subject entered wearing a mask ('pre'), gender, division of study (Humanities; Social Sciences; Mathematical, Physical & Life Sciences; Medical Sciences). The final column reports the p-value obtained from regressing the relevant variable on all treatment dummies and testing the hypothesis that the coefficients on all treatment dummies are equal to zero.

Table 4.1: Balance table (experiment 1)

that are available (including session and college fixed effects).

Figure 4.2 plots the results from our main specification (see Table 4.2 for the corresponding estimates, and Table C5 and C6 for the near identical results obtained by estimating probit and logit regressions; see also Figure C1 for the detailed figure with confidence intervals). The x-axis indicates the treatments, expressed as the fraction of experimenters wearing a mask (0, 0.25, 0.5, 0.75, 1). The y-axis displays the predicted share of individuals wearing a mask in each treatment. To obtain this predicted share, we set the three control variables (age, gender, and pre) equal to their mean values; so we are implicitly correcting for any imbalance in the pre variable. Thus, in the language of Section 4.2, Figure 4.2 displays our preferred estimates of $f(0)$, $f(0.25)$, $f(0.5)$, $f(0.75)$, and $f(1)$.

Several features of the data are apparent. First, we find evidence that the frequency of mask wearing is everywhere increasing in the share of experimenters who wear a mask. This pattern is evident in all the specifications that we estimate, regardless of whether they include controls, use logit or probit, etc. (again, see Tables 4.2, C5 and C6). From a statistical point of view, we can reject the hypothesis that lower treatments lead to the same levels of mask wearing as higher treatments for the large majority of treatment pairs (see Table C7), with the exception of the comparison of T0/T1 and the comparison of T1/T2. While we

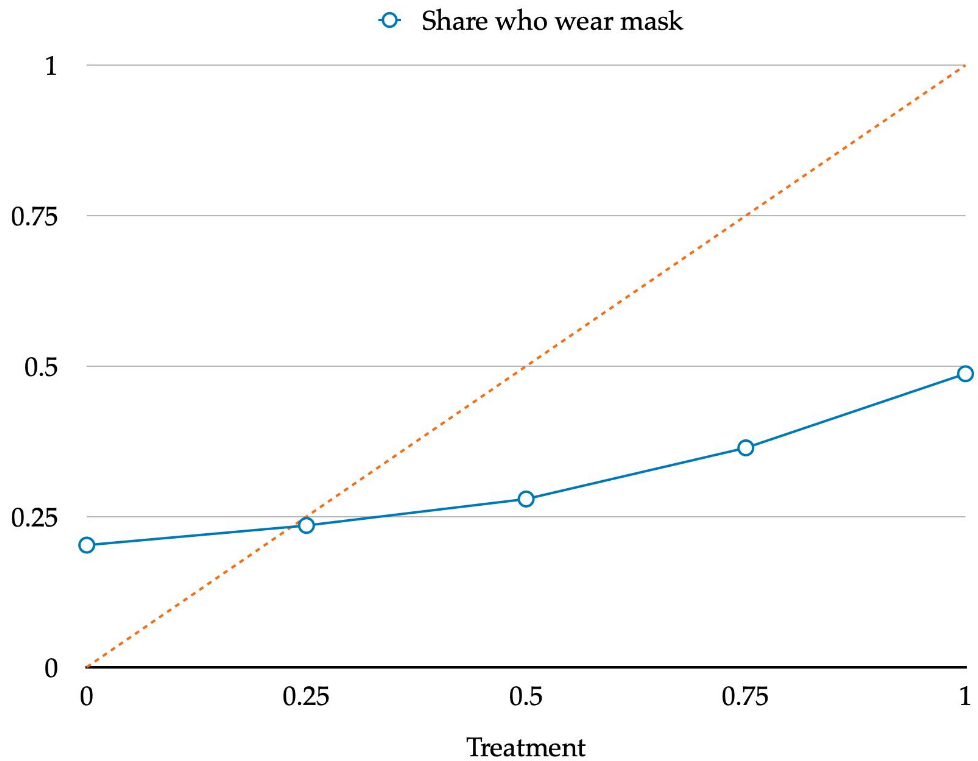

Notes. This figure shows how mask wearing varies by treatment group, after setting all covariates in the 'main specification' to their mean value.

Figure 4.2: Mask wearing by treatment group

discuss mechanisms later on, we note that this is consistent with a model in which higher rates of mask wearing lead to greater social pressure to wear a mask.

Second, we see that many individuals defy social pressure. In the treatment in which no experimenters wear a mask (T0), 20.0% of the participants nonetheless choose to wear a mask, a share which is statistically different from zero ($p < 0.0001$). In the language of Angrist et al. (1996), these people can be interpreted as 'always wearers', i.e. individuals who choose to wear a mask no matter how many others are doing the same (see Section 4.5 for elaboration). Similarly, in the treatment in which all experimenters wear a mask (T4), only 48.7% choose to wear a mask, which is again statistically different from 1 ($p < 0.0001$). The remaining 51.3% of individuals (who do not wear a mask) can be interpreted as 'never wearers', i.e. individuals who will never choose to wear a mask, no matter how many others are doing so (again, see Section 4.5 for a more formal discussion of this point).

Similar results are available if we look at changes. In treatment T0, out of the participants who entered the room wearing a mask, only 5.6% chose to take off their mask (see Table C8). Similarly, in treatment T4, out of the participants who entered the room not wearing a mask,

Variable	No controls	Main Specification	All Controls
Treatment 1	.044	.032	.020
	[.048]	[.029]	[.033]
Treatment 2	.171***	.078**	.075**
	[.053]	[.032]	[.035]
Treatment 3	.238***	.163***	.156***
	[.055]	[.039]	[.041]
Treatment 4	.331***	.284***	.289***
	[.054]	[.043]	[.046]
Pre		.757***	.741***
		[.029]	[.035]
Age		.002	.001
		[.005]	[.005]
Male		-.007	-.007
		[.026]	[.028]
Constant	.157***	.014	.130
	[.032]	[.107]	[.144]
n	646	646	646
R^2	0.070	0.494	0.517

Notes. This table reports our main regressions. To obtain the estimates in the first column, we regress whether subjects wore a mask on the treatment dummies. In the second column, we control for subject age, gender, and whether they entered wearing a mask. The third column also includes session and college fixed effects. Robust standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 4.2: Regressions (experiment 1)

only 36.8% chose to put on a mask. It is quite striking that the majority of those who entered without a mask in T4 decided to defy social pressure in this way, especially given that all four experimenters were clearly visible and that a box of masks was available.

Third, our estimated f function appears to be non-linear. Estimating a model with a quadratic term suggests some convexity ($p = 0.04$): see Table C9. Insofar as estimates appear non-linear, this is due to a large jump between the 3 and 4 treatments (the difference is 12 percentage points, as opposed to the average difference between treatments of 7 percentage points). This is indicative of a potential 'everybody effect', i.e. that a particularly large change in behaviour is induced by changing the share who are doing an action from 'most people' to 'everybody'.

Finally, we examine what our estimates imply when embedded in plausible dynamic models of the style discussed in Section 4.2. This delivers our fourth and perhaps most important finding: when embedded in such models, our estimates predict convergence to an interior equilibrium. In the very simple model discussed in Section 4.2, our model predicts global convergence to the fixed point of the estimated f function, which is about 23.3%.[9] We obtain very similar results when we use our estimates to calibrate our network model (see Appendix C.2), which predicts that around 23%-24% should wear a mask (with almost no dependence on the initial conditions). In these equilibria, around $0.20/0.23 \approx 87\%$ of mask wearers wear the mask because they always wear one; with the remainder wearing a mask due in part to copying behaviour. We discuss these interior equilibria in more detail in the next section (which obtains even more striking evidence of interiority).[10]

Relating Figure 4.2 to Figure 3 in Centola et al. (2018), we can highlight the differences between our findings and theirs.[11] In their figure, we can see that 25% of the population committed to doing a different action than the status quo is enough to eventually change the status quo in the "long-run". This effect is not possible in our setting (according to our estimates). Indeed 25% is approximately the percentage of people wearing the mask in the long-run. This effect is due to the fact that the social pressure induced by 25% of the people wearing the mask is not enough to induce the rest of the population to wear the mask.

Before moving to our second experiment, we briefly discuss the results of our online follow-up survey. As explained earlier, this survey directly asked participants how their decision to wear a face mask would vary with the number of individuals in the room who were also wearing a face mask. Given that individuals might not always know what they would do in a hypothetical situation, we do not emphasise the estimated f function obtained from this survey (although, reassuringly, it is also monotone increasing in the number of mask wearers). However, we use the follow-up survey to address two issues that our original experiment could not speak to, namely individual level f functions and mechanisms.

Our first finding from the online survey is that individual decision rules are plausibly monotone in the share of experimenters who are wearing a face mask. Indeed, over 99% of subjects report weakly increasing decision rules: if such subjects chose to wear a mask in some treatment Tk, they would also choose to wear the mask in all treatments Tk' for $k' > k$. This finding helps validate our assumption in Section 4.5 that individual preferences have a tipping point representation, which in turn provides an insightful decomposition of the

[9] This fixed point is obtained by linearly interpolating between $f(0)$ and $f(0.25)$. However, given that the fixed point is close to 0.25, near identical fixed points are obtained through other methods, e.g. quadratic interpolation.

[10] Although looking at Figure 4.2 it is easy to see that $f(0) \neq 0$ drives our result, this effect was not true a priori. It is easy to show that there could be also homogeneous convergence even if $f(0) \neq 0$ (e.g. when 1 is the only fixed point).

[11] Note that in that figure, the authors were estimating the effect of the size of a committed minority on the adoption of the action of the committed minority in the last round of the experiment. Although this figure is not directly the same as ours, the interpretations are similar, and therefore, we can compare the two of them.

observed aggregate behaviour. We should perhaps also stress that this finding cannot be obtained from the data from our main experiment, which is in principle consistent with the possibility that many individuals have decreasing decision rules.

Second, we obtain some suggestive evidence on why individuals are more likely to wear a mask if they see more mask wearing in their immediate environment. To do this, we consider only those individuals who reported that they would change their mask-wearing behaviour depending on the share of others wearing a mask. We then placed the explanation into various categories, including whether individuals were trying to avoid being judged, trying to put others at ease, or taking high rates of mask wearing as a sign of high COVID risk levels (see Appendix C.5 for a more detailed explanation of our categories along with examples). The main message from this exercise is that the health-based mechanism (i.e. that masks are used as a signal of COVID rates) is extremely unlikely to be driving our results: see Table C10 for details. Instead, the observed changes seem to be driven by a variety of social learning and social pressure mechanisms, although exactly identifying the relative importance of these mechanisms is challenging.[12]

4.4 Cameras

4.4.1 Experimental design

In order to study the generality of our results, we conducted a second experiment which used a near-identical methodology in a very different context. The basic idea of this second experiment was the following. Subjects joined a Zoom call knowing solely that they were taking part in some kind of economics experiment. Unbeknownst to them, the number of experimenters on the call with their video camera on had been randomised. We then observed whether each subject themselves chose to use their camera. Thus, this second experiment was essentially the same as the first, except with the subject of video-camera instead of face mask usage.[13]

This second experiment took place online in late July and early August of 2022. We conducted 16 two-hour sessions over the course of 8 days (with the help of 20 research assistants, some of whom participated in multiple sessions). On average, each session was attended by around 70 participants, leading to a sample size of 1,113 participants in total (see Table C11

[12] One particular issue is that individuals may not be entirely honest about the reasons for their behaviour. For example, they might overstate the extent to which their behaviour is driven by altruistic reasons (e.g. trying to put others at ease), as opposed to a fear of being judged.

[13] This experiment also received approval from the University of Oxford's Departmental Research Ethics Committee (ECONCIA21-22-44).

for the distribution of subjects across treatment groups). We recruited all participants from Prolific, and required all participants to have a working microphone and video camera.[14]

The structure of the experiment was as follows:

1. Subjects were asked to join a Zoom call at a particular time.
2. Before each subject joined the call, the number of the four experimenters in the meeting with their camera on (and which experimenters had their camera on) had been randomised. Thus, there were again five treatment groups: 0/4 cameras (denoted treatment T0), 1/4 cameras (T1), 2/4 cameras (T2), 3/4 cameras (T3), 4/4 cameras (T4).
3. Once a subject joined the call, all four experimenters introduced themselves by stating their name. The purpose of this was to ensure that each subject fully processed the number of experimenters whose cameras were on.
4. The subject was asked for their age, and whether they would hypothetically want to donate half of a bonus payment to the next subject on the call.
5. We then asked the subject to leave the call, and repeated the process for the next subject (again, see Appendix C.4 for a more detailed description of the experimental protocol).

Similarly to before, we recorded whether each subject had already turned their camera on when they joined the call; and whether they chose to turn their camera on after interacting with the experimenters. We also recorded their choice in the decision problem; as well as whether they asked if they ought to turn their camera on (in such cases, each was told that 'it's up to you'). Finally, if a subject had not turned their camera on at any point during the call, we asked them if there were any issues with their video camera.[15]

[14] The experiment was pre-registered here: `https://www.socialscienceregistry.org/trials/9829`. As for the masks experiment, we pre-registered the details about the design, the number of expected participants, and the ethics committee approval. As for the previous experiment, we did not pre-record any expected results. Since this second experiment was a follow-up to the previous one, the expected results were the results from the previous one; therefore, we felt even less of the need to pre-recording any expected result.

[15] Unsurprisingly, asking this question occasionally had the effect of prompting participants to turn their video camera on. In such cases, we still recorded such participants as having chosen to not use their camera (on the basis that they had chosen not to use their camera until effectively asked to do so).

4.4.2 Results

We now turn to our results, beginning again with a description of our sample. In contrast with the student population studied in our first experiment, the average participant in this experiment was around 42 years old, with a standard deviation of 13.9 years (see Table C12). Around 46% of the sample was male. As shown by Table 4.3, ages and genders were reasonably balanced across each of our five treatments. However, we again observe some imbalance in the share who joined the call with their camera on (the 'pre' variable), and so control for this variable in our main specification.

Variable	T0	T1	T2	T3	T4	p-value
Age	42.2	43.4	42.3	41.3	42.7	.615
	[.940]	[.931]	[.903]	[.906]	[.990]	
Pre	.116	.039	.058	.074	.070	.039
	[.021]	[.014]	[.016]	[.017]	[.018]	
Male	.472	.441	.439	.455	.516	.486
	[.033]	[.035]	[.033]	[.032]	[.034]	

Notes. This table shows the average value of various variables across the five treatments. The variables are age, whether the subject joined the call with their camera on ('pre'), and gender. The final column reports the p-value obtained from regressing the relevant variable on all treatment dummies and testing the hypothesis that the coefficients on all treatment dummies are equal to zero.

Table 4.3: Balance table (experiment 2)

Our regressions take the same form as Equation 4.1. That is, we regress whether an individual used their camera on the treatment dummies (using treatment T0 as the omitted category), and a vector of covariates. In our main specification, we control for participant age, gender, and whether they joined the call with their camera on. However, we once again also report uncontrolled regressions, as well as regressions that include the full set of possible controls (including session fixed effects).

Figure 4.3 plots the results from our main specification (see Table 4.4 for the corresponding estimates, and Tables C13 and C14 for the near identical results obtained by estimating probit and logit regressions; see also Figure C2 for the detailed figure with confidence intervals). Several points are apparent. First, similarly to the face mask experiment, we once again observe a monotone f function: the frequency of camera use is everywhere increasing in the share of experimenters who use a camera. This pattern arises in all of the specifications we estimate (see Tables 4.4, C13 and C14). In our main specification, we can reject the hypothesis that treatment i and treatment $i + 1$ lead to the same rates of camera usage ($p < 0.05$) for all i except $i = 3$; and we can always reject the hypothesis that treatment

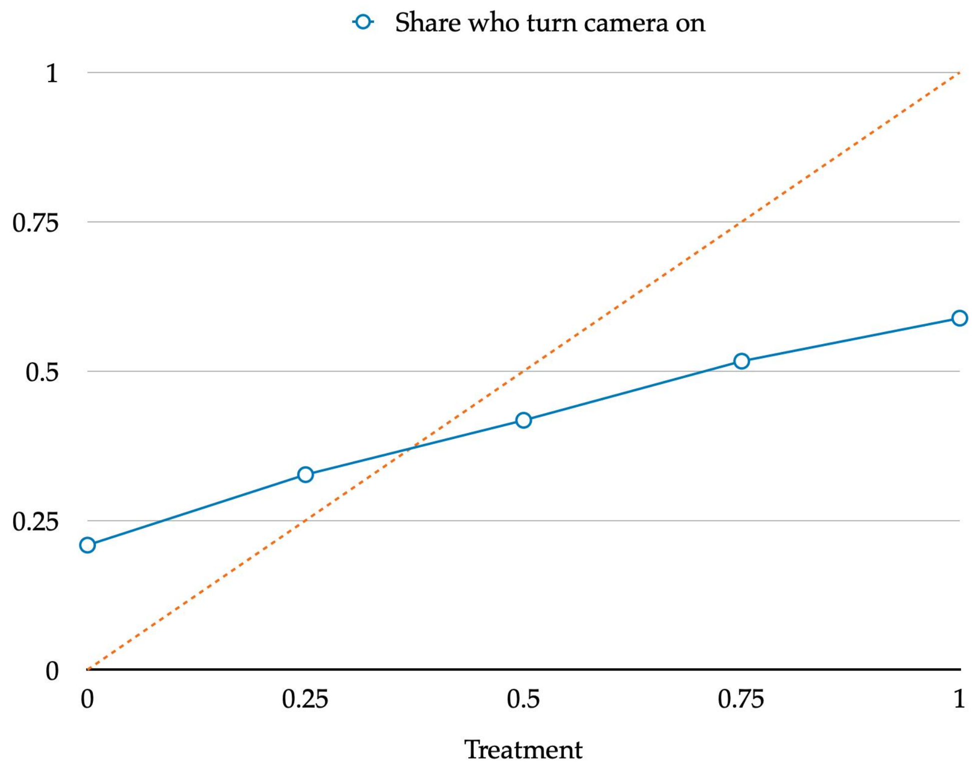

Notes. This figure shows how camera use varies by treatment group, after setting all covariates in the 'main specification' to their mean value.

Figure 4.3: Camera use by treatment group

i and treatment $i+2$ lead to the same rates of camera usage ($p < 0.01$) — see Table C15 for details. As before, this monotonicity is consistent with a model in which higher rates of camera use lead to greater social pressure to use a camera.

Second, we once again observe that many individuals defy social pressure. In the treatment in which no experimenters use a camera (T0), 20.9% of the participants nonetheless choose to use a camera, a share which is statistically different from zero ($p < 0.0001$). As explained in Section 4.5, such participants can be interpreted as 'always users', i.e. individuals who use a camera no matter how many others do the same. Similarly, in the treatment in which all experimenters use a camera (T4), only 58.7% choose to use a camera, which is again statistically different from 1 ($p < 0.0001$). The remaining 41.3% of individuals (who do not use a camera) can be interpreted as 'never users', i.e. individuals who will never choose to use a camera, no matter how many others are doing so. As before, similar results can be obtained by examining changes — see Table C16.

Third, the estimated f function in this context appears to be more linear. Statistically, we cannot reject a linear model: see Table C17. However, the jump between the 0 and 1

Variable	No controls	Main Specification	All Controls
Treatment 1	.077*	.118***	.125***
	[.043]	[.040]	[.041]
Treatment 2	.176***	.209***	.214***
	[.043]	[.039]	[.044]
Treatment 3	.281***	.308***	.320***
	[.043]	[.039]	[.049]
Treatment 4	.355***	.380***	.386***
	[.044]	[.041]	[.057]
Pre		.579***	.581***
		[.033]	[.034]
Age		.000	.000
		[.001]	[.001]
Male		.024	.023
		[.027]	[.027]
Constant	.241***	.155***	.094
	[.028]	[.047]	[.061]
n	1,113	1,111	1,109
R^2	0.069	0.161	0.183

Notes. This table reports our main regressions. To obtain the estimates in the first column, we regress whether subjects used a camera on the treatment dummies. In the second column, we control for subject age, gender, and whether they joined the call with their camera on. The third column also includes session fixed effects. Robust standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table 4.4: Regressions (experiment 2)

treatments (about 12 percentage points, in the main specification) is larger than the other 3 jumps (which are 9 percentage points, 10 percentage points, and 8 percentage points respectively). This provides some suggestive evidence on non-linearity, although one would need to obtain a larger sample to investigate this issue in greater detail. We should perhaps emphasise that, linear or not, our estimated f function is clearly different to that generated in standard evolutionary game theory models, which predict an S shape (see Section 4.2 for elaboration).

Fourth, and most importantly, our estimates once again predict convergence to an interior equilibrium when embedded in plausible dynamic models. When embedded in the model from Section 4.2, our estimates predict convergence to the fixed point of the estimated f function, which is about 37.0%.[16] We obtain similar results in our network model, which predicts that about 34.3% should turn the camera on (averaging across our models and initial

[16] This fixed point is calculated by linearly interpolating between $f(0.25)$ and $f(0.5)$. Given that the f function appears roughly linear, this seems like a sensible approach. Moreover, our network model delivers similar predictions without the need to rely on any kind of interpolation.

conditions). In the equilibrium of the simple model (for example), around $0.209/0.370 \approx 56\%$ of camera users use a camera because they always use one; with the remainder using a camera due in part to copying behaviour.

As for Figure 3 in Centola et al. (2018), we draw similar conclusions as in Section 4.3. We discuss the implication on canonical evolutionary game theory models (Young, 1993a, 2009). In these models the system typically converges to a situation in which either everybody or nobody does the relevant activity. In light of this, it may be worth explaining why our models predict interior equilibria (e.g. that 56% of individuals use a camera) despite the presence of copying-like behaviour. The explanation is as follows. According to our estimates, there is a substantial share of individuals who do the relevant activity no matter what others are doing. Given the behaviour of these individuals, others are induced to also do the activity, leading to a gradual increase in the share who do the activity (if the initial share is low). Eventually, the share of those doing the activity reaches the (unique) fixed point of our estimated f function, at which point the process stops. Likewise, if the share doing the activity is initially very high, then it gradually falls until it reaches the fixed point of our f function.

4.5 Discussion

In Section 4.2, we discuss what our estimates imply for long-run equilibria once embedded within dynamic models. However, we have not yet examined which models could give rise to our empirical findings; and what might explain the commonalities and differences in f functions across our two contexts. We turn to this question in the present section.

To provide a formal explanation for our results, let us assume that individual preferences have a tipping point representation. Formally, this means that, for every individual i, there exists a number $t_i \in [0, 1]$ such that the individual does the action ($a_i = 1$) if and only if $s \geq t_i$, where $s \in [0, 1]$ is the share of others who are doing it. Such an assumption seems very plausible in our two contexts; and it is given some empirical validation by the online survey discussed in Section 4.3.2.

Under this assumption,

$$f(s) = \frac{1}{n} \sum_i \mathbb{1}(a_i = 1|s) = \frac{1}{n} \sum_i \mathbb{1}(t_i \leq s), \tag{4.2}$$

where $\mathbb{1}(a_i = 1|s)$ is an indicator function that equals one if an individual takes the action (given that the share taking the action is s), and $\mathbb{1}(t_i \leq s)$ is an indicator function that equals one if an individual's tipping point exceeds the share. We thus see that, under our assumption, the f function is *exactly* the cumulative distribution of individual tipping points. Viewed in this way, our two field experiments can be seen as an experimental investigation of tipping point distributions.

Given that our experiments can be used to estimate the cumulative distribution of individual tipping points, it is straightforward to compute the probability distribution. To see how this works in practice, consider the data from the face mask experiment and define p_i as the share with a tipping point of i, for $i \in \{0, 1, 2, 3, 4\}$ (for convenience, we now use non-normalised tipping points, i.e. we do not divide by the population size). That is, for such values of i, p_i is the share of individuals who take the action if and only if they observe i or more of the four people in the room doing the same. Let us also define p_5 as the share who *never* do the action; and observe that $p_5 = 1 - \sum_{i=0}^{4} p_i$.

Table 4.5 reveals how the expected frequency of mask-wearing depends on the p_i parameters. In treatment T0, the only type who will do the action are the 'always doers', so the predicted share is p_0. In treatment T1, the types who do the action are the 'always doers' in combination with those who tip when they see one person doing the action. More generally, in treatment Tk, the expected share who will do the action is $\sum_{i=0}^{k} p_i$. Given this, one can estimate the p_i by matching the parameters with the sample frequencies (as suggested, e.g., by maximum likelihood). For example, we obtain the estimates $\hat{p_0} = 0.203$; and obtain $\hat{p}_1$, $\hat{p}_2$, $\hat{p}_3$, $\hat{p}_4$ by computing the difference of mask wearing between neighbouring treatments. Finally, our estimate for the 'never do-ers' is obtained using $\hat{p}_5 = 1 - \sum_{i=0}^{4} \hat{p}_i$.[17]

Treatment	Frequency	Predicted frequency
0	0.203	p_0
1	0.235	$p_0 + p_1$
2	0.279	$p_0 + p_1 + p_2$
3	0.364	$p_0 + p_1 + p_2 + p_3$
4	0.487	$p_0 + p_1 + p_2 + p_3 + p_4$

Table 4.5: Tipping points (experiment 1)

Figure 4.4 plots the results for both experiments. By construction, the distribution of tipping points plotted in the figure exactly generates the experimentally observed f functions. As a

[17] This approach implicitly treats the set of possible tipping points at discrete. It is also possible, however, to view (normalised) tipping points as continuously distributed on $[0, 1]$. In that case, our estimates can be used to calculate the share of normalised tipping points that are zero, the share that are between 0 and 0.25, etc.

result, it is possible to rationalise any observed differences across contexts by postulating a difference in the distribution of tipping points. For example, in the face mask experiment, a substantially lower share are estimated to have tipping points of 1 and 2 than in the Zoom experiment. This can be used to explain why the estimated f function is flatter over the $0-2$ range in the face mask experiment. Similarly, the observed non-linearity in the mask experiment can be rationalised by postulating that an especially large fraction have a tipping point of 4.

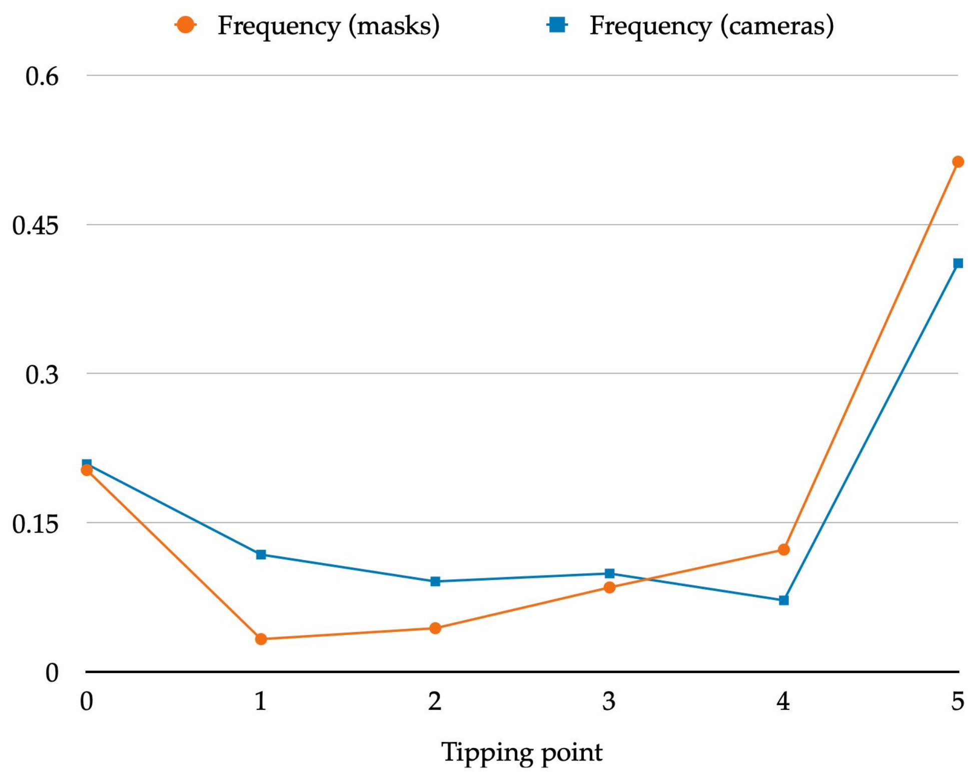

Notes. This figure shows the distributions of individual tipping points calculated from our two sets of experimental estimates.

Figure 4.4: Tipping point distributions

Although a model of heterogeneous tipping points is able to rationalise our results — and indeed can rationalise *any* non-decreasing f function — the explanation is rather mechanical. More precisely, while our results can be viewed as the cumulative distribution of tipping points, the question remains as to *why* the distributions take the form that they do. We now address this question using the simplest possible model of individual tipping points.

To this end, consider an individual i who is deciding whether to take the action or not given that a fraction s is already doing so. If they take the action, they obtain utility $u(a_i = 1) = \alpha_i + m(s)$, where $\alpha_i \in \mathbb{R}$ describes their intrinsic preference for taking the action and $m(s)$ describes their 'coordination payoff' from doing the same thing as a fraction s of their neighbours (assume m is differentiable). If they do not take the action, they get utility $u(a_i = 0) = m(1 - s)$. An (interior) tipping point $t_i \in (0, 1)$ is a share s that makes the individual indifferent between doing the action or not doing it, i.e. $\alpha_i + m(t_i) = m(1 - t_i)$.

To illustrate, consider the case in which of $m(s) = s$. In that case, it is easy to check that $t_i = 1$ for all individuals whose preferences satisfy $\alpha_i \leq -1$. Such individuals never take the action, no matter how many others are doing so. Similarly, we have $t_i = 0$ for all individuals for which $\alpha_i \geq 1$: such individuals always take the action. Interior tipping points satisfy the equation $\alpha_i + t_i = 1 - t_i$, or $t_i = 0.5(1 - \alpha_i)$ (see Figure 4.5). As a result, interior tipping points are strictly decreasing in the α_i parameter (and equal to 0.5 when the individual has no intrinsic reason to do the action, i.e. $\alpha_i = 0$). Intuitively, this is because individuals with a high intrinsic preference for taking the action will be indifferent between taking the action or not even when the share of others who are taking the action is very low.

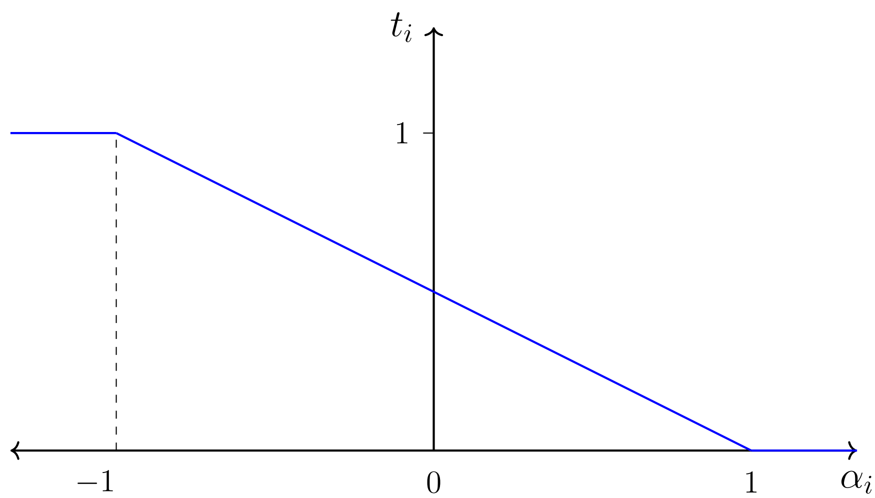

Notes. This figure shows how an individual's tipping point t_i varies with their intrinsic preference to take the action α_i when $m(s) = s$.

Figure 4.5: Tipping points when $m(s) = s$

We now verify that this result holds for *any* m function with an everywhere positive slope.

Proposition 3. *If $m'(s) > 0$ for all $s \in [0, 1]$, then each individual has a well defined tipping point $t_i \in [0, 1]$. Furthermore, if $t_i \in (0, 1)$, then $\frac{\partial t_i}{\partial \alpha_i} < 0$.*

Proposition 3 says that (interior) tipping points are strictly decreasing in an individual's intrinsic preference for taking the action. As a result, it provides us with an explanation for the differences we observe across our two contexts. In the mask experiment, the estimated tipping points tend to be higher (again, see Figure 4.5). One simple way to rationalise this, as suggested by Proposition 3, is to postulate that individuals are generally less willing to wear masks than to use their laptop cameras.

Although intrinsic preferences for taking actions do plausibly influence tipping points in the manner just discussed, we should emphasise that they are not the *only* factor that determines the shape of the f function. In our view, the f function is driven by the complicated interplay of both intrinsic preferences (captured by α_i) *and* potentially non-linear social pressure effects (captured by m). Future work could attempt to decompose these two channels.

4.6 Concluding remarks

In this paper, we conduct multi-treatment social norm experiments to obtain a quantitative understanding of how individuals' behaviour varies with the share doing an action in their immediate environment. Despite some differences between the estimates across our contexts (which we rationalise using a simple theory), we obtain many commonalities across the two experiments: increasing f functions, high levels of non-compliance, etc. Perhaps most importantly, when embedded in dynamic models, our estimates can explain how copying can plausibly lead to heterogenous behaviour. In our view, this is an important insight with applications to many settings (e.g. religious adoption, fashion trends, etc.): despite the ubiquity of social pressure and copying behaviour, different individuals nonetheless often end up doing very different things.

It may be worth briefly emphasising why our models generate convergence to interior equilibria, as opposed to the extreme equilibria predicted by canonical game theory models (e.g. Young, 1993a, 2009; Kandori et al., 1993). In these models, each individual effectively has the *same* tipping point, which in turn generates the S shaped f function discussed in Section 4.2. This in turn leads to three fixed points, of which only two (the extreme equilibria) are stable. In contrast, our empirical estimates suggest: i) reasonably high levels of non-compliance ii) substantial heterogeneity in individual tipping points. These two factors generate an f function with a unique fixed point in the interior, which is the globally stable equilibrium.

Despite the large number of social norm experiments, we believe that our findings open up several avenues for future research. First, it may be worthwhile to conduct more experiments with semi-continuous randomisation in additional contexts. In particular, this could provide further evidence on whether our key finding of interior equilibria is robust. Second, it may be worthwhile to conduct such experiments with an even larger number of treatment groups, thus allowing for a more fine grained estimate of the f function. Given the very large sample sizes required to do this, however, such experiments are likely to be even more logistically challenging to implement than the two field experiments whose results we report here.

Chapter 5

Does homophily impede human capital investments?

Abstract

We consider a game between multiple candidates and a decision-maker. There are two types of candidates: A and B. Each candidate chooses his social group and skill (high or low). The decision-maker must assign each candidate to four kinds of tasks. Type A candidates fit for certain kinds of tasks, and type B fits for others; only high skilled candidates can do specific tasks. The decision-maker observes each candidate's social group but needs to pay to observe his type and skill. We find that when homophily intensity is high enough, the unique possible kind of equilibrium is a separating equilibrium in social groups. In this equilibrium, social groups are informative about candidates' skills and types, and the decision-maker needs not to pay to know the candidates' types and skills. In such a case, the decision-maker never pays to see candidates' skills and types, and candidates never choose the high skill. If homophily intensity is weaker, there could be pooling equilibria on social groups; however, in none of them, the decision maker buys the information, and neither candidates invest in their skills.

Keywords: homophily; costly cognition; costly signal acquisition.
JEL Classification Codes: D82; D83; Z13

noindent *Notes.* This chapter is joint work together with Ennio Bilancini and Leonardo Boncinelli.

5.1 Introduction

In many situations, employers spend a lot of time and resources allocating employees to job tasks. This practice is challenging since individuals' abilities are hidden or not easily observable. However, social attitudes may support employers in this process. For example, it is well known that due to homophily, people with a similar feature, such as ability, are more likely to frequent the same social group (McPherson et al., 2001; Jackson, 2021). Since people from the same social group are distinguishable due to similar identity choices, employers may learn enough information about employees' abilities by looking at their social groups.

In this paper, we show a novel mechanism through which social attitudes could affect labor market outcomes. We show that if candidates with similar abilities hang out together (due to homophily), the employer never screens them to know their types or abilities, and therefore, they never invest in their skills. We illustrate our findings with the help of an example. Consider the human resources manager of a firm (HRM). She needs to allocate a pool of candidates to four categories of tasks: accountant, CTO, salesman, and CFO. Nature assigns each candidate to one of two types, i.e. types are mutually exclusive. Types represent innate abilities, e.g. relational and cognitive.

Candidates decide whether to invest in their human capital or not. If they do it, they become specialists in the ability they own by nature. Let us call high cognitive (relational) types those candidates with cognitive (relational) ability who invested in their human capital, and low cognitive (relational) types those with cognitive (relational) ability who did not invest in their human capital. Each candidate also joins one between two social groups, e.g. cool guys or nerds. Candidates all prefer the accountant job to the salesman one, CTO to the accountant one, and the CFO to the salesman one. Candidates choose the social group based on their social preferences. If there is homophily, candidates favor frequenting candidates of the same type and dislike frequenting candidates of the other type. If there is no homophily, candidates are inclined towards spending time with other candidates regardless of their type. As stated above, people belonging to the same social group make similar identity choices and are easily perceivable. For example, HRM could distinguish nerds from cool guys because the former wear plaid shirts and the latter wear suits.

The HRM assigns each candidate to one of the four tasks. She needs high cognitive types for the CTO position and high relational types for the CFO position. She needs low cognitive types for the accountant and low relational types for the salesman positions. The HRM always observes each candidate's social group, and she decides between two assignment policies: assigning tasks based on candidates' social groups or giving each candidate a trial period to observe his skill and ability. The first option is less time-consuming and cheaper for the firm, but only implementing the second, she learns candidates' abilities and skills.

Imagine that candidates all have the same capacity to invest in their human capital. The HRM will always have correct beliefs about their skills in equilibrium. However, she may be uncertain about their type of abilities. Two relevant scenarios can happen.

In the first scenario, all candidates frequent both relational and cognitive types. In this way, the HRM cannot know candidates' abilities from their social groups. She can only implement the trial period for everyone to know their abilities, provided that the cost of implementing it is low enough. In such a case, HRM screens the candidates, and they are not assigned to high-skilled tasks unless they invest in their human capital. Given that high-skilled tasks pay more than low-skilled ones, all candidates will invest in their human capital in this case. If the cost of implementing the trial period is too high, the HRM does not implement it, and candidates do not invest in their human capital. Therefore, she assigns each candidate to the same low-skilled task (salesman or accountant) depending on which ability is more frequent in the population.

In the second scenario, candidates of the same type hang out in the same group. For example, all cognitive types frequent nerds, and relational types frequent cool guys. In such a case, the HRM always knows candidates' abilities by looking at their social groups. She does not need to implement the trial period for each candidate. Consequently, The HRM does not test candidates' skills, and they do not invest in their human capital. Consequently, the HRM will assign all cool guys to the salesman position and all nerds to the accountant position, but she will assign no one to the CTO or the CFO.

Our main result is that the first scenario happens if there is no homophily or if there is homophily and the intensity of social preferences is low enough (*weak homophily*), while the second scenario only happens if there is homophily and social preferences are intense enough (*strong homophily*). The motivation behind these results is straightforward. If there is no homophily or *weak homophily*, candidates have the incentive to look like cognitive types since the cognitive type of jobs pay more. However, if there is *strong homophily*, candidates' incentive to all look like cognitive types is lower than their preferences for hanging out with similar types. Therefore, they will separate into two social groups, and the HRM can distinguish between the two types of candidates without the trial period.

Importantly, even under *weak homophily*, the case where candidates choose the same social group and invest in their skills is not stable. The reason is that if the HRM implement the trial period, candidates have the incentive to deviate and frequent their own types since the

HRM learns their types and skills.

Welfare implications follow directly from these results. When the information's cost is low enough, candidates are better off under no homophily since they are assigned to high-skilled tasks. The firm is better off under no homophily than under strong homophily, provided that the cost for the trial period policy is lower than the benefits of hiring candidates with high skills. When the information's cost is high enough, the firm is always better off under strong homophily than under no homophily. Indeed, candidates do not invest in their human capital in any case. However, only under strong homophily does the HRM correctly allocates candidates by observing their social groups. Candidates are better off under no homophily than under strong homophily if they all prefer the task they are assigned to under no homophily.

The remaining of the paper is organised as follows. In Section 5.1.1 we give a literature review. In Section 5.2 we present the model, in Section 5.3 we give the principal results, in Section 5.4 we discuss them, and in Section 5.5 we offer some concluding remarks. We prove the principal theorems in Appendix D.1 and we provide and prove the results in mixed strategies in Appendix D.2.

5.1.1 Related literature

Our model tackles three different fields of the literature: strategic information acquisition, endogenous group formation, and homophily.

Recently, the strategic information acquisition literature progressed in two directions: costly information acquisition and rational inattention. While there are some differences between the two approaches, they share the same founding idea: it takes a cost to learn information. The applications of these models vary from packaging decisions to job market signaling. The original idea of costly information is due to Grossman and Stiglitz (1980).

Since this seminal contribution, many models in this literature have analyzed the role of costly information in the markets' functioning. Some of these models focus on costly messages (Di Pei, 2015; Gentzkow and Kamenica, 2014; Argenziano et al., 2016), other focus on costly acquisition of signals (Caillaud and Tirole, 2007; Glazer and Rubinstein, 2004); Dewatripont and Tirole (2005) combines both approaches. Bilancini and Boncinelli (2018a) and Bilancini and Boncinelli (2018b) consider two variants of a sender-receiver game where a coarse reasoner receiver can pay to acquire hard information on the state of the world.

Concerning this literature, we study the coexistence of separating and pooling equilibria in a so-far unexplored way. In a similar model, Bilancini and Boncinelli (2018c) noticed that in a standard signaling model with costly acquisition of signals, there cannot exist any other pooling equilibrium that is not the one on the low signal. We consider a model with two kinds of signals (social group and skill). In this way, we can study the role of social connections in the screening process. We are the first to examine this mechanism in the literature to the best of our knowledge. We show that due to the costly nature of information, separating equilibria on social groups reduces the uncertainty faced by the decision-maker. This mechanism makes the purchase of the information less profitable. Hence, it also makes investment in human capital less attractive for candidates. On the other hand, pooling equilibria on social groups boost the decision-maker's uncertainty and information profitability. Therefore, under pooling equilibria on social groups, the decision-maker is more likely to purchase the information than under separating equilibria on social groups, and candidates are more likely to invest in their human capital.

Rational inattention contributions are built around the original model by Matějka and McKay (2015) (see Maćkowiak et al., 2023 for an exhaustive review of the field). Of the many existing models, we refer to Fosgerau et al. (2020) (FWS from now on) since it is closely related to our model. FWS considers a pool of candidates that differ in traits and categories. Candidates can invest in their qualifications to raise the probability of being hired, and a rational inattentive screener decides whether to hire candidates based on the information she acquires about them (in a similar model, Matveenko and Mikhalishchev, 2021 applies rational inattention to a labor market context, where candidates do not have an active role). In our model, we introduce an endogenous social group choice and the multiplicity of tasks' dimensions. Due to the endogenous social group selection, candidates may choose their social group strategically to signal their type. Due to the multiplicity of tasks, we can study homophily's role in driving information acquisition and human capital investment. Specifically, we show under which circumstances the presence of homophily reduces the profitability of buying the information and, therefore, the profitability of human capital investment for candidates.

The mathematics behind our model links our paper to the prejudice literature (see Phelps, 1972, Fryer and Jackson, 2008, Becker, 2010 or Bertrand and Mullainathan, 2004), and more importantly to the statistical discrimination literature (see Arrow, 1973, Spence, 1973, Coate and Loury, 1993 or Moro and Norman, 2004).

Closely related to our model, many works in economics and identity field focus on the effect of segregation on inequalities (see Benabou, 1996, Durlauf, 1996, Mookherjee et al., 2010 or Bowles et al., 2014). Even more related to our model, works like Fang (2001), Desmet and

Wacziarg (2018) or Kim and Loury (2019) treat group formation as an endogenous choice. These models study how endogenous group formation affects employers' stereotypes or people's performances in the labor market. Our purpose is different. Instead of studying ethnic/cultural group formation, we consider group formation based on the affinity of abilities. Furthermore, rather than focusing on hiring decisions, we focus on task assignments. Due to this difference, we can draw additional results from the previous works on endogenous group formation. Specifically, we can show how the existence of homophily may help employees allocate employees to different tasks at the cost of impairing the human capital investment of employees.

Lastly, we contribute to the homophily literature. Although homophily can have many forms (McPherson et al., 2001), in our model, we consider a specific form of homophily involving abilities. Homophily has been modeled in different ways: intrinsic preference for interacting with similar people (Currarini et al., 2009), higher probability to meet a similar person (Jackson and Rogers, 2007 or Bramoullé and Rogers, 2009), or it can mean that people better empathize with their similar (Kets and Sandroni, 2019). However, most of these approaches treat homophily as exogenously given (excluding the third): our approach is an endogenous version of the one adopted by Currarini et al. (2009).

For what concerns labor market outcomes, there are many channels through which homophily operates. As an example, the above mentioned statistical discrimination and prejudice may be driven by homophily (see Jacquemet and Yannelis, 2012, and Edo et al., 2019). Jackson (2021) gives a broad overview of other channels through which homophily operates. The first one is social connections: there is evidence that employers favor potential candidates that are linked with current workers in their firms (see Arrow and Borzekowski, 2004, Calvo-Armengol and Jackson, 2004 Patacchini and Zenou, 2012, Clauset et al., 2015, Beaman et al., 2018, Bolte et al., 2020, or Okafor, 2020). This difference in referrals can also lead to disincentive people with fewer contacts to invest in education (see Bowles et al., 2014 again). A second channel is the access to valuable information (see Golub and Jackson, 2012, Chetty et al., 2020, or Aybas and Jackson, 2021). A third channel concerns norms and peer pressure (see Jackson and Rogers, 2007, or Jackson and Storms, 2019).

We contribute to the homophily literature in two ways. Firstly, we study a novel mechanism through which homophily may influence labor market outcomes. Due to homophily, candidates with similar abilities congregate in the same social groups. Therefore, employers identify candidates' abilities more easily, even though this effect may impede the human capital investment of employees. Secondly, our model differs from most homophily models because, in our model, homophily is an endogenous social preference in candidates' utility. Through this introduction, we can capture for which degree of homophily, candidates with the same ability frequent the same social group. Indeed, we study a context where candidates perceive a trade-off between frequenting the same social group and polarizing into two different social groups. In such a context, polarization only happens if the intensity of social preferences is high enough compared to material preferences.

5.2 Model

There is a continuum of candidates of mass 1 and a decision-maker DM (sometimes referred to as "she"). Each candidate i is one of two types $t_i \in \{A, B\}$ (types are mutually exclusive): the proportion of A types (p_A) is greater than the proportion of B types (p_B).

Each candidate selects two actions: his skill $s_i \in \{L, H\}$, and his social group $k_i \in \{x, y\}$. The social group choice and the skill L are free, while the skill H is costly. $\zeta > 0$ is the cost of playing H. We say that candidates that choose H invest in their human capital. Candidates' social groups are common knowledge, while each candidate's type and skill are his private information. We name $K : \{x, y\} \to [0, 1]$ the distribution of candidates' social group choices (which represents the probability that a candidate chooses social group x), $T : \{A, B\} \to [0, 1]$ the distribution of candidates' types (the probability that a candidates is of type A) and $S : \{L, H\} \to [0, 1]$ the distribution of candidates' skill choices (the probability that a candidate selects skill H).

The decision-maker observes the social group of each candidate at no cost. She can acquire information about each candidate's type and skill. We denote this decision by $d \in \{d_0, d_1\}$, where d_0 is the choice of not acquiring the information and d_1 is the choice of acquiring it. If DM acquires the information, she pays a cost $c > 0$. When she plays d_1, she knows each candidate's skill and type. If she plays d_0, she only observes the social group of each candidate.

Furthermore, DM has to assign each candidate to one of four tasks: α_H for A types with H skill, α_L for A types with L skill, β_H for B types with H skill, and β_L for B types with L skill. We denote with $m \in \{\alpha_H, \alpha_L, \beta_H, \beta_L\}$ a generic task. If DM does not acquire the information, she can only assign tasks conditioning on the social group; if she acquires the information, she assigns tasks conditioning on the type and the skill. We denote with $\mathbf{m_k} = (m_x, m_y) \in \{\alpha_L, \alpha_H, \beta_L, \beta_H\}^2$ the action of DM if she does not acquire the information, and with $\mathbf{m_s} = (m_{AH}, m_{AL}, m_{BH}, m_{BL}) \in \{\alpha_L, \alpha_H, \beta_L, \beta_H\}^4$ the action when she acquires it. We call $\mathbf{m} = (\mathbf{m_k}, \mathbf{m_s}) \in \mathbf{M}$ the vector containing these choices. If DM acquires the information, she learns both the type and the skill of the candidate. We assume that DM always allocates the tasks correctly when she acquires the information. We call $\mathbf{m}^*_{\mathbf{s}} = (\alpha_H, \alpha_L, \beta_H, \beta_L)$ the optimal assignment for $d = d_1$. We define the common prior that a candidate is of type A as p_A and the common prior that a candidate is of type B as p_B. We assume that the common prior is correct; hence, $p_A > p_B$.

The prior beliefs held by DM are $\mathbf{p}(\mathbf{t}|\mathbf{k})$=$(p(A|x), p(A|y))$, and $\mathbf{p}(\mathbf{s}|\mathbf{k})$=$(p(H|x), p(H|y))$. The first couple are the probabilities that a randomly screened candidate with social group x or y is of type t, while the second are the probabilities that he has skill s.

The posterior beliefs of DM are the probabilities that a candidate i is of type t or of skill s, given the information choice of DM and the social group choice of i. They are $\mathbf{p}(\mathbf{t}|\mathbf{k};\mathbf{t_i},\mathbf{d})$=$(p(A|x;t_i,d), p(A|y;t_i,d))$, and $\mathbf{p}(\mathbf{s}|\mathbf{k};\mathbf{s_i},\mathbf{d})$=$(p(H|x;s_i,d), p(H|y;s_i,d))$. If DM plays d_1, the social group has no more informativeness. Moreover, in this case, her posterior beliefs are always correct; hence, they are either 1 or 0. If DM plays d_0, her posterior beliefs are equal to her priors.

A strategy for a candidate i is a pair (s_i, k_i). A strategy for DM is a vector $(d, \boldsymbol{\omega})$, where $d \in \{d_0, d_1\}$ and $\boldsymbol{\omega} = \begin{cases} \boldsymbol{\omega}_{d_0} & \text{if } d = d_0 \\ \boldsymbol{\omega}_{d_1} & \text{if } d = d_1 \end{cases}$ is a function that gives the vector of tasks assignments $(\mathbf{m})$, given DM decisions, and candidates' distributions of social group choices, skills and types. We denote with $\boldsymbol{\Omega}$ the set of all $\boldsymbol{\omega}$ functions. Note that $\boldsymbol{\omega}_{d_0} : [0,1] \to \mathbf{M}$ and $\boldsymbol{\omega}_{d_1} : [0,1]^2 \times \{d_0, d_1\} \to \mathbf{M}$.[1] We assume that when players are indifferent between two or more strategies, they randomize over these strategies.

The game is sequential, firstly each candidate i chooses (s_i, k_i), then DM chooses $(d, \boldsymbol{\omega})$.

The utility of candidate i depends on which task DM assigns to him and on how many candidates of his type and of the other choose his own social group. Let us define ϕ_m the payoff earned by a candidate from being assigned to task m. We assume that $\phi_{\alpha_L} > \phi_{\beta_L}$: such an assumption introduces a misalignment in preferences between candidates and DM, such that all candidates prefer to signal themselves as A types. We introduce this assumption to avoid cheap talk: this assumption can be justified by assuming that one of the two tasks pays a higher wage than the other (we discuss the implications of alternative assumptions in Section 5.4). We assume that $\phi_{\alpha_H} > \phi_{\alpha_L}$, and $\phi_{\beta_H} > \phi_{\beta_L}$: candidates prefer to be assigned to high-skilled tasks, because of higher job satisfaction for example.

[1] In the case of ω_{d0}, the domain of the function is the co-domain of the distribution of social group choices of candidates, while in the case of ω_{d1} it is the co-domain of the distributions of skill choices and types of candidates.

We name $\phi_{\alpha_L} - \phi_{\beta_L} = \underline{\phi}^m$. Where $\underline{\phi}^m$ determines the rigidity of candidates towards their favorite task. High levels of $\underline{\phi}^m$ denote strong rigidity of candidates, while low levels of $\underline{\phi}^m$ denote high flexibility (trivially, when $\underline{\phi}^m = 0$, candidates are indifferent between doing the two tasks).[2]

Assumption 3. $\phi_{\alpha_H} - \phi_{\alpha_L} = \phi_{\beta_H} - \phi_{\beta_L} = \overline{\phi}^s$.

$\overline{\phi}^s$ represents candidates' benefit from investing in their human capital. We assume the difference between α_H and α_L and the one between β_H and β_L to be equal for simplicity.

Assumption 4. $\overline{\phi}^s > \zeta$.

To keep our result as clear as possible, we assume that whenever a candidate invests in his skill and is assigned to a high-skilled task, the prize is worth the cost.[3]

We define $U_i(t_i, k_i, s_i, K_{-i}, S_{-i}, d, \boldsymbol{\omega})$ as the utility of candidate i, where K_{-i} is the distribution of the social group choices by every candidate that is not i and S_{-i} is the distribution of the skills choices by every candidate that is not i. Consider candidate i of type t playing strategy (s, k):

$$U_i(t, k, s, K_{-i}, d, \boldsymbol{\omega}) = MU_i(t, k, s, d, \boldsymbol{\omega}) + SU_i(t, k, K_{-i}). \tag{5.1}$$

We define $MU_i(t, k, s, d, \boldsymbol{\omega})$ as the material utility of candidate i. Such utility depends on the tasks a candidate is assigned to.

$$\begin{aligned} MU_i(t, k, s, d, \boldsymbol{\omega}) = {} & p\,(A, H|k; t, s, d)\,\phi_{\alpha_H} + p\,(A, L|k; t, s, d)\,\phi_{\alpha_L} + \\ & p\,(B, H|k; t, s, d)\,\phi_{\beta_H} + p\,(B, L|k; t, s, d)\,\phi_{\beta_L} - \zeta \mathbb{1}_i(s = H), \end{aligned} \tag{5.2}$$

where that $\mathbb{1}_i(s = H)$ is an indicator function that takes value 1 if and only if the candidate chooses H, and $p\,(t, s|k; t, s, d)$ is the joint posterior belief that a candidate is of type t and with skill s. $SU_i(t, k, K_{-i})$ is the social utility of candidate i. Such a utility depends on candidate's i social group choice, on his type, and on other candidates' social group choices, but it does not depend on DM choices. Before writing this kind of utility, we need to define

[2] When candidates are assigned to the high-skilled task, DM always buys the information, and candidates cannot choose between α_H and β_H. Therefore, we do not need to define $\phi_{\alpha_H} - \phi_{\beta_H} = \overline{\phi}^m$.

[3] Under this assumption, whenever DM buys the information, candidates invest in their human capital. If $\overline{\phi}^s < \zeta$, there could be an equilibrium where DM buys the information, and candidates don't invest in their human capital.

n_k^t and $n_k^{t'}$, i.e. the proportion of type $t \in A, B$ and $t' \neq t \in A, B$ in social group k.

$$SU_i(t, k, K_{-i}) = \eta f\left(n_k^{t'}, n_k^t\right). \tag{5.3}$$

The parameter $\eta > 0$ measures the degree of social preferences; $f\left(n_k^{t'}, n_k^t\right)$ assumes different forms depending on whether there is homophily or not.

Definition 6. *If* $f\left(n_k^{t'}, n_k^t\right) = \begin{cases} 1 & \text{if } n_k^{t'} + n_k^t \neq 0 \\ 0 & \text{if } n_k^{t'} + n_k^t = 0 \end{cases}$ *, there is no homophily.*

Definition 7. *If* $f\left(n_k^{t'}, n_k^t\right) = \begin{cases} \frac{n_k^t}{n_k^{t'} + n_k^t} & \text{if } n_k^{t'} + n_k^t \neq 0 \\ 0 & \text{if } n_k^{t'} + n_k^t = 0 \end{cases}$ *, there is homophily.*

In both cases, candidates prefer to choose a social group with other candidates rather than one with no candidate ($f\left(n_k^{t'}, n_k^t\right) = 0,$ if $n_k^{t'} + n_k^t = 0$). If there is homophily, the social utility of a candidate positively depends on how many candidates of the same type choose his same social group ($f\left(n_k^{t'}, n_k^t\right) = \frac{n_k^t}{n_k^{t'} + n_k^t},$ if $n_k^{t'} + n_k^t \neq 0$). If there is no homophily, the social utility of a candidate positively depends on how many candidates choose the same social group ($f\left(n_k^{t'}, n_k^t\right) = 1,$ if $n_k^{t'} + n_k^t \neq 0$).

We name m_i the task that DM assigns to candidate i. The utility for DM from assigning m_i to i is a function $V:\ \{A, B\} \times \{L, H\} \times \{d_0, d_1\} \times \{\alpha_L, \alpha_H, \beta_L, \beta_H\} \to \mathbb{R}$. Note that m_i depends on $\omega(K, S, T, d)$ which determines the decisions of DM based on her information's decision. With an abuse of notation we will write v as a function of m_i, which depends on $\omega(K, S, T, d)$. DM earns positive utility if she correctly matches candidates and tasks. A types should be assigned to α tasks, and B types should be assigned to β tasks, and trivially, candidates with H (L) skill should be assigned to high (low) skilled tasks.

$$v(t_i, s_i, d, m_i) = \mathbb{1}_t \tau + \mathbb{1}_{tH} \tau \delta_H + \mathbb{1}_{tL} \tau \delta_L - \mathbb{1}_d c, \tag{5.4}$$

where $\tau, \delta_H, \delta_L > 0$ and $\delta_H > \delta_L$. More precisely,

$$\mathbb{1}_t = \begin{cases} 1 & \text{if } ((m_i = \alpha_H \vee m_i = \alpha_L) \wedge t_i = A) \vee ((m_i = \beta_H \vee m_i = \beta_L) \wedge t_i = B) \\ 0 & \text{otherwise} \end{cases},$$

$$\mathbb{1}_{tL} = \begin{cases} 1 & \text{if } (m_i = \alpha_L \wedge t_i = A \wedge s_i = L) \vee (m_i = \beta_L \wedge t_i = B \wedge s_i = L) \\ 0 & \text{otherwise} \end{cases},$$

$$\mathbb{1}_d = \begin{cases} 1 & \text{if } d = d_1 \\ 0 & \text{otherwise} \end{cases} .$$

$\mathbb{1}_{tH}$ depends on weather there is substitutability between skills and types or complementarity between them.

Definition 8. *There is substitutability between skills and types if*

$$\mathbb{1}_{tH} = \begin{cases} 1 & \textit{if } ((m_i = \alpha_H \vee m_i = \beta_H) \wedge s_i = H) \\ 0 & \textit{otherwise} \end{cases} .$$

Definition 9. *There is complementarity between skills and types if*

$$\mathbb{1}_{tH} = \begin{cases} 1 & \textit{if } ((m_i = \alpha_H \wedge t_i = A \wedge s_i = H) \vee ((m_i = \beta_H \wedge t_i = B \wedge s_i = H) \\ 0 & \textit{otherwise} \end{cases} .$$

Note that DM earns a payoff for each candidate she allocates according to Equation (5.4), her total utility is the sum of all the payoffs she earns for each candidate. Therefore, the total utility of DM is $V(T, S, d, \mathbf{m}) = \int_0^1 v(t_i, s_i, d, m_i)di$.

The interpretation of the above utility is as follows. If the decision-maker does not guess either a candidate's type or skill correctly, she earns 0 utility. In case DM correctly guesses a candidate's skill but not his type, and the candidate has skill L, the decision-maker earns 0 (i.e. $\mathbb{1}(tsL) = 0$ in case DM assigns β_L to an A type with skill L). The decision-maker always earns a positive payoff τ if she correctly guesses a candidate's type. Moreover, if she correctly guesses a candidate's skill but not his type, and the candidate has skill H, the decision-maker earns $\tau\delta_H$ if skills and types are substitutes (Definition 8), or she earns 0 if skills and types are complements (Definition 9). She earns $\tau(1 + \delta_H)$ if she correctly guesses a candidate's type and skill, and the candidate has skill H, and she earns $\tau(1 + \delta_L)$ if she correctly guesses a candidate's type and skill, and the candidate has skill L.

Given that we are studying a context where the type of ability matters, we think that it is reasonable to reward the decision-maker whenever she matches a task with the correct type of candidate. If the decision-maker does not correctly guess a candidate's type, she always earns 0 if candidates have skill L. However, she may earn a positive payoff if candidates have skill H and there is substitutability between skills and types, i.e. if candidates with skill H excel in both types of abilities. She earns 0 if there is complementarity between skills and types, i.e. a candidate with skill H only excels in one type of ability. This can be the case in competitive sectors where workers' specialization matters the most.

Situation	DM utility
Task assignment does not match type or skills	0
Task assignment matches L skill but not type	0
Task assignment matches H skill but not type	$\tau\delta_H$ (t and s substitutes) 0 (t and s complements)
Task assignment matches type but not skills	τ
Task assignment matches H skill and type	$\tau(1+\delta_H)$
Task assignment matches L skill and type	$\tau(1+\delta_L)$

Table 5.1: Recap of DM's utility.

Parameter	Meaning
ζ	Cost of playing H
c	Information's cost
$\phi_{m'}$	Candidates' Utility from m' task
$\underline{\phi}^m$	Candidates' rigidity towards α_L
$\overline{\overline{\phi}}^s$	Premium from high-skilled tasks
η	Intensity of Social preferences
τ	DM premium for matching types and tasks
δ_H	DM multiplier for H skills candidates
δ_L	DM multiplier for matching L skills candidates

Table 5.2: Parameters of the model.

5.3 Results

Firstly, we use the *Perfect Bayesian-Nash Equilibrium* (PBE from now on) as the solution concept, and secondly, we use the *Coalition-Proof Nash Equilibrium* as a refinement (CPE from now on). For a formal definition, see Bernheim et al. (1987). The structure of our model allows us to restrict the attention to coalitions between candidates (see Battaglini, 2002 for a similar approach). In short, the decision-maker never has the incentive to deviate from an equilibrium where she buys the information, and given the misalignment of preferences between the candidates and the decision-maker, it cannot exist any self enforcing coalition between them if the decision-maker does not buy the information.

We provide the results in pure strategies. In Appendix D.2, we show that the intuition of our results still holds including mixed strategies. We begin by giving the results for the benchmark case.

5.3.1 Benchmark case

In this section, we consider $f\left(n_k^{t'}, n_k^t\right)$ as indicated in Definition 6.

Theorem 8.

Consider a game where there is not homophily, and types and skills are either complement or substitute.

- *If $c < p_B\tau$, there only exist pooling equilibria on social groups, where all candidates play H, and DM buys the information. These equilibria are CPE.*
- *If $c > p_B\tau(1+\delta_H)$, there only exist pooling equilibria on social groups, where all candidates play L, DM does not buy the information, and she assigns α_L to each candidate. These equilibria are CPE.*

From Theorem 8, we conclude that the absence of homophily is beneficial for both candidates and the decision-maker provided that the information's cost is low enough. However, the decision-maker faces a trade-off: she hires candidates with high skills, but on the other hand, she pays a cost to know their types and skills. When the information's cost is high enough, the absence of homophily is adverse to the decision-maker since she does not allocate B types correctly.

For simplicity, we only show results where equilibria are unique; a detailed map of what happens for every level of c can be found in Appendix D.1 (Corollary 11). Despite some quantitative differences between complementarity and substitutability, the qualitative message stays the same. We prove the theorem in Appendix D.1 and give an intuition here.

If there is no homophily, there can only be pooling equilibria on social groups. Indeed, candidates maximize their social utility by choosing whichever social group with at least a positive mass of candidates, and they maximize their material utility by signaling themselves as A types. As a consequence, social groups are not informative about candidates' types in equilibrium. Additionally, due to Assumption 3 and 4, all candidates chose the same skill, and DM always knows candidates' skills in equilibrium.

In the light of these considerations, we are only interested in two situations. In the first, the decision-maker buys the information, candidates invest in their human capital, all A types are assigned to α_H, and all B types are assigned to β_H. Such a situation is the unique possible one if and only if $c < p_B\tau$. In the second, the decision-maker does not buy the information, candidates do not invest in their human capital and all candidates are assigned to α_L. This situation is the unique possible one if and only if $c > p_B\tau(1+\delta_H)$.

5.3.2 Homophily

In this section, we consider $f\left(n_k^{t'}, n_k^t\right)$ as indicated in Definition 7.

Theorem 9.

Consider a game where there is homophily, and types and skills are either complement or substitute.

- *If $c < p_B\tau$, there exist infinite pooling equilibria on social groups, where all candidates play H, and DM buys the information. These equilibria are not CPE.*
- *If $c > p_B\tau(1+\delta_H)$, there exist infinite pooling equilibria on social groups, where all candidates play L, DM does not buy the information, and $\mathbf{m}^*_{\mathbf{k}} = (\alpha_L, \alpha_L)$. These equilibria are not CPE.*
- *If $c > p_B\tau(1+\delta_H)$, there exists one pooling equilibria on x (and one on y), where all candidates play L, DM does not buy the information, and $\mathbf{m}^*_{\mathbf{k}} = (\alpha_L, \beta_L)$ ($\mathbf{m}^*_{\mathbf{k}} = (\beta_L, \alpha_L)$). These equilibria are CPE if and only if $\eta < \frac{1}{p_A}\underline{\phi}^m$.*
- *$\forall c > 0$, if $\eta > \underline{\phi}^m$, there exist two separating equilibria on social groups, where all candidates play L, DM does not buy the information, and she assigns α_L to A types and β_L to B types. These equilibria are CPE if and only if $\eta > \frac{p_A+\epsilon p_B}{p_A}\underline{\phi}^m$.*

We depict all the possible equilibria as a function of η in Figure 5.1, and 5.2 to clarify this result. We say that there is *strong homophily* when separating equilibria on social groups are the unique CPE equilibria; *weak homophily* represents the case when only pooling equilibria on social groups are CPE.

From Theorem 9, we conclude that *strong homophily* helps the decision-maker in taking her decisions. Indeed, under *strong homophily*, she always knows the candidates' types and skills in equilibrium without buying the information. However, due to this selection mechanism, candidates never invest in their human capital. This result is detrimental for the candidates and for the decision-maker: for candidates since they are not assigned to high-skilled tasks and for the decision-maker because she cannot hire candidates with high skill. Interestingly, we observe pooling equilibria on social groups under *weak homophily*. However, none of them lead to CPE with skill acquisition from candidates. This is one of the most important results of the paper: although we observe pooling equilibria on social groups even under homophily, no equilibrium with candidates' skill acquisition is CPE under homophily, whatever the intensity of social preferences is.

As for Theorem 8, there are other possible pooling equilibria on social groups for c taking middle values (see Corollary 12 in Appendix D.1). None of them is CPE under *strong homophily*; therefore, we omit them in the main paper. We prove the theorem in Appendix D.1 and give an intuition here.

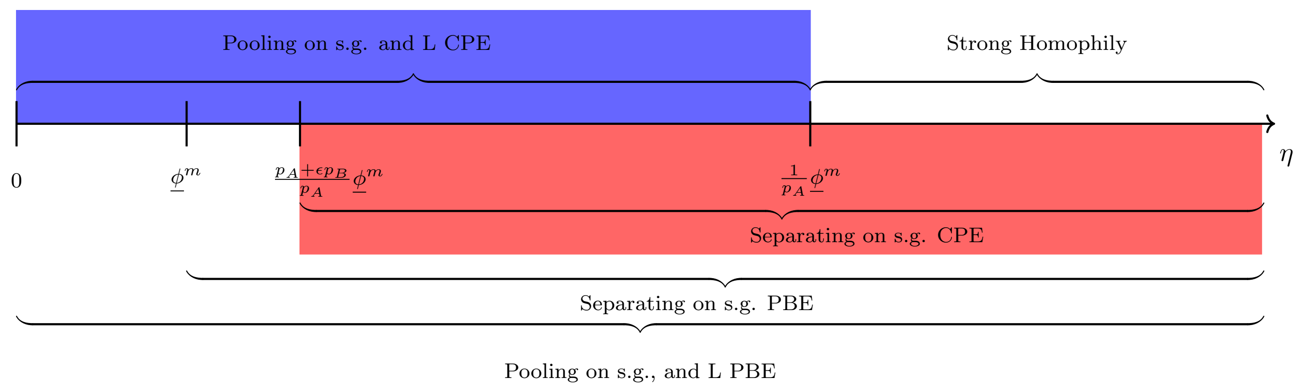

Figure 5.1: Equilibria under Homophily as a function of η, for $c > p_B\tau(1+\delta_H)$.

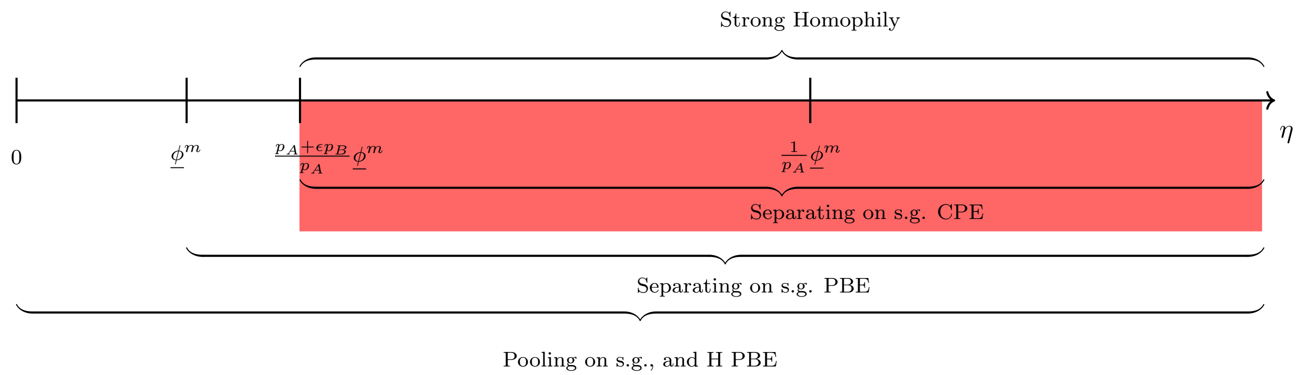

Figure 5.2: Equilibria under Homophily as a function of η, for $c < p_B\tau$.

The results in Theorem 9 can be interpreted as follows. Under homophily, there is an additional conflict of interest between the candidates and the decision-maker: all candidates want to be assigned to α_L, but the decision-maker wants to assign β_L to B types. Therefore, the existence of pooling and separating equilibria on social groups depends on the trade-off between the social and the material utility of candidates. The uniqueness of the separating equilibria is attained by the intensity of social preferences.

If $\eta > \underline{\phi}^m$, separating equilibria on social groups exist for all $c > 0$. Let us take an example where all A types choose x, and all B types choose y. In this case, DM has no uncertainty on candidates' types and skills, and she will never buy the information. Moreover, she will assign all candidates of social group x to α tasks and all candidates of social group y to β tasks. Given this scenario, two situations can happen. In the first, candidates invest in their human capital, all candidates with social group x are assigned to α_H, and all candidates with social group y are assigned to β_H. In the second, candidates do not invest in their human capital, all candidates with social group x are assigned to α_L, and all candidates with social group y are assigned to β_L. Given that DM plays d_0, only the second situation is possible in equilibrium. In this case, each B type has a trade-off between choosing x and y. On the one hand, if he chooses x, he earns a material utility of ϕ_{α_L} but a social utility of 0. On the other hand, if he chooses y, he earns a material utility of ϕ_{β_L} but a social utility of η. Therefore, all B types choose y if and only if $\eta > \phi_{\alpha_L} - \phi_{\beta_L} = \underline{\phi}^m$.

Our results show how homophily could lead employees to non-invest in their skills, resulting in a sub-optimal equilibrium for society. Homophily has been historically used to explain the persistence or the rising of inequalities. However, other mechanisms describe homophily as a potential source of inefficiency in productivity Bolte et al. (2020); Jackson (2021). These works show that homophily could trigger inefficient productivity through unequally spread referrals. We show that homophily could trigger inefficient equilibria through the information conveyed to the decision-maker (employer). Similar to other works in the homophily literature, the results in our paper comes from the separation of the two types into two different social groups. As shown in previous works, due to this separation (that may be seen as segregation), inequalities can arise if one of the two groups is somehow disadvantaged (Austen-Smith and Fryer Jr, 2005; Kim and Loury, 2012; Bowles et al., 2014; Kim and Loury, 2019). We show how the separation into social groups can lead the system to a sub-optimal equilibrium for society, not driven by inequalities but by the under-investment in the skills of employees. Another source of inequalities through (homophilous) group formation is access to valuable information Golub and Jackson (2012); Lobel and Sadler (2016); Aybas and Jackson (2021). According to these models, people tend to listen more to their peers, i.e. their connections, and due to homophily, their connections can be biased and lead them towards bad choices. Therefore, inequalities may emerge between two social groups because one group has better sources of information (e.g. on school choices). Compared to these works, we show that group formation can be detrimental to society due to the amount of information delivered to employers.

5.3.3 Welfare

In this section, we compare the social welfare under *strong homophily* and the benchmark. We made this decision since our paper concerns situations in which social preferences matter compared to material ones, and only under *strong homophily* does this happen. For the sake of our comparison, we keep omitting middle values of c. This decision is coherent with the one we took for theorems. As it is not clear whether homophily corresponds to preferring the company of similar types or disliking the company of dissimilar ones, we also omit from the welfare analysis the social utility component.

In Figure 5.3 and 5.4, we depict the different human capital investments under the benchmark case and the *strong homophily* one if there is substitutability and complementarity. As we can see from these pictures, candidates never invest in their human capital under homophily, while they invest in their human capital under the benchmark case, provided that the cost is low enough.

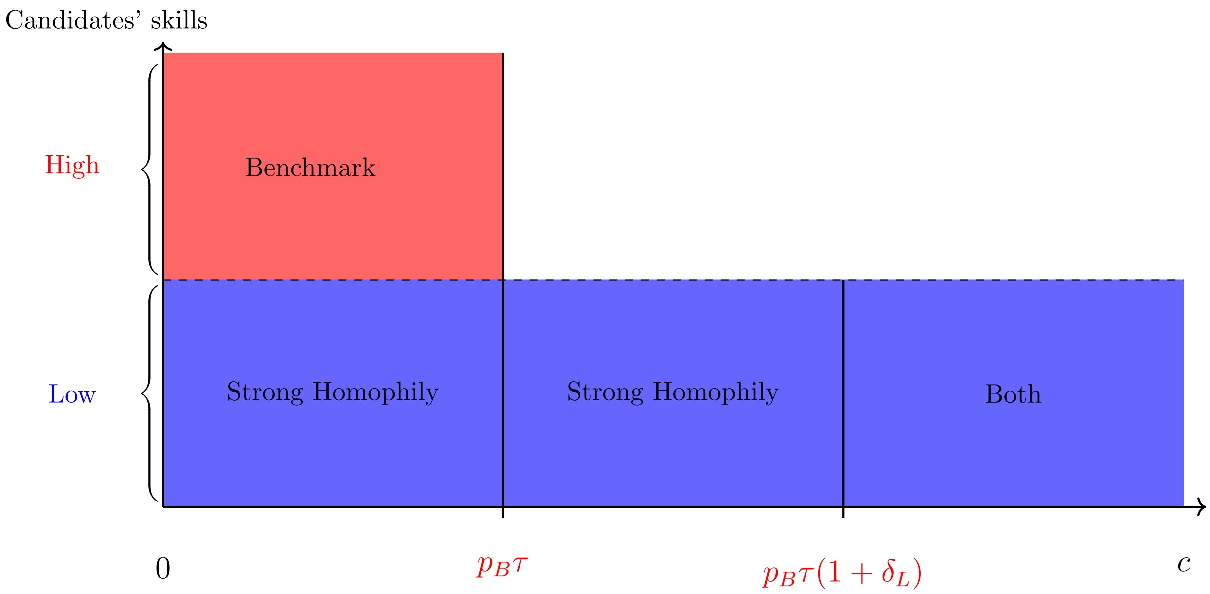

Figure 5.3: Cost's effect on human capital investments (under substitutability between skills and types).

Firstly, we consider candidates' welfare. We call W^{ben}_{can} candidates welfare under the benchmark and W^{hom}_{can} the one under *strong homophily*.

If $c < p_B\tau$,

$$W^{ben}_{can} = p_A\phi_{\alpha_H} + p_B\phi_{\beta_H} - \zeta,$$

$$W^{hom}_{can} = p_A\phi_{\alpha_L} + p_B\phi_{\beta_L}.$$

Clearly, $W^{ben}_{can} > W^{hom}_{can}$. The intuition behind this result is simple: if the information's cost is low enough, candidates always invest in their human capital under the benchmark, and consequently, their welfare is higher under this case than under *strong homophily*.

If $c > p_B\tau(1+\delta_H)$,

$$W^{ben}_{can} = \phi_{\alpha_L},$$

$$W^{hom}_{can} = p_A\phi_{\alpha_L} + p_B\phi_{\beta_L}.$$

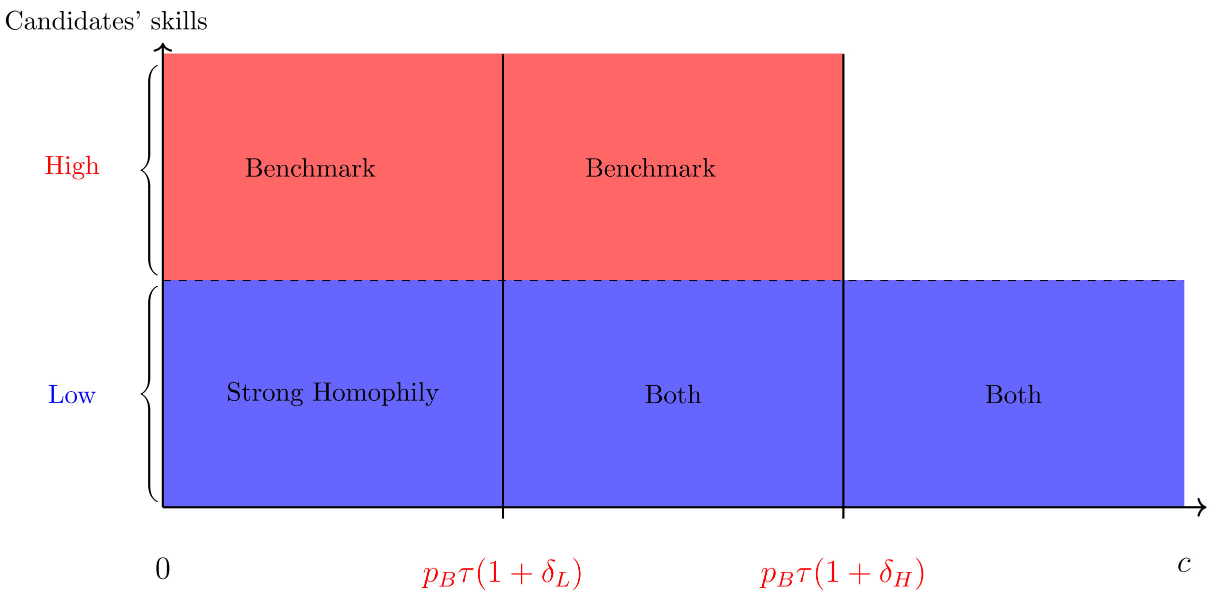

Figure 5.4: Cost's effect on human capital investments (under complementarity between skills and types).

Again, $W^{ben}_{can} > W^{hom}_{can}$. However, this result is arbitrary since candidates' favorite tasks coincide with the task that better fits with the majority type.

Secondly, we consider the decision-maker's welfare. We call W^{ben}_{DM} the decision-maker's welfare under the benchmark and W^{hom}_{DM} the one under *strong homophily*.

If $c < p_B\tau$,

$$W^{ben}_{DM} = \tau(1+\delta_H) - c,$$

$$W^{hom}_{DM} = \tau(1+\delta_L).$$

Therefore, $W^{ben}_{DM} > W^{hom}_{DM}$ if and only if $c < \delta_H - \delta_L$. The interpretation is straightforward. Under the benchmark, the decision-maker hires candidates with high skill, but she pays for the information, while under *strong homophily*, she does not buy the information, but she does not hire candidates with high skill. Consequently, if the cost of buying the information is paid off by the benefit of having candidates with high skill, the decision-maker's welfare is higher under the benchmark.

If $c > p_B\tau(1+\delta_H)$,

$$W^{ben}_{DM} = p_A\tau(1+\delta_L),$$

$$W_{DM}^{hom} = \tau(1 + \delta_L).$$

Hence, $W_{DM}^{ben} < W_{DM}^{hom}$. In this case, the decision-maker does not buy the information under the benchmark, and the information she has without the purchasing is higher under *strong homophily.* Consequently, the decision-maker's welfare is always higher under *strong homophily* if the information's cost is high enough.

Lastly, we consider the total social welfare as the sum of the candidates' welfare and the decision-maker's one. We call W^{ben} the total social welfare under the benchmark and W^{hom} the one under *strong homophily.*

If $c < p_B\tau$,

$$W^{ben} = \tau(1 + \delta_H) - c + p_A\phi_{\alpha_H} + p_B\phi_{\beta_H} - \zeta,$$
$$W^{hom} = \tau(1 + \delta_L) + p_A\phi_{\alpha_L} + p_B\phi_{\beta_L}.$$

$W^{ben} > W^{hom}$ if and only if $c + \zeta < \overline{\phi}^s + \delta_H - \delta_L$. In other words, if the social costs from information acquisition and human capital investment do not overwhelm the social benefits from having high candidates with high skills, the total social welfare is better under the benchmark when the cost is sufficiently low.

If $c > p_B\tau(1 + \delta_H)$,

$$W^{ben} = p_A\tau(1 + \delta_L)\phi_{\alpha_L},$$
$$W^{hom} = \tau(1 + \delta_L) + p_A\phi_{\alpha_L} + p_B\phi_{\beta_L}.$$

$W^{ben} > W^{hom}$ if and only if $\tau(1 + \delta_L) < \underline{\phi}^m$. We can interpret this inequality in the following way. When the cost is sufficiently high, there are no (direct) social costs since the decision-maker does not buy the information, and candidates do not invest in their human capital. Moreover, we know for sure that candidates' welfare is always better under the benchmark, while the decision-maker's welfare is always better under homophily. Hence, if the candidates' benefits are higher than the decision-maker's one, the total social welfare is higher under the benchmark and vice-versa.

5.4 Discussion

In this section we discuss some points concerning alternative assumptions, and follow-up analysis.

5.4.1 Heterophily and intrinsic values for social groups

In our model, we do not consider heterophily, i.e. the love for the different, since it would give the same result as the benchmark. Nevertheless, heterophily could give interesting results when incorporated with more realistic assumptions, such as introducing an intrinsic value for the social groups. Indeed, in our model, cognitive types perceive the same utility in joining nerds or cool guys a priori. However, the opposite may be true in the real world.

Introducing this assumption may eradicate all the results under the benchmark but reinforce those under homophily since an intrinsic value for social groups would accentuate candidates' desire to polarize. Under homophily, we would observe separating equilibria on social groups for even lower values of η. Under no homophily, even a little intrinsic value for social groups would destroy all equilibria where the decision-maker buys the information and candidates invest in their human capital. The intuition is as follows. If the decision-maker buys the information, candidates earn the same material utility no matter which social group they choose. They also earn the same social utility no matter which social group they choose, provided that there is at least one person in each social group (see Definition 6). Therefore, even a little intrinsic value for social groups would make candidates polarize in the two social groups they prefer, destroying all these pooling equilibria. Under heterophily, candidates would have material and social incentives to pool in the same group, and they would have a reason to split due to the different intrinsic preferences of the social groups. Therefore, sufficiently strong social preferences would sustain a pooling equilibrium on social groups.

5.4.2 Candidates with heterogeneous preferences

As we assume in Section 5.2, all candidates prefer α_L to β_L. We could introduce heterogeneity in candidates' preferences between types, meaning A types prefer α_L to β_L, but B types prefer β_L to α_L. In this case, when the decision-maker does not buy the information, each candidate could claim to be his own type, having no incentive to lie. In other words, candidates could send cheap signals rather than choosing social groups to signal their type to the decision-maker. As a result, the decision-maker would always prefer cheap talk signals to social groups to infer candidates' types, and there would no longer be a homophily effect.

Nevertheless, it would be sufficient to introduce a small heterogeneity in candidates' preferences to return to our results. As a matter of fact, if some candidates' preferences were heterogeneous within the same type (some candidates of the same type prefer α_L, and some β_L), cheap talk signals would be no longer effective, and the decision-maker would rely again on social groups to infer candidates' types and skills under homophily. Therefore, there would be again the homophily effect.

Intuitively, candidates who prefer α_L would claim to be A types no matter their real type, and candidates who prefer β_L would claim to be B types no matter their real type. The decision-maker could not rely on these cheap talk signals, and she would look at candidates' social groups as in our version. Therefore, our results would be preserved.

5.4.3 Alternative assumptions on the decision-maker

In our model, we assume that there are enough tasks for every candidate and that there is no difference between the number of high and low tasks. Secondly, we assume a unique decision-maker, ignoring possible competitors of the firm.

Concerning the second case, we can argue that our analysis applies to situations when there is only one big firm in the market (e.g., a monopoly). However, it is reasonable to think that part of our results would still hold even under competition between two decision-makers. A competition between decision-makers could raise the wages in equilibrium (e.g., no matter their beliefs, the two decision-makers assign all the high-skilled tasks in equilibrium), but it would not change the relevant mechanism behind our model. Separating or pooling equilibria on social groups would still happen depending on candidates' social preferences. Therefore, the two decision-makers would still face the same uncertainties, and they would buy the information under the same conditions as in our model.

Concerning the first case, it is arguably more reasonable to assume that the number of tasks is limited (specifically high-skilled ones). However, these assumptions would neither affect candidates' incentives to invest in their human capital nor their incentives to signal their type through social groups. Firstly, when the decision-maker does not buy the information, candidates would have no incentive to invest in their human capital, even in this case. Secondly, when the decision-maker buys the information, only a few candidates will be assigned to high-skilled tasks, but the only way for them to have a chance is to invest in their human capital. Since candidates cannot signal their types differently, the decision-maker would still face the same uncertainty levels under homophily or the benchmark. Therefore, our results would remain untouched by this assumption.

5.4.4 Testable Implications

To the best of our knowledge, we are the first to study this mechanism referred to the specific setting of our model. For this reason, there is a lack of empirical validation or examples as a reference in the literature. However, we believe that our hypothesis can be tested in different frameworks. Two main implications can be tested: the first concerns the role of homophily as a tool for decision-makers, while the second concerns the implied mechanism such that candidates should invest less in their human capital if the decision-makers invest fewer resources in the screening process.

Testable Implication 1. *A higher level of homophily correlates with a lower level of resources invested in the screening process of the decision-maker.*

Testable Implication 2. *A higher level of homophily correlates with a lower level of human capital investments of candidates.*

Whichever road is chosen to test our implications, there should be three key variables: one on homophily, one on the screening process (information acquisition), and one on human capital investments. Note that the underlying assumption beyond these implications is that the number of tasks coincides with the number of types of candidates.

Empirical data

Our hypothesis can be empirically validated through existing data or ad hoc questionnaires. Ideal data should come from big firms with different figures to fill (preferably, at least four, as in our model). The analysis can be done within the same firm, comparing two different departments or between two firms.

The degree of homophily among employees can be studied through variables concerning the employees' social networks. Alternatively, homophily can be studied through ad hoc questionnaires. Such questionnaires should be delivered to the employees and concern their relations inside or outside the workplace. Otherwise, homophily can be calculated by measuring the diversity within the different departments (for example, looking if there are recurrent characteristics among employees of each department). The variable regarding screening methods of the firm could be directly asked to the firm or captured from existing data: e.g., the amount of resources spent in interviews. The variable concerning human capital investment could concern, but may not be limited to, educational attainments. Some measures that can be used are grades at school, the number of degrees, or the number of refresher courses done by each employee.

Experimental investigation

Our implication could also be tested through an experiment: the aim of such an experiment should be to generate the variables described in the previous section to test Implications 1 and 2. There are two possible roads that the experimenter can adopt. The experiment could be designed in a way that subjects would play a simplified version of our model, with some subjects being the decision-maker and some subjects being candidates. Alternatively, the experiment could be designed in a way to test the interpretation of our results. In our paper, we give a specific interpretation of the mathematical structure, but any other interpretation that a possible experimenter could think about could be used as well. As an example, the pooling and the separating equilibria can be forced among subjects, and the experimenter could focus mainly on the decision-maker role.

The key variables that the experiment should generate are the same as in the previous section. Homophily should be the treatment, and it could either be generated endogenously or given exogenously. The easiest way to design this experiment is to use two treatments: one with no/low homophily and the second with high homophily. The easiest way to induce homophily is to reward subjects, as in our model, but any alternative way could be used as well. Endogenous homophily could give more relevance to the analysis but may be harder to induce.

5.4.5 Policy implications

Our results report a malfunction of the job assignment process when firms can see the social groups of potential employees. Such an issue is closely related to problems associated with genders or ethnicities. Both of them are visible to employers, and they are often causes of discrimination against employees. As a consequence of these similarities in the problems, there could also be similarities in the solutions. Although a consensus has not been reached yet, anonymous resumes are often used to dampen the effect of gender or ethnic discrimination (Goldin and Rouse, 2000; Behaghel et al., 2015; Derous and Ryan, 2019; Lacroux and Martin-Lacroux, 2020). These kinds of interventions can also be functional to reduce the negative effect of homophily found in our paper.

Nevertheless, implementing blind interviews or blind CV policies can be a redundant cost if there is no homophily. Therefore, firms need to know whether there is homophily between potential/current employees and how intense are the social preferences of employees. Among the many techniques that can be used to study the intensity of homophily, we find particularly suggesting the one in Currarini and Mengel (2016). In that experiment, subjects can choose with whom to interact between in-group or out-group members, and their willingness to play with an in-group member is also elicited through their willingness to pay. With similar methods, firms can elicit homophily and the intensity of homophily of their candidates/current employees.

A final consideration about these policies is their expensive nature, both in terms of resources and cognitive costs. In both cases, firms need to pay a cost that is reminiscent of an information cost; therefore, these policies are only implementable when the information's cost is considerably low. In such a case, the comparison between homophily and the benchmark case is still important, given the results in Section 5.3.3.

5.5 Conclusions

In this paper, we have shown a novel mechanism through which homophily can affect the labor markets. Overall, we have shown that homophily favors decision-makers because they can retrieve all the relevant information about candidates without buying it. However, this tool is a double-edged sword since candidates never invest in their human capital under homophily due to this effect. We also show that our results hold both under substitutability and complementarity between skills and types.

We started from the case where candidates have a social desire to aggregate indiscriminately from their types. In Theorem 8, we showed that effortlessly observable signals are no longer informative about the relevant pieces of information if this is the case. Therefore, the decision-maker buys the information (for low cost levels), and candidates invest in their human capital. Afterward, we consider the homophily case, where candidates have a social desire to aggregate only with candidates of the same type. In Theorem 9, we showed that the only reasonable equilibria under (strong) homophily are the ones where effortlessly observable signals (such as social group traits) are informative about the relevant pieces of information in the model. Due to this effect, the decision-maker never needs to buy the information under *strong homophily*. Consequently, candidates never invest in their human capital. Importantly, even though pooling equilibria on social groups are still possible under *weak homophily*, the pooling equilibria on social groups, where candidates invest in their skills and the decision-maker buys the information are never CPE under homophily, for any intensity of social preferences.

Bibliography

Agerström, J., Carlsson, R., Nicklasson, L., and Guntell, L. (2016). Using descriptive social norms to increase charitable giving: The power of local norms. *Journal of Economic Psychology*, 52:147–153.

Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics*, 95(9-10):1082–1095.

Allcott, H. and Kessler, J. B. (2019). The welfare effects of nudges: A case study of energy use social comparisons. *American Economic Journal: Applied Economics*, 11(1):236–76.

Alós-Ferrer, C. and Buckenmaier, J. (2020). Behavioral rules. In *Handbook of Experimental Game Theory*, pages 289–309. Edward Elgar Publishing.

Alós-Ferrer, C. and Ritschel, A. (2021). Multiple behavioral rules in cournot oligopolies. *Journal of Economic Behavior & Organization*, 183:250–267.

Andreoni, J., Nikiforakis, N., and Siegenthaler, S. (2021). Predicting social tipping and norm change in controlled experiments. *Proceedings of the National Academy of Sciences*, 118(16):e2014893118.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455.

Anufriev, M. and Hommes, C. (2012a). Evolution of market heuristics. *The Knowledge Engineering Review*, 27(2):255–271.

Anufriev, M. and Hommes, C. (2012b). Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(4):35–64.

Archer, J. (1988). *The behavioural biology of aggression*, volume 1. Cup Archive.

Argenziano, R., Severinov, S., and Squintani, F. (2016). Strategic information acquisition and transmission. *American Economic Journal: Microeconomics*, 8(3):119–55.

Arieli, I., Babichenko, Y., Peretz, R., and Young, H. P. (2020). The speed of innovation diffusion in social networks. *Econometrica*, 88(2):569–594.

Arigapudi, S., Heller, Y., and Schreiber, A. (2021). Sampling dynamics and stable mixing in hawk-dove games. *arXiv preprint arXiv:2107.08423*.

Arrow, K. J. (1973). *O. Ashenfelter and A. Rees (eds), Discrimination in Labor Markets*. Princeton University Press.

Arrow, K. J. and Borzekowski, R. (2004). Limited network connections and the distribution of wages. *FEDS working papers*.

Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. *Organizational Influence Processes*, 58:295–303.

Austen-Smith, D. and Fryer Jr, R. G. (2005). An economic analysis of “acting white”. *The Quarterly Journal of Economics*, 120(2):551–583.

Aybas, Y. C. and Jackson, M. O. (2021). Homophily in information networks: Sample-based herding and its implications for immobility and inequality. *mimeo Stanford*.

Ayres, I., Raseman, S., and Shih, A. (2013). Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. *The Journal of Law, Economics, and Organization*, 29(5):992–1022.

Baliga, S. and Sjöström, T. (2012). The strategy of manipulating conflict. *American Economic Review*, 102(6):2897–2922.

Baliga, S. and Sjöström, T. (2020). The strategy and technology of conflict. *Journal of Political Economy*, 128(8):3186–3219.

Barceló, J. and Sheen, G. C.-H. (2020). Voluntary adoption of social welfare-enhancing behavior: Mask-wearing in spain during the covid-19 outbreak. *PLoS one*, 15(12):e0242764.

Barile, J. P., Guerin, R. J., Fisher, K. A., Tian, L. H., Okun, A. H., Vanden Esschert, K. L., Jeffers, A., Gurbaxani, B. M., Thompson, W. W., and Prue, C. E. (2021). Theory-based behavioral predictors of self-reported use of face coverings in public settings during the covid-19 pandemic in the united states. *Annals of Behavioral Medicine*, 55(1):82–88.

Battaglini, M. (2002). Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401.

Beaman, L., Keleher, N., and Magruder, J. (2018). Do job networks disadvantage women? evidence from a recruitment experiment in malawi. *Journal of Labor Economics*, 36(1):121–157.

Becker, G. S. (2010). *The economics of discrimination*. University of Chicago press.

Bedenlier, S., Wunder, I., Gläser-Zikuda, M., Kammerl, R., Kopp, B., Ziegler, A., and Händel, M. (2021). "generation invisible?. higher education students'(non) use of webcams in synchronous online learning. *International Journal of Educational Research Open*, 2:100068.

Behaghel, L., Crépon, B., and Le Barbanchon, T. (2015). Unintended effects of anonymous resumes. *American Economic Journal: Applied Economics*, 7(3):1–27.

Bellemare, M. F., Novak, L., and Steinmetz, T. L. (2015). All in the family: Explaining the persistence of female genital cutting in west africa. *Journal of Development Economics*, 116:252–265.

Belloc, M. and Bowles, S. (2013). The persistence of inferior cultural-institutional conventions. *American Economic Review*, 103(3):93–98.

Benabou, R. (1996). Equity and efficiency in human capital investment: the local connection. *The Review of Economic Studies*, 63(2):237–264.

Benndorf, V. and Martinez-Martinez, I. (2017). Perturbed best response dynamics in a hawk–dove game. *Economics Letters*, 153:61–64.

Benndorf, V., Martinez-Martinez, I., and Normann, H.-T. (2016). Equilibrium selection with coupled populations in hawk–dove games: Theory and experiment in continuous time. *Journal of Economic Theory*, 165:472–486.

Berger, U. and De Silva, H. (2021). Evolution of deterrence with costly reputation information. *PLoS one*, 16(6):e0253344.

Bergin, J. and Lipman, B. L. (1996). Evolution with state-dependent mutations. *Econometrica*, 64(4):943–956.

Bernheim, B. D., Peleg, B., and Whinston, M. D. (1987). Coalition-proof nash equilibria i. concepts. *Journal of Economic Theory*, 42(1):1–12.

Bertrand, M. and Mullainathan, S. (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013.

Bhaskar, V. and Vega-Redondo, F. (2004). Migration and the evolution of conventions. *Journal of Economic Behavior & Organization*, 55(3):397–418.

Bicchieri, C. (2005). *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press.

Bicchieri, C., Jiang, T., and Lindemans, J. W. (2014). A social norms perspective on child marriage: The general framework.

Bilancini, E. and Boncinelli, L. (2018a). Rational attitude change by reference cues when information elaboration requires effort. *Journal of Economic Psychology*, 65:90–107.

Bilancini, E. and Boncinelli, L. (2018b). Signaling to analogical reasoners who can acquire costly information. *Games and Economic Behavior*, 110:50–57.

Bilancini, E. and Boncinelli, L. (2018c). Signaling with costly acquisition of signals. *Journal of Economic Behavior & Organization*, 145:141–150.

Bilancini, E. and Boncinelli, L. (2018d). Social coordination with locally observable types. *Economic Theory*, 65(4):975–1009.

Bilancini, E. and Boncinelli, L. (2020). The evolution of conventions under condition-dependent mistakes. *Economic Theory*, 69(2):497–521.

Bilancini, E. and Boncinelli, L. (2021). When market unraveling fails and mandatory disclosure backfires: Persuasion games with labeling and costly information acquisition. *Journal of Economics & Management Strategy*, 30(3):585–599.

Bilancini, E., Boncinelli, L., Ille, S., and Vicario, E. (2022). Memory retrieval and harshness of conflict in the hawk–dove game. *Economic Theory Bulletin*, 10(2):333–351.

Bilancini, E., Boncinelli, L., and Nax, H. H. (2021). What noise matters? experimental evidence for stochastic deviations in social norms. *Journal of Behavioral and Experimental Economics*, 90:101626.

Bilancini, E., Boncinelli, L., and Wu, J. (2018). The interplay of cultural intolerance and action-assortativity for the emergence of cooperation and homophily. *European Economic Review*, 102:1–18.

Binmore, K. and Samuelson, L. (1999). Evolutionary drift and equilibrium selection. *The Review of Economic Studies*, 66(2):363–393.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5(3):387–424.

Blume, L. E. (2003). How noise matters. *Games and Economic Behavior*, 44(2):251–271.

Bokemper, S. E., Cucciniello, M., Rotesi, T., Pin, P., Malik, A. A., Willebrand, K., Paintsil, E. E., Omer, S. B., Huber, G. A., and Melegaro, A. (2021). Experimental evidence that changing beliefs about mask efficacy and social norms increase mask wearing for covid-19 risk reduction: Results from the united states and italy. *PLoS one*, 16(10):e0258282.

Bolte, L., Immorlica, N., and Jackson, M. O. (2020). The role of referrals in inequality, immobility, and inefficiency in labor markets. *SSRN https://doi. org/10.2139/ssrn*, 3512293.

Bond, R. and Smith, P. B. (1996). Culture and conformity: A meta-analysis of studies using asch's (1952b, 1956) line judgment task. *Psychological Bulletin*, 119(1):111.

Borsari, B. and Carey, K. B. (2003). Descriptive and injunctive norms in college drinking: a meta-analytic integration. *Journal of Studies on Alcohol*, 64(3):331–341.

Bott, K. M., Cappelen, A. W., Sørensen, E. Ø., and Tungodden, B. (2020). You've got mail: A randomized field experiment on tax evasion. *Management Science*, 66(7):2801–2819.

Bowles, S., Loury, G. C., and Sethi, R. (2014). Group inequality. *Journal of the European Economic Association*, 12(1):129–152.

Bramoullé, Y. and Rogers, B. W. (2009). Diversity and popularity in social networks. *CIRPEE Working Paper*.

Brams, S. J. and Kilgour, D. M. (1987). Threat escalation and crisis stability: A game-theoretic analysis. *American Political Science Review*, 81(3):833–850.

Brams, S. J. and Kilgour, D. M. (2001). Fallback bargaining. *Group Decision and Negotiation*, 10(4):287–316.

Brock, W. A. and Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8-9):1235–1274.

Brock, W. A. and Hommes, C. H. (2001). A rational route to randomness. In *Growth Theory, Nonlinear Dynamics and Economic Modelling*, pages 402–438. Edward Elgar Publishing.

Bursztyn, L., Ederer, F., Ferman, B., and Yuchtman, N. (2014). Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica*, 82(4):1273–1301.

Butler, J. V., Guiso, L., and Jappelli, T. (2013). Manipulating reliance on intuition reduces risk and ambiguity aversion. *CEPR Discussion Paper No. DP9461*.

Caillaud, B. and Tirole, J. (2007). Consensus building: How to persuade a group. *American Economic Review*, 97(5):1877–1900.

Calvo-Armengol, A. and Jackson, M. O. (2004). The effects of social networks on employment and inequality. *American Economic Review*, 94(3):426–454.

Carvalho, J.-P. (2017). Coordination and culture. *Economic Theory*, 64(3):449–475.

Cason, T. N. and Mui, V.-L. (1998). Social influence in the sequential dictator game. *Journal of Mathematical Psychology*, 42(2-3):248–265.

Castelli, F. R. and Sarvary, M. A. (2021). Why students do not turn on their video cameras during online classes and an equitable and inclusive plan to encourage them to do so. *Ecology and Evolution*, 11(8):3565–3576.

Centola, D. and Baronchelli, A. (2015). The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*, 112(7):1989–1994.

Centola, D., Becker, J., Brackbill, D., and Baronchelli, A. (2018). Experimental evidence for tipping points in social convention. *Science*, 360(6393):1116–1119.

Chetty, R., Friedman, J. N., Saez, E., Turner, N., and Yagan, D. (2020). Income segregation and intergenerational mobility across colleges in the united states. *The Quarterly Journal of Economics*, 135(3):1567–1633.

Chiarella, C. and He, X.-Z. (2002). Heterogeneous beliefs, risk and learning in a simple asset pricing model. *Computational Economics*, 19(1):95–132.

Chiarella, C. and He, X.-Z. (2003). Dynamics of beliefs and learning under al-processes—the heterogeneous case. *Journal of Economic Dynamics and Control*, 27(3):503–531.

Cialdini, R. B. (2007). Descriptive social norms as underappreciated sources of social control. *Psychometrika*, 72(2):263–268.

Cialdini, R. B., Reno, R. R., and Kallgren, C. A. (1990). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology*, 58(6):1015.

Clauset, A., Arbesman, S., and Larremore, D. B. (2015). Systematic inequality and hierarchy in faculty hiring networks. *Science Advances*, 1(1):e1400005.

Coate, S. and Loury, G. C. (1993). Will affirmative-action policies eliminate negative stereotypes? *American Economic Review*, 83(5):1220–1240.

Cortez, M. H. and Ellner, S. P. (2010). Understanding rapid evolution in predator-prey interactions using the theory of fast-slow dynamical systems. *The American Naturalist*, 176(5):E109–E127.

Costa, D. L. and Kahn, M. E. (2013). Energy conservation "nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment. *Journal of the European Economic Association*, 11(3):680–702.

Cressman, R., Ansell, C., and Binmore, K. (2003). *Evolutionary dynamics and extensive form games*, volume 5. MIT Press.

Cressman, R., Garay, J., and Hofbauer, J. (2001). Evolutionary stability concepts for n–species frequency-dependent interactions. *Journal of Theoretical Biology*, 211(1):1–10.

Currarini, S., Jackson, M. O., and Pin, P. (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77(4):1003–1045.

Currarini, S. and Mengel, F. (2016). Identity, homophily and in-group bias. *European Economic Review*, 90:40–55.

Damm, A. P. and Dustmann, C. (2014). Does growing up in a high crime neighborhood affect youth criminal behavior? *American Economic Review*, 104(6):1806–32.

Dekel, E., Ely, J. C., and Yilankaya, O. (2007). Evolution of preferences. *The Review of Economic Studies*, 74(3):685–704.

Derous, E. and Ryan, A. M. (2019). When your resume is (not) turning you down: Modelling ethnic bias in resume screening. *Human Resource Management Journal*, 29(2):113–130.

Desmet, K. and Wacziarg, R. (2018). The cultural divide. Technical report, National Bureau of Economic Research.

Dewatripont, M. (2006). Costly communication and incentives. *Journal of the European Economic Association*, 4(2-3):253–268.

Dewatripont, M. and Tirole, J. (2005). Modes of communication. *Journal of Political Economy*, 113(6):1217–1238.

Di Pei, H. (2015). Communication with endogenous information acquisition. *Journal of Economic Theory*, 160:132–149.

Dindo, P. and Tuinstra, J. (2011). A class of evolutionary models for participation games with negative feedback. *Computational Economics*, 37(3):267–300.

Dixit, A., McAdams, D., and Skeath, S. (2019). 'we haven't got but one more day': The cuban missile crisis as a dynamic chicken game. *Available at SSRN 3406265*.

Dodds, P. and Watts, D. J. (2011). Threshold models of social influence. *The Oxford Handbook of Analytical Sociology*.

Dokumacı, E. and Sandholm, W. H. (2011). Large deviations and multinomial probit choice. *Journal of Economic Theory*, 146(5):2151–2158.

Dridi, S. and Lehmann, L. (2015). A model for the evolution of reinforcement learning in fluctuating games. *Animal Behaviour*, 104:87–114.

Durlauf, S. N. (1996). A theory of persistent income inequality. *Journal of Economic Growth*, 1(1):75–93.

Edo, A., Jacquemet, N., and Yannelis, C. (2019). Language skills and homophilous hiring discrimination: Evidence from gender and racially differentiated applications. *Review of Economics of the Household*, 17(1):349–376.

Efferson, C. and Vogt, S. (2018). Behavioural homogenization with spillovers in a normative domain. *Proceedings of the Royal Society B: Biological Sciences*, 285(1879):20180492.

Efferson, C., Vogt, S., Elhadi, A., Ahmed, H. E. F., and Fehr, E. (2015). Female genital cutting is not a social coordination norm. *Science*, 349(6255):1446–1447.

Efferson, C., Vogt, S., and Fehr, E. (2020). The promise and the peril of using social influence to reverse harmful traditions. *Nature Human Behaviour*, 4(1):55–68.

Ellison, G. (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *The Review of Economic Studies*, 67(1):17–45.

Ellison, G., Fudenberg, D., and Imhof, L. A. (2016). Fast convergence in evolutionary models: A lyapunov approach. *Journal of Economic Theory*, 161:1–36.

Fang, H. (2001). Social culture and economic performance. *American Economic Review*, 91(4):924–937.

Farrow, K., Grolleau, G., and Ibanez, L. (2017). Social norms and pro-environmental behavior: A review of the evidence. *Ecological Economics*, 140:1–13.

Ferraro, P. J. and Price, M. K. (2013). Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment. *Review of Economics and Statistics*, 95(1):64–73.

Foley, M., Forber, P., Smead, R., and Riedl, C. (2018). Conflict and convention in dynamic networks. *Journal of the Royal Society Interface*, 15(140):20170835.

Foley, M., Smead, R., Forber, P., and Riedl, C. (2021). Avoiding the bullies: The resilience of cooperation among unequals. *PLoS computational biology*, 17(4):e1008847.

Fortin, B., Lacroix, G., and Villeval, M.-C. (2007). Tax evasion and social interactions. *Journal of Public Economics*, 91(11-12):2089–2112.

Fosgerau, M., Sethi, R., and Weibull, J. W. (2020). Categorical screening with rational inattention. *Available at SSRN*.

Foster, D. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38(2):219–232.

Freidin, E., Acera Martini, L., Senci, C. M., Duarte, C., and Carballo, F. (2022). Field observations and survey evidence to assess predictors of mask wearing across different outdoor activities in an argentine city during the covid-19 pandemic. *Applied Psychology: Health and Well-Being*, 14(1):81–100.

Freidlin, M. I. and Wentzell, A. D. (1998). *Random perturbations of dynamical systems*. Springer.

Frey, B. S. and Meier, S. (2004). Social comparisons and pro-social behavior: Testing" conditional cooperation" in a field experiment. *American Economic Review*, 94(5):1717–1722.

Friedman, D. (1991). Evolutionary games in economics. *Econometrica*, 59(3):637–666.

Friedman, D. (1996). Equilibrium in evolutionary games: Some experimental results. *The Economic Journal*, 106(434):1–25.

Friedman, D. (1998). On economic applications of evolutionary game theory. *Journal of Evolutionary Economics*, 8(1):15–43.

Fryer, R. and Jackson, M. O. (2008). A categorical model of cognition and biased decision making. *The BE Journal of Theoretical Economics*, 8(1):1–44.

Galesloot, B. M. and Goyal, S. (1997). Costs of flexibility and equilibrium selection. *Journal of Mathematical Economics*, 28(3):249–264.

Gentzkow, M. and Kamenica, E. (2014). Costly persuasion. *American Economic Review*, 104(5):457–62.

Gerber, A. S. and Rogers, T. (2009). Descriptive social norms and motivation to vote: Everybody's voting and so should you. *The Journal of Politics*, 71(1):178–191.

Gherheș, V., Șimon, S., and Para, I. (2021). Analysing students' reasons for keeping their webcams on or off during online classes. *Sustainability*, 13(6):3203.

Glazer, J. and Rubinstein, A. (2004). On optimal rules of persuasion. *Econometrica*, 72(6):1715–1736.

Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137.

Goerg, S. J. and Selten, R. (2009). Experimental investigation of stationary concepts in cyclic duopoly games. *Experimental Economics*, 12(3):253–271.

Goldberg, M., Gustafson, A., Maibach, E., van der Linden, S., Ballew, M. T., Bergquist, P., Kotcher, J., Marlon, J. R., Rosenthal, S., and Leiserowitz, A. (2020). Social norms motivate covid-19 preventive behaviors. *PsyArXiv*.

Goldin, C. and Rouse, C. (2000). Orchestrating impartiality: The impact of" blind" auditions on female musicians. *American Economic Review*, 90(4):715–741.

Goldstein, N. J., Cialdini, R. B., and Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of consumer Research*, 35(3):472–482.

Golub, B. and Jackson, M. O. (2012). How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics*, 127(3):1287–1338.

Goyal, S., Hernández, P., Martínez-Cánovas, G., Moisan, F., Muñoz-Herrera, M., and Sánchez, A. (2021). Integration and diversity. *Experimental Economics*, 24:387–413.

Goyal, S. and Janssen, M. C. (1997). Non-exclusive conventions and social coordination. *Journal of Economic Theory*, 77(1):34–57.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408.

Gulesci, S., Jindani, S., La Ferrara, E., Smerdon, D., Sulaiman, M., and Young, P. H. (2021). A stepping stone approach to understanding harmful norms. *CEPR Discussion Paper No. DP15776*.

Güth, W. and Kliemt, H. (1994). Competition or co-operation: On the evolutionary economics of trust, exploitation and moral attitudes. *Metroeconomica*, 45(2):155–187.

Hamilton, W. D. (1964). The genetical evolution of social behaviour. ii. *Journal of Theoretical Biology*, 7(1):17–52.

Hamilton, W. D. (1967). Extraordinary sex ratios: A sex-ratio theory for sex linkage and inbreeding has new implications in cytogenetics and entomology. *Science*, 156(3774):477–488.

He, S. (2019). Minority advantage and disadvantage in competition and coordination. *Journal of Economic Behavior & Organization*, 163:464–482.

He, S. and Wu, J. (2020). Compromise and coordination: An experimental study. *Games and Economic Behavior*, 119:216–233.

Heemeijer, P., Hommes, C., Sonnemans, J., and Tuinstra, J. (2009). Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics and Control*, 33(5):1052–1072.

Heldt, T. (2005). Conditional cooperation in the field: cross-country skiers' behavior in sweden. In *Economic Science Association, European Meeting, Alessandria, Italien, September 15-18, 2005.*

Heller, Y. (2015). Three steps ahead. *Theoretical Economics*, 10(1):203–241.

Herold, F. and Kuzmics, C. (2020). The evolution of taking roles. *Journal of Economic Behavior & Organization*, 174:38–63.

Hommes, C., Sonnemans, J., Tuinstra, J., and Van de Velden, H. (2005). Coordination of expectations in asset pricing experiments. *The Review of Financial Studies*, 18(3):955–980.

Hommes, C., Sonnemans, J., Tuinstra, J., and Van De Velden, H. (2007). Learning in cobweb experiments. *Macroeconomic Dynamics*, 11(S1):8–33.

Howard, J. A. and Gibson, M. A. (2017). Frequency-dependent female genital cutting behaviour confers evolutionary fitness benefits. *Nature Ecology & Evolution*, 1(3):1–6.

Huntingford, F. A. and Turner, A. K. (1987). The consequences of animal conflict. In *Animal Conflict*, pages 227–250. Springer.

Hwang, S.-H., Lim, W., Neary, P. R., and Newton, J. (2018). Conventional contracts, intentional behavior and logit choice: Equality without symmetry. *Games and Economic Behavior*, 110:273–294.

Hwang, S.-H., Naidu, S., and Bowles, S. (2016). Social conflict and the evolution of unequal conventions. In *Santa Fe Institute Working Paper*.

Ichino, A. and Maggi, G. (2000). Work environment and individual background: Explaining regional shirking differentials in a large italian firm. *The Quarterly Journal of Economics*, 115(3):1057–1090.

Jackson, M. O. (2021). Inequality's economic and social roots: The role of social networks and homophily. *Available at SSRN 3795626*.

Jackson, M. O. and Rogers, B. W. (2007). Meeting strangers and friends of friends: How random are social networks? *American Economic Review*, 97(3):890–915.

Jackson, M. O. and Storms, E. (2019). Behavioral communities and the atomic structure of networks. *Available at SSRN 3049748*.

Jackson, M. O. and Watts, A. (2002). On the formation of interaction networks in social coordination games. *Games and Economic Behavior*, 41(2):265–291.

Jackson, M. O. and Yariv, L. (2007). Diffusion of behavior and equilibrium properties in network games. *American Economic Review*, 97(2):92–98.

Jacquemet, N. and Yannelis, C. (2012). Indiscriminate discrimination: A correspondence test for ethnic homophily in the chicago labor market. *Labour Economics*, 19(6):824–832.

Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory*, 123(2):81–104.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.

Kets, W. and Sandroni, A. (2019). A belief-based theory of homophily. *Games and Economic Behavior*, 115:410–435.

Khan, A. and Peeters, R. (2014). Cognitive hierarchies in adaptive play. *International Journal of Game Theory*, 43(4):903–924.

Kim, Y. C. and Loury, G. C. (2012). Stereotypes and inequality: A'signaling'theory of identity choice. *KDI Journal of Economic Policy*, 34(2):1–15.

Kim, Y.-C. and Loury, G. C. (2019). To be, or not to be: Stereotypes, identity choice and group inequality. *Journal of Public Economics*, 174:36–52.

Kimbrough, E. O., Laughren, K., and Sheremeta, R. (2020). War and conflict in economics: Theories, applications, and recent trends. *Journal of Economic Behavior & Organization*, 178:998–1013.

Klaus, B. and Newton, J. (2016). Stochastic stability in assignment problems. *Journal of Mathematical Economics*, 62:62–74.

Kreindler, G. E. and Young, H. P. (2013). Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior*, 80:39–67.

Krupka, E. and Weber, R. A. (2009). The focusing and informational effects of norms on pro-social behavior. *Journal of Economic Psychology*, 30(3):307–320.

Kuran, T. and Sandholm, W. H. (2008). Cultural integration and its discontents. *The Review of Economic Studies*, 75(1):201–228.

Lacroux, A. and Martin-Lacroux, C. (2020). Anonymous résumés: An effective preselection method? *International Journal of Selection and Assessment*, 28(1):98–111.

Ladley, D. (2012). Zero intelligence in economics and finance. *The Knowledge Engineering Review*, 27(2):273–286.

Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, 32(1):4–14.

Lee-Rife, S., Malhotra, A., Warner, A., and Glinski, A. M. (2012). What works to prevent child marriage: a review of the evidence. *Studies in Family Planning*, 43(4):287–303.

Lefebvre, M., Pestieau, P., Riedl, A., and Villeval, M. C. (2015). Tax evasion and social information: an experiment in belgium, france, and the netherlands. *International Tax and Public Finance*, 22(3):401–425.

Lewis, D. (2008). *Convention: A philosophical study.* John Wiley & Sons.

LiCalzi, M. and Mühlenbernd, R. (2019). Categorization and cooperation across games. *Games*, 10(1):5.

Lim, W. and Neary, P. R. (2016). An experimental investigation of stochastic adjustment dynamics. *Games and Economic Behavior*, 100:208–219.

Linek, M. and Traxler, C. (2021). Framing and social information nudges at wikipedia. *Journal of Economic Behavior & Organization*, 188:1269–1279.

Lipnowski, I. and Maital, S. (1983). Voluntary provision of a pure public good as the game of 'chicken'. *Journal of Public Economics*, 20(3):381–386.

Lobel, I. and Sadler, E. (2016). Preferences, homophily, and social learning. *Operations Research*, 64(3):564–584.

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531–542.

Martin, R. and Randal, J. (2008). How is donation behaviour affected by the donations of others? *Journal of Economic Behavior & Organization*, 67(1):228–238.

Mäs, M. and Nax, H. H. (2016). A behavioral study of "noise" in coordination games. *Journal of Economic Theory*, 162:195–208.

Mascagni, G. (2018). From the lab to the field: A review of tax experiments. *Journal of Economic Surveys*, 32(2):273–301.

Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98.

Matros, A. (2003). Clever agents in adaptive learning. *Journal of Economic Theory*, 111(1):110–124.

Matveenko, A. and Mikhalishchev, S. (2021). Attentional role of quota implementation. *Journal of Economic Theory*, 198:105356.

McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444.

Michaeli, M. and Spiro, D. (2017). From peer pressure to biased norms. *American Economic Journal: Microeconomics*, 9(1):152–216.

Mohlin, E. (2012). Evolution of theories of mind. *Games and Economic Behavior*, 75(1):299–318.

Mookherjee, D., Ray, D., and Napel, S. (2010). Social interactions and segregation in skill accumulation. *Journal of the European Economic Association*, 8(2-3):388–400.

Moro, A. and Norman, P. (2004). A general equilibrium model of statistical discrimination. *Journal of Economic Theory*, 114(1):1–30.

Myatt, D. P. and Wallace, C. (2003a). A multinomial probit model of stochastic evolution. *Journal of Economic Theory*, 113(2):286–301.

Myatt, D. P. and Wallace, C. (2003b). Sophisticated play by idiosyncratic agents. *Journal of Evolutionary Economics*, 13(3):319–345.

Naidu, S., Hwang, S.-H., and Bowles, S. (2010). Evolutionary bargaining with intentional idiosyncratic play. *Economics Letters*, 109(1):31–33.

Naidu, S., Hwang, S.-H., and Bowles, S. (2017). The evolution of egalitarian sociolinguistic conventions. *American Economic Review*, 107(5):572–77.

Nax, H. H. and Newton, J. (2022). Deep and shallow thinking in the long run. *Theoretical Economics*, 17(4):1501–1527.

Neary, P. R. (2012). Competing conventions. *Games and Economic Behavior*, 76(1):301–328.

Neary, P. R. and Newton, J. (2017). Heterogeneity in preferences and behavior in threshold models. *Journal of Mechanism and Institution Design*, 1(1):1.

Newton, J. (2018). Evolutionary game theory: A renaissance. *Games*, 9(2):31.

Newton, J. (2021). Conventions under heterogeneous behavioural rules. *The Review of Economic Studies*, 88(4):2094–2118.

Novak, L. (2020). Persistent norms and tipping points: The case of female genital cutting. *Journal of Economic Behavior & Organization*, 177:433–474.

Nyborg, K., Anderies, J. M., Dannenberg, A., Lindahl, T., Schill, C., Schlüter, M., Adger, W. N., Arrow, K. J., Barrett, S., Carpenter, S., et al. (2016). Social norms as solutions. *Science*, 354(6308):42–43.

Okafor, C. O. (2020). All things equal? social networks as a mechanism for discrimination. *arXiv preprint arXiv:2006.15988*.

Oprea, R., Henwood, K., and Friedman, D. (2011). Separating the hawks from the doves: Evidence from continuous time laboratory games. *Journal of Economic Theory*, 146(6):2206–2225.

Patacchini, E. and Zenou, Y. (2012). Ethnic networks and employment outcomes. *Regional Science and Urban Economics*, 42(6):938–949.

Peski, M. (2010). Generalized risk-dominance and asymmetric dynamics. *Journal of Economic Theory*, 145(1):216–248.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review*, 62(4):659–661.

Quilter, T. et al. (2007). Noise matters in heterogeneous populations. *ESE Discussion Papers*, 169.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review*, 83(5):1281–1302.

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., and Laland, K. N. (2010). Why copy others? insights from the social learning strategies tournament. *Science*, 328(5975):208–213.

Rudert, S. C. and Janke, S. (2022). Following the crowd in times of crisis: Descriptive norms predict physical distancing, stockpiling, and prosocial behavior during the covid-19 pandemic. *Group Processes & Intergroup Relations*, 25(7):1819–1835.

Sáez-Martı, M. and Weibull, J. W. (1999). Clever agents in young's evolutionary bargaining model. *Journal of Economic Theory*, 86(2):268–279.

Sandholm, W. H. (2010). *Population games and evolutionary dynamics*. MIT press.

Sawa, R. and Wu, J. (2018a). Prospect dynamics and loss dominance. *Games and Economic Behavior*, 112:98–124.

Sawa, R. and Wu, J. (2018b). Reference-dependent preferences, super-dominance and stochastic stability. *Journal of Mathematical Economics*, 78:96–104.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2):143–186.

Schelling, T. C. (1980). *The Strategy of Conflict: with a new Preface by the Author*. Harvard university press.

Schipper, B. C. (2009). Imitators and optimizers in cournot oligopoly. *Journal of Economic Dynamics and Control*, 33(12):1981–1990.

Schmitt, N., Tuinstra, J., and Westerhoff, F. (2017). Side effects of nonlinear profit taxes in an evolutionary market entry model: abrupt changes, coexisting attractors and hysteresis problems. *Journal of Economic Behavior & Organization*, 135:15–38.

Sederevičiūtė-Pačiauskienė, Ž., Valantinaitė, I., and Asakavičiūtė, V. (2022). 'should i turn on my video camera?'the students' perceptions of the use of video cameras in synchronous distant learning. *Electronics*, 11(5):813.

Shell-Duncan, B. and Herniund, Y. (2006). Are there "stages of change" in the practice of female genital cutting? qualitative research finding from senegal and the gambia. *African Journal of Reproductive Health*, 10(2):57–71.

Shell-Duncan, B., Wander, K., Hernlund, Y., and Moreau, A. (2011). Dynamics of change in the practice of female genital cutting in senegambia: testing predictions of social convention theory. *Social Science & Medicine*, 73(8):1275–1283.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118.

Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(5427):15–18.

Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge university press.

Smith, S., Windmeijer, F., and Wright, E. (2015). Peer effects in charitable giving: Evidence from the (running) field. *The Economic Journal*, 125(585):1053–1071.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355.

Staudigl, M. (2012). Stochastic stability in asymmetric binary choice coordination games. *Games and Economic Behavior*, 75(1):372–401.

Staudigl, M. and Weidenholzer, S. (2014). Constrained interactions and social coordination. *Journal of Economic Theory*, 152:41–63.

Tanaka, C. M., Lee, J.-H., and Iwasa, Y. (2018). The persistence of a local dialect when a national standard language is present: An evolutionary dynamics model of cultural diversity. *Bulletin of Mathematical Biology*, 80(10):2761–2786.

Thöni, C. and Gächter, S. (2015). Peer effects and social preferences in voluntary cooperation: A theoretical and experimental analysis. *Journal of Economic Psychology*, 48:72–88.

Tirole, J. (2009). Cognition and incomplete contracts. *American Economic Review*, 99(1):265–94.

Tuinstra, J., Wegener, M., and Westerhoff, F. (2014). Positive welfare effects of trade barriers in a dynamic partial equilibrium model. *Journal of Economic Dynamics and Control*, 48:246–264.

van Gerwen, N. and Buskens, V. (2018). Partner-specific behavior in social networks: Coordination among actors with heterogeneous preferences. *Social Networks*, 52:68–80.

Van Huyck, J., Battalio, R., Mathur, S., Van Huyck, P., and Ortmann, A. (1995). On the origin of convention: Evidence from symmetric bargaining games. *International Journal of Game Theory*, 24(2):187–212.

Vogt, S., Mohmmed Zaid, N. A., El Fadil Ahmed, H., Fehr, E., and Efferson, C. (2016). Changing cultural attitudes towards female genital cutting. *Nature*, 538(7626):506–509.

Wallace, C. and Young, H. P. (2015). Stochastic evolutionary game dynamics. In *Handbook of Game Theory with Economic Applications*, volume 4, pages 327–380. Elsevier.

Weibull, J. W. (1997). *Evolutionary game theory*. MIT press.

Woodcock, A. and Schultz, P. W. (2021). The role of conformity in mask-wearing during covid-19. *PLoS one*, 16(12):e0261321.

Young, H. P. (1993a). The evolution of conventions. *Econometrica*, 61(1):57–84.

Young, H. P. (1993b). An evolutionary model of bargaining. *Journal of Economic Theory*, 59(1):145–168.

Young, H. P. (1996). The economics of convention. *Journal of Economic Perspectives*, 10(2):105–122.

Young, H. P. (2009). Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American Economic Review*, 99(5):1899–1924.

Young, H. P. (2020). Individual strategy and social structure. In *Individual Strategy and Social Structure*. Princeton University Press.

Appendix A

Chapter 1 Appendix

Proof of Lemma 1. We have to formally show that each strategy inside the same behavior is behaviorally and payoff-equivalent for each player. Consider a player $i \in K$. Define $g_i^K(s_{-i})$ and $g_i^{K'}(s_{-i})$ as the frequencies of successful coordination for i on action a with group K and K' given strategy profile s_{-i}.

$$U_s^i((0,a,a,a),s_{-i}) = U_s^i((0,a,b,b),s_{-i}) =$$
$$U_s^i((0,a,b,a),s_{-i}) = U_s^i((0,a,a,b),s_{-i}) = \frac{N_K - 1}{N-1} g_i^K(s_{-i})\pi_a^K + \frac{N_{K'}}{N-1} g_i^{K'}(s_{-i})\pi_a^K.$$

Therefore, if $(0,a,a,a)$ is the maximizer, then also $(0,a,a,b)$, $(0,a,b,a)$, and $(0,a,b,b)$ are so. Hence, in this case, i maximizes her/his payoff by choosing behavior a. Moreover, consider $s'_{-i} = (0,a,a,b)^{N-1}$ and $s''_{-i} = (0,a,a,a)^{N-1}$. In this case $g_i^K(s'_{-i}) = g_i^K(s''_{-i})$, so for $g_i^{K'}$. Contrarily, if $s'''_{-i} = (0,b,a,a)^{N-1}$, $g_i^K(s'_{-i}) \neq g_i^K(s'''_{-i})$, so for $g_i^{K'}$. Therefore, $U_s^i(a,s'_{-i}) = U_s^i(a,s''_{-i}) = U_z^i(a,a^{N-1})$. Thanks to symmetry in payoff matrix, the argument stands for all strategies and behaviors. This passage completes the proof. □

Proof of Lemma 2. Consider a player $i \in K$ currently playing behavior a that is given a revision opportunity at period t. $g_i^K(z_{-i}(t))$ is the frequency of successful coordinations of player i on action a with group K at period t, given $z_{-i}(t)$. In this case, $U_z^i(a, z_{-i}(t)) = \frac{N_K-1}{N-1} g_i^K(z_{-i}(t))\pi_a^K + \frac{N_{K'}}{N-1} g_i^{K'}(z_{-i}(t))\pi_a^K$. Note that $g_i^K(z_{-i}(t)) = g_i^K(z_i(t),\omega_t)$ and $g_i^{K'}(z_{-i}(t)) = g_i^{K'}(z_i(t),\omega_t)$, where $g_i^K(z_i(t),\omega_t)$ is the frequency of successful coordinations of player i on action a with group K at period t, given ω_t and that player i is currently playing $z_i(t)$. Therefore, $U_z^i(a, z_{-i}(t)) = U_{z_i(t)}^i(a,\omega_t)$, with $z_i(t) = a$ in our case.

Note that $g_i^K(a,\omega_t) = \frac{n_t^{KK}-1}{N_K-1}$, and $g_i^K(b,\omega_t) = \frac{n_t^{KK}}{N_K-1}$. Moreover, $g_i^K(a,\omega_t) = g_i^K(aa,\omega_t) =$

$g_i^K(ab,\omega_t)$, and $g_i^K(b,\omega_t) = g_i^K(bb,\omega_t) = g_i^K(ba,\omega_t)$. Contrarily, $g_i^{K'}(z_i(t),\omega_t) = g_i^{K'}(\omega_t) = \frac{n_t^{K'K}}{N_{K'}}$, $\forall z_i(t) \in Z$.

Therefore, $U_a^i(a,\omega_t) = U_{aa}^i(a,\omega_t) = U_{ab}^i(a,\omega_t)$. Equally, $U_b^i(a,\omega_t) = U_{bb}^i(a,\omega_t) = U_{ba}^i(a,\omega_t)$. Thanks to symmetry in payoff matrix, the argument stands for all strategies and behaviors.

□

A.1 Proofs of Section 2.3

Proof of Lemma 3. Consider a player $i \in K$ currently playing aa who is given the revision opportunity at period t. On the one hand, $\forall n_t^{KK}$, $U_a^i(ab,\omega_t) = U_a^i(aa,\omega_t)$. On the other hand, $\forall n_t^{K'K}$, $U_a^i(ba,\omega_t) = U_a^i(aa,\omega_t)$. Therefore, player i chooses aa or ab depending on $n_t^{K'K}$, and ba or aa depending on n_t^{KK}.

Moreover, if player i chooses ab instead of aa, $n_{t+1}^{KK} = n_t^{KK}$, but $n_{t+1}^{K'K} < n_t^{K'K}$. If player i chooses ba instead of aa, $n_{t+1}^{KK} < n_t^{KK}$, but $n_{t+1}^{K'K} = n_t^{K'K}$. This passage completes the proof.

□

With abuse of notation, we call best reply (BR) the action optimally taken by a player in one of the three dynamics. For example, if a player of group A earns the highest payoff by playing a against a player of group B, we say that a is her/his BR. We do this in the context of complete information because of the separability of the dynamics.

Proof of Lemma 4. Thanks to Lemma 3, we can consider the three separated dynamics: n_t^{AA}, n_t^{BB}, and n_t^I.

Inside-group interactions.

Firstly, we prove the result for n_t^{AA}, and then the argument stands for n_t^{BB} thanks to symmetry of payoff matrix. We have to show that all the states in ω^R have an absorbing component for n_t^{AA}, that is, 0 or N_A. When $n^{AA} = N_A$, $\forall i \in A$, a is BR against group A at period t. Hence, $F_1(N_A,\theta_{t+1}) = N_A$. Symmetrically, if $n^{AA} = 0$, b is always BR, and so $F_1(0,\theta_{t+1}) = 0$. Therefore, N_A and 0 are fixed points for n_t^{AA}.

We need to show that these states are absorbing, that all the other states are transient, and that there are no cycles. Consider player $i \in A$, who is given the revision opportunity at period t. We define $\bar{n}^A$ as the minimum number of A players such that a is BR, and $\underline{\mathrm{n}}^A$ as the maximum number of A players such that b is BR. From Equations (2.1)–(2.4), we know that $\bar{n}^A = \frac{N_A\pi_A+\Pi_A}{\Pi_A+\pi_A}$, and that $\underline{\mathrm{n}}^A = \frac{N_A\pi_A-\pi_A}{\Pi_A+\pi_A}$. Assume $n_t^{AA} \geq \bar{n}^A$. There is always a positive probability that a player not playing a is given the revision opportunity. Hence, $F_1(n_t^{AA}, \theta_{t+1}) \geq n_t^{AA}$. Symmetrically, we can say that if $n_t^{AA} < \underline{\mathrm{n}}^A$, $F_1(n_t^{AA}, \theta_{t+1}) \leq n_t^{AA}$.

We now prove that if $n_t^{AA} \leq \underline{\mathrm{n}}^A \neq 0$,

$$Pr\left(\lim_{t\to\infty} F_1(n_t^{AA}, \theta_{t+1}) = n_t^{AA}\right) = 0.$$

Equally, if $n_t^{AA} \geq \bar{n}^A \neq N_A$,

$$Pr\left(\lim_{t\to\infty} F_1(n_t^{AA}, \theta_{t+1}) = n_t^{AA}\right) = 0.$$

We prove the first case, and the result stands for the second, thanks to symmetry in payoff matrices. Consider a period s in a state $n_s^{AA} < \underline{\mathrm{n}}^A \neq 0$. For every player, b is BR. Define $Pr\left(n_{s+1}^{AA} = n_s^{AA}\right) = p$. Such a probability represents the event that only players playing b are given the revision opportunity. $Pr\left(n_{s+2}^{AA} = n_s^{AA}\right) = p^2$, $Pr\left(n_{s+k}^{AA} = n_s^{AA}\right) = p^k$. If $k \to \infty$, $Pr\left(n_{s+k}^{AA} = n_s^{AA}\right) = 0$. Therefore,

$$\text{If } n_0^{AA} \leq \underline{\mathrm{n}}^A, \quad Pr\left(\lim_{t\to\infty} F_1(n_t^{AA}, \theta_{t+1}) = 0\right) = 1,$$

$$\text{If } n_0^{AA} \geq \bar{n}^A, \quad Pr\left(\lim_{t\to\infty} F_1(n_t^{AA}, \theta_{t+1}) = N_A\right) = 1.$$

Next, consider $\underline{\mathrm{n}}^A < n_0^{AA} < \bar{n}^A$. For every i playing a, b is BR, while, for every i' playing b, a is BR. There are no absorbing states between these states. If only players playing a are given the revision opportunity, they all choose b, and if enough of them are given the revision opportunity, $n_1^{AA} < \underline{\mathrm{n}}^A$. The opposite happens if only players playing b are given the revision opportunity.

Inter-group interactions.

We now pass to the analysis of n_t^I. We define four important values for n^{AB} and n^{BA}: $T_A = min\left\{n^{BA}|n^{BA} > \frac{\pi_A N_B}{\Pi_A + \pi_A}\right\}$, $T_B = min\left\{n^{AB}|n^{AB} > \frac{\Pi_B N_A}{\Pi_B + \pi_B}\right\}$,
$D_A = max\left\{n^{BA}|n^{BA} < \frac{\pi_A N_B}{\Pi_A + \pi_A}\right\}$, and $D_B = max\left\{n^{AB}|n^{AB} < \frac{\Pi_B N_A}{\Pi_B + \pi_B}\right\}$.

Given these values, we also define two sets of states, Ω_I^b and Ω_I^a:
$\Omega_I^a = \left\{n^I|n^{BA} \geq T_A \text{ and } n^{AB} \geq T_B\right\}$ and $\Omega_I^b = \left\{n^I|n^{BA} \leq D_A \text{ and } n^{AB} \leq D_B\right\}$.

With similar computation as for n_t^{AA}, we can say that $(0,0)$ and (N_A, N_B) are two fixed points for n_t^I. Are they absorbing states?

Consider the choice of a player $i \in A$ against player $j \in B$ and vice-versa. There can be four possible combinations of states: states in which a is BR for every player, states in which b is BR for every player, states, in which $\forall i \in A$, a is the best reply and b is the best reply $\forall j \in B$, and states for which the opposite is true. Let us call the third situation as Ω_I^{ab} and the fourth as Ω_I^{ba}.

Firstly, we prove that Ω_I^a and Ω_I^b are the regions where a and b are BR for every player. Secondly, we prove that there is no other absorbing state in Ω_I^a than (N_A, N_B), and no other absorbing state in Ω_I^b than $(0,0)$.

Assume that player $i \in A$ is given a revision opportunity at period t. From Equations (2.1)–(2.4), a is the BR against group B if $n_t^{BA} > \frac{\pi_A N_B}{\Pi_A + \pi_A}$. Since T_A is defined as the minimum value s.t., the latter holds, $\forall n_t^{BA} \geq T_A$, $\forall i \in A$, a is BR against B groups. Now, assume that $j \in B$ is given the revision opportunity, a is the BR against group A if $n_t^{AB} > \frac{\Pi_B N_A}{\Pi_B + \pi_B}$. Since T_B is defined as the minimum value s.t., this relation is true, $\forall n_t^{AB} \geq T_B$, a is the best reply $\forall j \in B$. Therefore, if $n_0^I \in \Omega_I^a$, $n_s^I \in \Omega_I^a, \forall s \geq 0$. Similarly, if $n_0^I \in \Omega_I^b$, $n_s^I \in \Omega_I^b, \forall s \geq 0$.

Consider being in a generic state $(T_B + d, T_A + d') \in \Omega_I^a$ at period t, with $d \in [0, N_A - T_B)$ and $d' \in [0, N_B - T_A)$. In such a state, there is always a probability p that a player not playing a is given the revision opportunity.

Therefore, if $n_t^I \in \Omega_I^a \setminus (N_A, N_B), Pr\left(F_{2,3}\left(n_t^I, \theta_{t+1}\right) \geq n_t^I\right) > p$ (meaning that $n_t^{'I} > n_t^{''I}$ if either $n_t^{'AB} > n_t^{''AB}$ and $n_t^{'BA} = n_t^{''BA}$ or $n_t^{'BA} > n_t^{''BA}$ and $n_t^{'AB} = n_t^{''AB}$ or both $n_t^{'BA} > n_t^{''BA}$ and $n_t^{'AB} > n_t^{''AB}$). Similar to what we proved before,

$$\text{if } n_t^I \in \Omega_I^a \setminus (N_A, N_B), \quad Pr\left(\lim_{t\to\infty} F_{2,3}(n_t^I, \theta_{t+1}) = n_t^I\right) = 0,$$

$$\text{if } n_t^I \in \Omega_I^b \setminus (0,0), \quad Pr\left(\lim_{t\to\infty} F_{2,3}(n_t^I, \theta_{t+1}) = n_t^I\right) = 0.$$

Consequently,

$$\text{If } n_0^I \in \Omega_I^a \quad Pr\left(\lim_{t\to\infty} F_{2,3}(n_t^I, \theta_{t+1}) = (N_A, N_B)\right) = 1,$$

$$\text{if } n_0^I \in \Omega_I^b, \quad Pr\left(\lim_{t\to\infty} F_{2,3}(n_t^I, \theta_{t+1}) = (0,0)\right) = 1.$$

We now consider Ω_I^{ab} and Ω_I^{ba}. Take an $n_0^I \in \Omega_I^{ab}$: at each period, there is a positive probability that only players of group A are given the revision opportunity, since for them a is the best reply, in the next period, there will be more or equal players in A playing a. Hence, if enough players of A that are currently playing b are given the revision opportunity, $n_1^I \in \Omega_I^a$. By the same reasoning, there is also a positive probability that only players from B are given the revision opportunity; hence, that $n_1^I \in \Omega_I^b$. The same can be said for every state in Ω_I^{ba}. Hence, starting from every state in $\Omega_I^{ab} \bigcup \Omega_I^{ba}$, there is always a positive probability to end up in Ω_I^a or Ω_I^b. □

Lemma 9. *Under complete information,*
$Pr\left(\lim_{t\to\infty} n_t^I = (N_A, N_B)\right) = 1 - Pr\left(\lim_{t\to\infty} n_t^I = (0,0)\right)$.
$Pr\left(\lim_{t\to\infty} n_t^{AA} = N_A\right) = 1 - Pr\left(\lim_{t\to\infty} n_t^{AA} = 0\right)$.
$Pr\left(\lim_{t\to\infty} n_t^{BB} = N_B\right) = 1 - Pr\left(\lim_{t\to\infty} n_t^{BB} = 0\right)$.

Proof. We prove the result for n_t^I, and the argument stands for the two other dynamics thanks to symmetry in the payoff matrix. Firstly, note that whenever the process starts in $\Omega_I^a \cup \Omega_I^b$, the lemma is always true thanks to the proof of Lemma 4. We need to show that this is the case, as well as when the process starts inside $\Omega_I^{ab} \bigcup \Omega_I^{ba}$. We prove the result for Ω_I^{ab} using the same logic, and the result stands for Ω_I^{ba} for symmetry of payoff matrix.

Take $n_0^I \in \Omega_I^{ab}$, define as p_a the probability of extracting m players from A that are currently playing b, and who would change action a if given a revision opportunity. Define as p_b the probability of picking m players from B currently choosing a who would change action to b if given a revision opportunity. The probability $1 - p_a - p_b$ defines all the other possibilities.

Let us take k steps forward in time:

$$Pr\left(n_k^I \in \Omega_I^a\right) \geq (p_a)^k$$

$$Pr\left(n_k^I \in \Omega_I^b\right) \geq (p_b)^k$$

$$Pr\left(n_k^I \in \Omega_I^{ab} \bigcup \Omega_I^{ba}\right) \leq (1 - p_a - p_b)^k.$$

Consider period $k + d$:

$$Pr\left(n_{k+d}^I \in \Omega_I^a\right) \geq (p_a)^k$$

$$Pr\left(n^I_{k+d} \in \Omega^b_I\right) \geq (p_b)^k$$

$$Pr\left(n^I_{k+d} \in \Omega^{ab}_I \bigcup \Omega^{ba}_I\right) \leq (1 - p_a - p_b)^{k+d}.$$

Clearly, the probability of being in $\Omega^a_I(\Omega^b_I)$ is now greater than or equal to $(p_a)^k((p_b)^k)$: we know that once in $\Omega^a_I(\Omega^b_I)$, the system stays there. The probability of being in $\Omega^{ab}_I \bigcup \Omega^{ba}_I$ consequently, is lower than $(1 - p_a - p_b)^{k+d}$.

Taking the limit for d that goes to infinity

$$\lim_{d\to\infty}\left(Pr\left(n^I_{k+d} \in \Omega^{ab}_I \bigcup \Omega^{ba}_I\right)\right) = 0.$$

This means that if we start in a state in Ω^{ab}_I there is no way of ending up in $\Omega^{ab}_I \bigcup \Omega^{ba}_I$ in the long-run; hence, the system ends up either in Ω^a_I or in Ω^b_I, but given this, we know that it ends up either in $(0, 0)$ or in (N_A, N_B). $\square$

Corollary 8. *Under complete information,*
$Pr\left(\lim_{t\to\infty} n^I_t = (N_A, N_B)\right) = 1$ *IFF* $n^I_0 \in \Omega^a_I$.
$Pr\left(\lim_{t\to\infty} n^I_t = (0, 0)\right) = 1$ *IFF* $n^I_0 \in \Omega^b_I$.
$Pr\left(\lim_{t\to\infty} n^{AA}_t = N_A\right) = 1$ *IFF* $n^{AA}_0 \in \left[\bar{n}^A, N_A\right]$, *and*
$Pr\left(\lim_{t\to\infty} n^{AA}_t = 0\right) = 1$ *IFF* $n^{AA}_0 \in \left[0, \underline{n}^A\right]$.
$Pr\left(\lim_{t\to\infty} n^{BB}_t = N_B\right) = 1$ *IFF* $n^{BB}_0 \in \left[\bar{n}^B, N_B\right]$, *and*
$Pr\left(\lim_{t\to\infty} n^{BB}_t = 0\right) = 1$ *IFF* $n^{BB}_0 \in \left[0, \underline{n}^B\right]$.

This result is a consequence of the previous lemmas, and therefore, the proof is omitted. Since the only two absorbing states in the dynamics of n^I_t are $(0, 0)$ and (N_A, N_B), they are the only two candidates to be stochastically stable states. From now on we call $(0, 0)$ as I^b_n and (N_A, N_B) as I^a_n. We define as 0_A the state where all players of group A play b with group A and 0_B the state where all players of group B play b with group B.

Let us call E_A and E_B the two values for which players in A and in B are indifferent in playing a or b in inter-group interactions. $E_A = \left\lceil \frac{N_B \pi_A}{\Pi_A + \pi_A} \right\rceil$ and $E_B = \left\lceil \frac{N_A \Pi_B}{\Pi_B + \pi_B} \right\rceil$. From now on, we often use values of N large enough to compare the arguments inside ceiling functions.

Lemma 10. *Under free information acquisition, for large enough* N, $R(I^b_n) = CR(I^a_n) = E_A$

for all values of payoffs and sizes of groups, while

$$R(I_n^a) = CR(I_n^b) = \begin{cases} N_A - E_B & \text{if } \frac{\pi_B}{\Pi_A} < \frac{N_B}{N_A} \\ N_B - E_A & \text{if } \frac{\pi_B}{\Pi_A} > \frac{N_B}{N_A} \end{cases}$$

Proof. Firstly, we know from Ellison (2000) that if there are just two absorbing states, the radius of one is the coradius of the other and vice-versa. Hence, $R(I_n^b) = CR(I_n^a)$, and $R(I_n^a) = CR(I_n^b)$. Moreover, from the proof of Lemma 4, we know that $D(I_n^a) = \Omega_I^a$ and $D(I_n^b) = \Omega_I^b$.

We prove that the minimum resistance path to exit the basin of attraction of I_n^b is the one that reaches $(E_B, 0)$ or $(0, E_A)$, and that the one to exit the basin of attraction of I_n^a is the one that reaches either (E_B, N_B) or (N_A, E_A). To prove this statement for I_n^b, firstly, note that once inside Ω_I^b, every step that involves a passage to a state with more people playing a requires a mistake. Secondly, note that in a state that is out of Ω_I^b, at least one of the two groups is indifferent in playing b or a, in other words, in a state where either $n^{AB} = E_B$ or $n^{BA} = E_A$ or both. Hence, the minimum resistance path to exit from I_n^b is the one either to $(E_B, 0)$ or to $(0, E_A)$. It is straightforward to show that all the other paths have greater resistance than the two above. Since we use uniform mistakes, every mistake counts the same value, and without loss of generality, we can count each of them as 1. Since every resistance counts as 1, then $R(I_n^b) = \min\{E_B; E_A\} = E_A$. Similarly, $R(I_n^a) = \min\{N_A - E_B; N_B - E_A\}$, and

$$N_A - E_B < N_B - E_A \Longleftrightarrow \frac{\pi_B}{\Pi_A} < \frac{N_B}{N_A}.$$

$\square$

Lemma 11. *Under free information acquisition, for large enough N, $R(0_A) = \left\lceil \frac{N_A\pi_A - \pi_A}{\Pi_A + \pi_A} \right\rceil$, $R(N_A) = \left\lceil \frac{N_A\Pi_A + \Pi_A}{\Pi_A + \pi_A} \right\rceil$, $R(0_B) = \left\lceil \frac{N_B\Pi_B + \Pi_B}{\Pi_B + \pi_B} \right\rceil$ and $R(N_B) = \left\lceil \frac{N_B\pi_B - \pi_B}{\Pi_B + \pi_B} \right\rceil$.*

Proof. The proof is straightforward; indeed, the minimum resistance path in terms of mistakes required to reach one absorbing state starting from the other one is the cost of exit from the basin of attraction of the first. As a matter of fact, let us consider $R(0_A)$; we know from the proof of Lemma 4 that we are out of the basin of attraction of 0_A when we reach the state $\underline{n}^A$. Hence, $R(0_A) = \left\lceil \frac{N_A\pi_A - \pi_A}{\Pi_A + \pi_A} \right\rceil$. The same applies to the other states. $\square$

Proof of Theorem 1. We divide the proof for the three dynamics described so far: for what concerns n_t^{AA}, N_A is uniquely stochastically stable, and for what concerns n_t^{BB}, 0_B is uniquely stochastically stable; this proof follows directly from Lemma 11 and therefore is omitted.

Let us pass to n_t^I. We know from Lemma 10 that $R(I_n^b) = E_A$ and that the value of $R(I_n^a)$ depends on payoffs and group size. Let us firstly consider the case when $\frac{\pi_B}{\Pi_A} < \frac{N_B}{N_A}$ and $R(I_n^a) = N_A - E_B$. It is sufficient that $E_A > N_A - E_B$ for I_n^b to be uniquely stochastically stable. Indeed, if this happens, $R(I_n^b) > CR(I_n^b)$. This is the case IFF

$$\frac{\pi_A N_B}{\Pi_A + \pi_A} > \frac{\pi_B N_A}{\Pi_B + \pi_B} \iff \frac{\pi_B}{\pi_A} < \frac{N_B}{N_A}. \tag{A.1}$$

To complete the proof, we show that whenever $\frac{\pi_B}{\pi_A} > \frac{N_B}{N_A}$, then I_n^a is the uniquely stochastically stable state. Firstly, note that $\frac{\pi_B}{\Pi_A} < \frac{\pi_B}{\pi_A}$; hence, for $\frac{\pi_B}{\pi_A} > \frac{N_B}{N_A} > \frac{\pi_B}{\Pi_A}$, $R(I_n^a) = N_A - E_B$ and $E_A = R(I_n^b)$. However, Equation (A.1) is reversed, so, I_n^a is uniquely stochastically stable. For $\frac{\pi_B}{\pi_A} > \frac{\pi_B}{\Pi_A} > \frac{N_B}{N_A}$, $R(I_n^a) = N_B - E_A$ and still $R(I_n^b) = E_A$. In this case, I_n^a is the uniquely stochastically stable if $E_A < N_B - E_A$, hence, IFF

$$\frac{\pi_A N_B}{\Pi_A + \pi_A} < \frac{\Pi_A N_B}{\Pi_A + \pi_A}.$$

This happens for every value of the payoffs (given that $\Pi_A > \pi_A$) and of the group size. We conclude that whenever $\frac{\pi_B}{\pi_A} < \frac{N_B}{N_A}$, PS_b is uniquely stochastically stable and when $\frac{\pi_B}{\pi_A} > \frac{N_B}{N_A}$, PS_a is uniquely stochastically stable. □

A.2 Proofs of Section 2.4

For convenience, we call behavior τ_1 the optimal behavior when a player decides to acquire the information: $\tau_1 = \max(ab, ba, aa, bb)$.

We will use in some proofs the concept of Modified Coradius from Ellison (2000). We write here the formal definition. Suppose $\bar{\omega}$ is an absorbing state and $(\omega_1, \omega_2, \dots \omega_T)$ is a path from state ω' to $\bar{\omega}$. Let $L_1, L_2, \dots, L_r = \bar{\omega}$ be the sequence of limit sets through which the path passes consecutively. The modified resistance is the original resistance minus the radius of the intermediate limit sets through which the path passes,

$$r^*(\omega_1, \omega_2, \dots \omega_T) = r(\omega_1, \omega_2, \dots \omega_T) - \sum_{i=2}^{r-1} R(L_i).$$

Define

$$r^*(\omega', \bar{\omega}) = \min_{(\omega_1,\omega_2,\dots\omega_T)\in\Upsilon(\omega',\bar{\omega})} r^*(\omega_1, \omega_2, \dots \omega_T),$$

the Modified Coradius is defined as follows

$$CR^*(\bar{\omega}) = \max_{\omega' \neq \bar{\omega}} r^*(\omega', \bar{\omega}).$$

Note that $CR^*(\bar{\omega}) \leq CR(\bar{\omega})$. Thanks to Theorem 2 in Ellison (2000), we know that when $R(\bar{\omega}) > CR^*(\bar{\omega})$, $\bar{\omega}$ is uniquely stochastically stable.

Proof of Lemma 5. We first show that the nine states are effectively strict equilibria, that there is no other possible equilibrium, and that a state is absorbing if and only if it is a strict equilibrium.

Monomorphic States.

It is easy to show that (N_A, N_A, N_B, N_B) and $(0, 0, 0, 0)$ are two strict equilibria. We take the first case, and the argument stands also for the second, thanks to the symmetry of the payoff matrix. Consider player $i \in K$ who is given the revision opportunity at period t:

$$U_a^i(a, \omega_t) = \frac{N_K + N_{K'} - 1}{N - 1}\pi_a^K = \pi_a^K,$$

$$U_a^i(b, \omega_t) = \frac{N - N_K - N_{K'}}{N - 1}\pi_b^K = 0,$$

$$U_a^i(\tau_1, \omega_t) = \frac{N_K + N_{K'} - 1}{N - 1}\pi_a^K - c = \pi_a^K - c.$$

(N_A, N_A, N_B, N_B) is a strict equilibrium since $\pi_a^K > 0$ and $c > 0$.

Polymorphic States.

Firstly let us consider the case of $(N_A, 0, 0, N_B)$. Since in this case, every player is playing ab, the state is a strict equilibrium IFF $\max z_i^o = ab$, $\forall i \in N$. If player $i \in K$ is given the revision opportunity at period t,

$$U_a^i(a, \omega_t) = \frac{N_K - 1}{N - 1}\pi_a^K,$$

$$U_a^i(b, \omega_t) = \frac{N_{K'}}{N - 1}\pi_b^K,$$

$$U_a^i(\tau_1, \omega_t) = \frac{N_K - 1}{N-1}\pi_a^K + \frac{N_{K'}}{N-1}\pi_b^K - c.$$

For group A players, $U_a^i(a, \omega_t) > U_a^i(b, \omega_t)$ since $\frac{N_B}{N-1}\pi_A < \frac{N_A - 1}{N-1}\Pi_A$. For group B players, $U_a^i(b, \omega_t) > U_a^i(a, \omega_t)$ as $\frac{N_B - 1}{N-1}\pi_B < \frac{N_A}{N-1}\Pi_B$. $U_a^i(\tau_1, \omega_t)$ is the highest of the three $\forall i \in N$ IFF $c < \frac{1}{N-1}\min\{N_B\pi_A, N_B - 1\pi_B\}$.

Consider the case of $(0, N_A, N_B, 0)$; since every player is playing ba, it must be that $\max z_i^o = ba$. $i \in K$ faces the following payoffs

$$U_b^i(a, \omega_t) = \frac{N_{K'}}{N-1}\pi_a^K,$$

$$U_b^i(b, \omega_t) = \frac{N_K - 1}{N-1}\pi_b^K,$$

$$U_b^i(\tau_1, \omega_t) = \frac{N_K - 1}{N-1}\pi_b^K + \frac{N_{K'}}{N-1}\pi_a^K - c.$$

Note that $U_b^i(a, \omega_t) > U_b^i(b, \omega_t)$ IFF $\frac{\pi_b^K}{\pi_a^K} < \frac{N_{K'}}{N_K - 1}$, and therefore ba is the best reply behavior in this case if $c < \frac{N_K - 1}{N-1}\pi_b^K$. When the opposite happens, and so $\frac{\pi_b^K}{\pi_a^K} > \frac{N_{K'}}{N_K - 1}$, ba is the best reply behavior if $c < \frac{N_{K'}}{N-1}\pi_a^K$. These conditions take the form of the ones in Table A1.

Consider the remaining four PS; they are characterized by the following fact: $BR(n^{KK}) = BR(n^{K'K})$ but $BR(n^{K'K'}) \neq BR(n^{KK'})$. In this case, it must be that $\tau_i = 0$ is optimal for $i \in K$ while $\tau_j = 1$ is optimal for $j \in K'$. Thanks to the symmetry in payoff matrices, we can say that the argument to prove the results for these four states is similar to the one for $(N_A, 0, 0, N_B)$ and $(0, N_A, N_B, 0)$. All the conditions are listed in Table A1.

Type Monomorphic State.

$(N_A, N_A, 0, 0)$ is a strict equilibrium if a is the BR $\forall i \in A$ and b, $\forall j \in B$. Consider a player $i \in A$, who is given the revision opportunity at period t:

$$U_a^i(a, \omega_t) = \frac{N_A - 1}{N-1}\Pi_A,$$

$$U_a^i(b, \omega_t) = \frac{N_B}{N-1}\pi_A,$$

$$U_a^i(\tau_1, \omega_t) = \frac{N_A - 1}{N - 1}\Pi_A + \frac{N_B}{N - 1}\pi_A - c.$$

Given that $U_a^i(a, \omega_t) > U_a^i(b, \omega_t)$, a is the best reply behavior IFF $c > \frac{N_B}{N-1}\pi_A$. Consider player $j \in B$:

$$U_b^j(a, \omega_t) = \frac{N_A}{N - 1}\pi_B,$$

$$U_b^j(b, \omega_t) = \frac{N_B - 1}{N - 1}\Pi_B,$$

$$U_b^j(\tau_1, \omega_t) = \frac{N_A}{N - 1}\pi_B + \frac{N_B - 1}{N - 1}\Pi_B - c.$$

In this case, when $\frac{\pi_B}{\Pi_B} > \frac{N_B - 1}{N_A}$, b is never the best reply and a is the best reply; hence, the state can not be a strict equilibrium. When $\frac{\pi_B}{\Pi_B} < \frac{N_B - 1}{N_A}$, $U_b^j(b, \omega_t) > U_b^j(a, \omega_t)$, and $U_b^j(b, \omega_t) > U_b^j(\tau_1, \omega_t)$ IFF $c > \frac{N_A}{N-1}\pi_B$.

No other state is a strict equilibrium.

For what concerns states where not all players of a group are playing the same action with the same group, this is easy to prove. Indeed, by definition, in these states, either not all players are playing their best reply action, or players are indifferent between two or more behaviors. In the first case, the state is not a strict equilibrium by definition; in the second case, there is no strictness of the equilibrium since there is not one best reply, but more behaviors can be best reply simultaneously. Hence, such states can not be strict equilibria. We are left with the seven states where every player of one group is doing the same thing against the same group. Such states are: $(0, 0, N_B, N_B)$, $(0, N_A, 0, N_B)$, $(N_A, 0, N_B, 0)$, $(0, 0, N_B, 0)$, $(N_A, N_A, 0, N_B)$, $(0, N_A, 0, 0)$, and $(N_A, 0, N_B, N_B)$. It is easy to prove that these states enter in the category of states where not every player is playing her/his best reply. Therefore, they can not be strict equilibria.

Strict equilibria are always absorbing states.

We first prove the sufficient and necessary conditions to be a fixed point, and second that every fixed point is an absorbing state. To prove the sufficient part, we rely on the definition of strict equilibrium. In a strict equilibrium, every player is playing her/his BR, and no one has the incentive to deviate. Whoever is given the revision opportunity does not change her/his behavior. Therefore, $F(\omega_t, \theta_{t+1}) = \omega_t$. To prove the necessary condition, think about being in a state that is not a strict equilibrium; in this case, by definition, we know that not all the players are playing their BR. Among them, there are states in which there are no indifferent players. In this case, with positive probability, one or more players who are not playing their BR are given the revision opportunity and they change action; therefore, $F(\omega_t, \theta_{t+1}) \neq \omega_t$ for some realization of θ_{t+1}. In states where some players are indifferent between two or more behaviors, thanks to the tie rule, there is always a positive probability that the indifferent player changes her/his action since s/he is randomizing her/his choice. Moreover, there is also a positive probability to select a player indifferent between two or more behaviors. In this case, s/he changes the one that is currently playing with a positive probability too. Knowing this, we are sure that no state outside strict equilibria can be a fixed point. In our case, a fixed point is also an absorbing state by definition. Indeed, every fixed point absorbs at least one state: the one where all players except one are playing the same behavior. In this case, if that player is given the revision opportunity, s/he changes for sure her/his behavior into the one played by every player. $\square$

State	**Conditions**	**Conditions on c**
MS_a	none	none
MS_b	none	none
TS	$\frac{\pi_B}{\Pi_B} < \frac{N_B-1}{N_A}$	$c > \max\left\{\frac{N_B}{N-1}\pi_A, \frac{N_A}{N-1}\pi_B\right\}$
PS_b	none	$c < \frac{N_B}{N-1}\pi_A$
PS_a	(1) $\frac{\pi_B}{\Pi_B} > \frac{N_B-1}{N_A}$ (2) $\frac{\pi_B}{\Pi_B} < \frac{N_B-1}{N_A}$	(1) $c < \frac{N_B-1}{N-1}\Pi_B$ (2) $c < \frac{N_A}{N-1}\pi_B$
$(0, N_A, N_B, N_B)$	(1) $\frac{\pi_A}{\Pi_A} < \frac{N_B}{N_A-1}$ (2) $\frac{\pi_A}{\Pi_A} > \frac{N_B}{N_A-1}$	(1) $c < \frac{N_A-1}{N-1}\pi_A$ (2) $c < \frac{N_B}{N-1}\Pi_A$
$(N_A, 0, 0, N_B)$	none	$c < \min\left\{\frac{N_B}{N-1}\pi_A, \frac{N_B-1}{N-1}\pi_B\right\}$
$(0, N_A, N_B, 0)$	(1) $\frac{\pi_A}{\Pi_A} < \frac{N_B}{N_A-1}$ and $\frac{\pi_B}{\Pi_B} > \frac{N_B-1}{N_A}$ (2) $\frac{\pi_A}{\Pi_A} > \frac{N_B}{N_A-1}$ and $\frac{\pi_B}{\Pi_B} > \frac{N_B-1}{N_A}$ (3) $\frac{\pi_A}{\Pi_A} < \frac{N_B}{N_A-1}$ and $\frac{\pi_B}{\Pi_B} < \frac{N_B-1}{N_A}$ (4) $\frac{\pi_A}{\Pi_A} > \frac{N_B}{N_A-1}$ and $\frac{\pi_B}{\Pi_B} < \frac{N_B-1}{N_A}$	(1) $c < \min\left\{\frac{N_A-1}{N-1}\pi_A, \frac{N_B-1}{N-1}\Pi_B\right\}$ (2) $c < \min\left\{\frac{N_B}{N-1}\Pi_A, \frac{N_B-1}{N-1}\Pi_B\right\}$ (3) $c < \min\left\{\frac{N_A-1}{N-1}\pi_A, \frac{N_A}{N-1}\pi_B\right\}$ (4) $c < \min\left\{\frac{N_B}{N-1}\Pi_A, \frac{N_A}{N-1}\pi_B\right\}$
$(0, 0, 0, N_B)$	none	$c < \frac{N_B-1}{N-1}\pi_B$

Table A1: Necessary and sufficient conditions for absorbing states.

Here, we prove the results of the stochastic stability analysis of Section 2.4.

Proof of Theorem 2. We split the absorbing states into two sets and then apply Theorem 1 by Ellison (2000). Define the following two sets of states: $M_1 = \{PS_a, PS_b\}$ and $M_2 = (PS \setminus M_1) \cup MS$. Similarly, define $M_1' = PS_b$ and $M_2' = MS \cup (PS \setminus M_1')$.

Analysis with M_1 and M_2.

$R(M_1)$ is the minimum number of mistakes to escape the basins of attraction of both PS_a and PS_b. The dimension of these basins of attraction is determined by the value of c. In a state inside $D(PS_a)$, ba is BR for B, and a is BR for A. Similarly, ab is optimal for A inside $D(PS_b)$ and b is optimal for B. The minimum resistance paths that start in PS_a, and PS_b and exit from their basins of attraction involve ε mistakes.

We calculate the dimension of these basins of attraction for $0 < c < \frac{1}{N-1}\min\{\pi_A, \pi_B\}$. We start from PS_a and the argument stands for the other states in PS for symmetry of payoffs matrix.

Firstly, we consider the minimum number of mistakes that makes a BR for B players. Consider the choice of a B player inside a category of states where $n^{BB} \in \left[0, \frac{N_B\Pi_B - \Pi_B}{\Pi_B+\pi_B}\right)$ and $n^{AB} \in \left(\frac{N_A\Pi_B}{\Pi_B+\pi_B}, N_A\right]$. Referring to Equations (2.5)–(2.8), the optimal level of c s.t. 1 is the best reply, and for B players, it is

$$c < \min\left\{\frac{N_B\Pi_B - n^{BB}(\Pi_B+\pi_B) - \Pi_B}{N-1}, \frac{n^{AB}(\Pi_B+\pi_B) - N_A\Pi_B}{N-1}\right\}.$$

If $0 < c < \frac{1}{N-1}\min\{\pi_A, \pi_B\}$, whenever $n^{BB} \in \left[0, \frac{N_B\Pi_B-\Pi_B}{\Pi_B+\pi_B}\right)$ and $n^{AB} \in \left(\frac{N_A\Pi_B}{\Pi_B+\pi_B}, N_A\right]$, 1 is the BR for B. Therefore, a path towards a state where $n^{BB} \geq \frac{N_B\Pi_B+\pi_B}{\Pi_B+\pi_B}$ is a transition out of the basin of attraction of PS_a. Starting from $n^{BB} = 0$, the cost of this transition is $\frac{N_B\Pi_B+\pi_B}{\Pi_B+\pi_B}$. This cost is determined by ε mistakes, since once in PS_a, it is sufficient that a number of B plays by mistake b. Another possible path is to make ba BR for A. The cost of this transition is $\frac{N_A\Pi_A+\pi_A}{\Pi_A+\pi_A}$. With similar arguments, it is possible to show that the cost of exit from M_1 starting from PS_b is the same. For this reason, $R(M_1) = \min\left\{\frac{N_B\Pi_B+\pi_B}{\Pi_B+\pi_B}, \frac{N_A\Pi_A+\pi_A}{\Pi_A+\pi_A}\right\}$.

We can show that the minimum resistance path to exit from the basin of attraction of M_2 reaches either PS_a from MS_a or PS_b from MS_b. Therefore, $R(M_2) = \min\left\{\frac{N_A\pi_A+\Pi_A}{\Pi_A+\pi_A}, \frac{N_B\pi_B+\Pi_B}{\Pi_B+\pi_B}\right\}$. $R(M_1) > R(M_2)$ for every value of payoffs and group size: the stochastically stable state must be in M_1.

Analysis with M_1' and M_2'

Let us consider the path that goes from M_1' to PS_a. Starting in PS_b, it is sufficient that $\frac{N_B\pi_A}{\Pi_A+\pi_A}$ players from A play a for a transition from PS_b to $D(PS_a)$ to happen.

Since $\frac{N_B\pi_A}{\Pi_A+\pi_A} < \min\left\{\frac{N_B\Pi_B+\pi_B}{\Pi_B+\pi_B}, \frac{N_A\Pi_A+\pi_A}{\Pi_A+\pi_A}\right\}$, we can say that $R(M_1') = \frac{N_B\pi_A}{\Pi_A+\pi_A}$. With a similar argument, it can be shown that $R(M_2') = \frac{N_A \pi_B}{\Pi_B + \pi_B}$. When $R(M_2') > R(M_1')$, PS_a is uniquely stochastically stable. When $R(M_1') > R(M_2')$, PS_b is uniquely stochastically stable.

$R(M_2') \lesseqgtr R(M_1')$ when $\frac{N_B}{N_A} \lesseqgtr \frac{\pi_B}{\pi_A}$. □

Proof of Theorem 3. In this case, $R(MS_a) = CR(MS_b)$ and $R(MS_b) = CR(MS_a)$. Therefore, we just need to calculate the two Radius.

Radius of each state.

Let us consider $R(MS_a)$. Since the basin of attraction of MS_a is a region where a is the best reply behavior for both groups, many players should make a mistake such that b becomes BR for one of the two groups. For b to be BR for B players, it must be that $n^{AB}+n^{BB} \leq \frac{N\Pi_B - \Pi_B}{\Pi_B + \pi_B}$. This state can be reached with ε mistakes at cost $\frac{N\pi_B + \Pi_B}{\Pi_B + \pi_B}$. In a state where $n^{AA} + n^{BA} \leq \frac{N\pi_A - \pi_A}{\Pi_A + \pi_A}$, b is BR for A, this path happens at cost $\frac{N\Pi_A + \pi_A}{\Pi_A + \pi_A}$. In principle, $\frac{N\Pi_A + \pi_A}{\Pi_A + \pi_A} > \frac{N\pi_B + \Pi_B}{\Pi_B + \pi_B}$, hence, $R(MS_a)$ should be $\frac{N\pi_B + \Pi_B}{\Pi_B + \pi_B}$. However, it may not be sufficient to reach such a state.

Consider reaching a state s.t. $n^{AB} + n^{BB} = \frac{N\Pi_B - \Pi_B}{\Pi_B + \pi_B}$, since $\frac{N\Pi_B - \Pi_B}{\Pi_B + \pi_B} > \frac{N\pi_A - \pi_A}{\Pi_A + \pi_A}$, it must be that a is still the best reply $\forall i \in A$, and therefore, there is a path of zero resistance to MS_a. Nevertheless, once in that state, it can happen that only B players are given the revision opportunity, and that they all choose behavior b. This creates a path of zero resistance to a state $(\bar{n}^{AA}, \bar{n}^{AB}, 0, 0)$. Once in this state, if $\bar{n}^{AA} < \frac{N\pi_A - \pi_A}{\Pi_A + \pi_A}$, the state is in the basin of attraction of MS_b. This happens only if $\frac{N\pi_A - \pi_A}{\Pi_A + \pi_A} + N_B = \frac{N\Pi_B - \Pi_B}{\Pi_B + \pi_B}$. More generally, considering $k \geq 0$, this happens if $\frac{N\pi_A - \pi_A}{\Pi_A + \pi_A} + N_B = \frac{N\Pi_B - \Pi_B}{\Pi_B + \pi_B} - k$. Fixing payoffs and groups size, $k = \frac{N\Pi_B - \Pi_B}{\Pi_B + \pi_B} - \frac{N\pi_A - \pi_A}{\Pi_A + \pi_A} - N_B$; hence, the cost of this path would be

$$\frac{N\pi_B + \Pi_B}{\Pi_B + \pi_B} + \frac{N\Pi_B - \Pi_B}{\Pi_B + \pi_B} - \frac{N\pi_A - \pi_A}{\Pi_A + \pi_A} - N_B = N_A - \frac{N\pi_A - \pi_A}{\Pi_A + \pi_A}.$$

With a similar reasoning, $R(MS_b) = \frac{N\pi_A + \Pi_A}{\Pi_A + \pi_A}$.

We prove that all the other paths with η mistakes are costlier than ones with ε. We know that a is the BR for every state inside the basin of attraction of MS_a, nobody in the basin of attraction of MS_a optimally buys the information, and every player who once bought the information (by mistake) plays behavior aa. Every path with an η mistake also involves an ε mistake, and hence is double that of the one described above.

Conditions for stochastically stable states.

MS_a is stochastically stable IFF $N_A - \frac{N\pi_A - \pi_A}{\Pi_A + \pi_A} > \frac{N\pi_A + \Pi_A}{\Pi_A + \pi_A}$, this is verified when $N_A > \frac{2N\pi_A + \Pi_A - \pi_A}{\Pi_A + \pi_A}$. Therefore, we conclude that MS_a is stochastically stable in the above scenario, while if the opposite happens, MS_b is stochastically stable. □

Proof of Theorem 4. We first calculate radius, coradius, and modified coradius for the three states we are interested in, and then we compare them to draw inference about stochastic stability.

Radius of each state.

The Radius of MS_a is the minimum number of mistakes that makes b BR for B players. This number is $\frac{N\pi_B + \Pi_B}{\Pi_B + \pi_B}$. The alternative is to make b BR for A: hence, a path to state $(0, 0, N_B, N_B)$, and then to $(0, 0, 0, 0)$. The number of ε mistakes for this path is $\frac{N\Pi_A + \pi_A}{\Pi_A + \pi_A}$. Therefore, $R(MS_a) = \frac{N\pi_B + \Pi_B}{\Pi_B + \pi_B}$. With a similar reasoning, we can conclude that $R(MS_b) = \frac{N\pi_A + \Pi_A}{\Pi_A + \pi_A}$.

Consider TS: the minimum-resistance path to exit from its basin of attraction reaches either MS_a or MS_b, depending on payoffs. In other words, the minimum number of mistakes to exit from $D(TS)$ is the one that makes either a or b as BR. Consider the path from TS to MS_a: in this case, some mistakes are needed to make a BR for B. The state in which a is BR for B depends on payoffs and group size. In a state (N_A, N_A, k', k'), a is BR for every player in B if $(N_A + k' - 1)\pi_B > (N - N_A - k')\Pi_B$. This inequality is obtained by declining Equations (2.5)–(2.8), comparing B playing a/ab or b/ba. Fixing payoffs, we can calculate the exact value of k' that is $\frac{N_B\Pi_B - N_A\pi_B + \pi_B}{\Pi_B + \pi_B}$; this would be the cost of the minimum mistake transition from TS to MS_a. With a similar argument, the cost of the minimum mistake transition from TS to MS_b is $\frac{N_A\Pi_A - N_B\pi_A + \pi_A}{\Pi_A + \pi_A}$.

There are no paths involving η mistakes that are lower than the two proposed above. The intuition is the following. Consider a situation in which m players of A are given the revision opportunity at one period, and they all choose to buy the information. In this case, they all optimally choose behavior ab. This means that at the cost of n, there is a path to a state in which $N_A - m$ players are playing b against B, in this state, b is still the BR for group B, while a is still the BR for A. Hence, from that state, there is a path of zero resistance to TS. The same happens when B players choose by mistake to buy the information. Therefore, $R(TS) = \min\left\{\frac{N_B\Pi_B - N_A\pi_B + \pi_B}{\Pi_B + \pi_B}, \frac{N_A\Pi_A - N_B\pi_A + \pi_A}{\Pi_A + \pi_A}\right\}$.

Coradius of each state.

We start from TS: in this case, we have to consider the two minimum-resistance paths to reach it from MS_a and MS_b. Therefore, $\frac{N\pi_A+\Pi_A}{\Pi_A+\pi_A}$ and $\frac{N\pi_B+\Pi_B}{\Pi_B+\pi_B}$. Firstly, the argument to prove that these two are the minimum resistance paths to reach TS from MS_a and MS_b are given by the previous part of the proof. Secondly, we have to prove that this path is the maximum among the minimum resistance paths starting from any other state and ending in TS. There are three regions from which we can start and end up in TS: the basin of attraction of MS_b, the one of MS_a, and all the other states that are not in the basins of attraction of the three states considered. We can say that from this region, there is always a positive probability of ending up in MS_a, MS_b, or TS. Hence, we can consider as 0 the cost to reach TS from this region. The other two regions are the one considered above, and since we are taking the maximum path to reach TS from any other state, we have to take the sum of this two. Hence, $CR(TS) = \frac{N\pi_A+\Pi_A}{\Pi_A+\pi_A} + \frac{N\pi_B+\Pi_B}{\Pi_B+\pi_B}$.

Let us think about MS. Similarly to the two previous proofs, we can focus only on ε paths. Note that in this case, TS is always placed between the two MS. Let us start from MS_b: in this case we can consider three different paths starting from any state and arriving to MS_b. The first one starts in TS, the second starts in every state outside the basin of attractions of the three absorbing states, and the last starts in MS_a. In the second case, there is at least one transition of zero resistance to MS_b. Next, assume starting in TS: the minimum number of mistakes to reach MS_b from TS is the one that makes b BR for A players. Therefore, $\frac{N_A\Pi_A-N_B\pi_A+\pi_A}{\Pi_A+\pi_A}$.

Now, we need to consider the case of starting in MS_a. Firstly, consider the minimum number of mistakes to make b BR for A players. This number is $\frac{N\Pi_A+\pi_A}{\Pi_A+\pi_A}$. Secondly, consider the minimum number of mistakes to make b BR for B players, and then once TS is reached, the minimum number of mistakes that makes b BR for A players.

$$\min r(MS_a, MS_b) = \min\left\{\frac{N\Pi_A+\pi_A}{\Pi_A+\pi_A}, \frac{N\pi_B+\Pi_B}{\Pi_B+\pi_B} + \frac{N_A\Pi_A-N_B\pi_A+\pi_A}{\Pi_A+\pi_A}\right\}.$$

Since the two numbers in the expression are all greater than $\frac{N_A\Pi_A-N_B\pi_A+\pi_A}{\Pi_A+\pi_A}$, we can say that $CR(MS_b) = \min r(MS_a, MS_b)$.

Reaching a state where b is BR for group A from TS is for sure less costly than reaching it from MS_a, since in TS there are more people playing b. Therefore, $\frac{N\Pi_A+\pi_A}{\Pi_A+\pi_A} \geq \frac{N\pi_B+\Pi_B}{\Pi_B+\pi_B} + \frac{N_A\Pi_A-N_B\pi_A+\pi_A}{\Pi_A+\pi_A}$, hence, $CR(MS_b) = \frac{N\pi_B+\Pi_B}{\Pi_B+\pi_B} + \frac{N_A\Pi_A-N_B\pi_A+\pi_A}{\Pi_A+\pi_A}$. With a similar reasoning, $CR(MS_a) = \frac{N\pi_A+\Pi_A}{\Pi_A+\pi_A} + \frac{N_B\Pi_B-N_A\pi_B+\pi_B}{\Pi_B+\pi_B}$.

Modified Coradius of each state.

Firstly, note that $CR(TS) = CR^*(TS)$, since between MS and TS, there are no intermediate states. Formally,

$$CR^*(TS) = \min r^*(MS_a, \ldots, TS) + \min r^*(MS_b, \ldots, TS) = \frac{N\pi_A + \Pi_A}{\Pi_A + \pi_A} + \frac{N\pi_B + \Pi_B}{\Pi_B + \pi_B}.$$

The maximum path of minimum resistance from each MS to the other MS passes through TS. Hence, for each MS, we need to subtract from the coradius the cost of passing from TS to the other MS. Let us consider $CR^*(MS_a)$; we need to subtract from the coradius the cost of passing from TS to MS_b: this follows from the definition of modified coradius. Hence,

$$CR^*(MS_a) = \frac{N\pi_A + \Pi_A}{\Pi_A + \pi_A} + \frac{N_B\Pi_B - N_A\pi_B + \pi_B}{\Pi_B + \pi_B} - \frac{N_A\Pi_A - N_B\pi_A + \pi_A}{\Pi_A + \pi_A}.$$

Similarly,

$$CR^*(MS_b) = \frac{N\pi_B + \Pi_B}{\Pi_B + \pi_B} + \frac{N_A\Pi_A - N_B\pi_A + \pi_A}{\Pi_A + \pi_A} - \frac{N_B\Pi_B - N_A\pi_B + \pi_B}{\Pi_B + \pi_B}.$$

Note that $CR^*(MS_a) < CR(MS_a)$ and $CR^*(MS_b) < CR(MS_b)$.

Conditions for stochastically stable states.

By comparing all the possibilities, it is possible to verify that if $R(MS_a) > CR(MS_a)$, both $R(MS_b) < CR(MS_b)$ and $R(TS) < CR(TS)$. Similar for $R(MS_b) > CR(MS_b)$ or $R(TS) > CR(TS)$. When $R(MS_a) \leq CR(MS_a)$, $R(MS_b) \leq CR(MS_b)$, and $R(TS) \leq CR(TS)$, we need to use Modified Coradius. Given that $CR(TS) = CR^*(TS)$ it will never be that $R(TS) > CR^*(TS)$. We can show that when $R(MS_a) > CR^*(MS_a)$, then $R(MS_b) < CR^*(MS_b)$ and vice-versa.

When $R(MS_a) = CR^*(MS_a)$, it is also possible that $R(MS_b) = CR^*(MS_b)$. Thanks to Theorem 3 in Ellison (2000), we know that either both states are stochastically stable, or neither of the two is. Note that for the ergodicity of our process, the second case is impossible; hence, it must be that when both $R(MS_a) = CR^*(MS_a)$ and $R(MS_b) = CR^*(MS_b)$, both MS_b and MS_a are stochastically stable. □

Proof of Lemma 6. Recall from Section 2.3 that $\omega^R = \{PS, MS\}$. Firstly, if $\frac{1}{N-1} \max\{\pi_A, \pi_B\} < c < \frac{1}{N-1} \min\{\Xi_{PS}\}$, TS is not an absorbing state (see Corollary 2), all PS are absorbing states (see Corollary 1), and MS are absorbing states (see Table A1). Secondly, consider the set $M = \{PS_a, PS_b, MS_a, MS_b\}$ and the set $\omega^R \setminus M$ containing all the PS not in M. If $R(M) > R(\omega^R \setminus M)$, then the stochastically stable state must be in M. Since the level of the cost is not fixed, the radius of these two sets depend on the cost level. Following the same logic as in Theorem 2 but computing the result as a function of c, we can calculate the two Radii:

$$R(M) = \min\left\{\frac{N_A\pi_B + c(N-1)}{\Pi_A + \pi_A}, \frac{N_B\pi_A + c(N-1)}{\Pi_A + \pi_A}, \frac{N\Pi_A + \pi_A}{\Pi_A + \pi_A}, \frac{N\Pi_B + \pi_B}{\Pi_B + \pi_B}\right\}$$

$$R(\omega^R \setminus M) = \min\left\{\frac{N_B\pi_B - c(N-1) + \Pi_B}{\Pi_B + \pi_B}, \frac{N_B\pi_A - c(N-1)}{\Pi_A + \pi_A}, \frac{N_A\pi_B - c(N-1)}{\Pi_B + \pi_B}\right\}.$$

By comparing all the twelve possibilities case by case, it is possible to show that for every value of payoffs, group size, and cost, $R(M) > R(\omega^R \setminus M)$. Therefore the stochastically stable state must be in the set M. $\square$

Appendix B

Chapter 2 Appendix

B.1 Proofs

Before proceeding with the proofs, I clarify some terms and phrases that will be recurrent throughout the proofs. First of all, referring to Young (1993a), the resistance between two DS is the minimum fraction of mistakes to go from one DS to another one (the reader can interpret this term as the cost for going from one DS to another, such a cost is weighted by the fraction of errors necessary to reach one DS starting from the other one). Secondly, I will often use the phrase "induce X to play Y"; such a phrase refers to the fraction of mistakes necessary to make Y the best reply strategy for type X, starting from a DS where Y was not the best reply for X. This fraction often coincides with the resistance between two DS.

The following lemma states the results fro the unperturbed dynamics part.

Lemma 12.

If conflict is mild,

- $\forall \alpha(t) \in [0, \frac{1}{2})$, *$HM$ is the only stable distribution of strategies.*
- $\forall \alpha(t) \in [\frac{1}{2}, \frac{v}{c})$, *$HM$ and DH are the only stable distributions of strategies.*
- $\forall \alpha(t) \in [\frac{v}{c}, 1]$, *$HD$ and DH are the only stable distributions of strategies.*

If conflict is harsh,

- $\forall \alpha(t) \in [0, \frac{1}{2})$, *$DM$ is the only stable distribution of strategies.*
- $\forall \alpha(t) \in [\frac{1}{2}, \frac{c-v}{c})$, *$DM$ and HD are the only stable distributions of strategies.*

- $\forall \alpha(t) \in [\frac{c-v}{c}, 1]$, *$HD$ and DH are the only stable distributions of strategies.*

Proof of Lemma 12.

I prove the results for $\frac{v}{c} > \frac{1}{2}$, and the results for $\frac{v}{c} < \frac{1}{2}$ hold for symmetry in payoffs. I proceed to prove the result by steps. Firstly, I show the absorbing DS for $\alpha \in [0, \frac{1}{2})$, secondly, I show the absorbing DS in $\alpha \in [\frac{1}{2}, \frac{v}{c})$, and thirdly, I show the absorbing DS for $\alpha \in [\frac{v}{c}, 1]$.

$\alpha \in [0, \frac{1}{2})$

To show that HM is stable, firstly, acknowledge that when $n_H(\tau) = \frac{v}{c}$, according to Equation (3.1), $U_M(H, \frac{v}{c}) = U_M(D, \frac{v}{c})$. Secondly, note that if $n_H(\tau) = \frac{v}{c}$, $U_C(H, \frac{v}{c}) > U_C(D, \frac{v}{c})$. Therefore, HM is a fixed point. That is, if the system ends up in HM, it will stay there. It is easy to prove referring to Equation (3.1), that HM is unique for this range of levels of α.

To see that this DS is also absorbing, consider to start in a DS s.t. $n_H(0) = \frac{v}{c} + \epsilon$, with $n_H^C(0) = 1$. In this DS, $U_C(H, \frac{v}{c} + \epsilon) > U_C(D, \frac{v}{c} + \epsilon)$, and $U_M(D, \frac{v}{c} + \epsilon) > U_M(H, \frac{v}{c} + \epsilon)$ for all $\epsilon > 0$. This means that if conformists are given the revision opportunity, they do not change their strategy, while if M types are given the revision opportunity, they choose H. In other words, with probability 1, $n_H(m) \leq \frac{v}{c} + \epsilon$ with $m > 0$, and for $m' > m$ big enough, $n_H(m') = \frac{v}{c}$ with probability 1. A similar reasoning can be made for a DS $n_H(0) = \frac{v}{c} - \epsilon$, with $n_H^C(0) = 1$.

To see why, the fraction of M types playing H is exactly $\frac{v-\alpha c}{(1-\alpha)c}$, consider the case where $n_H^C(\tau) = 1$. In this case, $n_H(\tau) \geq \alpha$: since $\alpha < \frac{1}{2} < \frac{v}{c}$, the fraction of hawks grows until it reaches $\frac{v}{c}$, and since all conformists are already playing H, myopic best repliers will be the one fulfilling this hole. Since conformists choose deterministically, we know that when $n_H(\tau) = \frac{v}{c}$, α of it will be conformists; therefore, the remaining part must be myopic best repliers, that is $\frac{v}{c} - \alpha = \frac{v-\alpha c}{c}$. This number is the fraction in the population of myopic best repliers playing H. To see the fraction of hawks within M types, it is sufficient to divide $\frac{v-\alpha c}{c}$ by $1 - \alpha$, obtaining $\frac{v-\alpha c}{(1-\alpha)c}$.

$\alpha \in [\frac{1}{2}, \frac{v}{c})$

HM is still a stable DS, and the argument is the same for the case $\alpha \in [0, \frac{1}{2})$. Moreover, When $\alpha > \frac{1}{2}$ and $n_H(\tau) = 1 - \alpha$, referring to Equation (3.1), $U_C(D, n_H(\tau)) > U_C(H, n_H(\tau))$, and $U_M(H, n_H(\tau)) > U_M(D, n_H(\tau))$. Hence, all agents are playing a best reply strategy to the current DS, i.e. DH is a fixed point.

To show that DH is also absorbing, consider to start in a state $n_H(0) = 1 - \alpha - \epsilon$, where $n_H^C(0) = 0$. In this case, $U_C(D, 1 - \alpha - \epsilon) > U_C(H, 1 - \alpha - \epsilon)$, and $U_M(H, 1 - \alpha - \epsilon) >$

$U_M(D, 1-\alpha-\epsilon)$, for all $\epsilon > 0$. Therefore, $n_H(m) \geq 1-\alpha-\epsilon$, for $m > 0$, with probability 1, and for $m' > m$, $n_H(m') = 1-\alpha$ with probability 1. A similar reasoning can be done for $n_H(0) = 1-\alpha+\epsilon$ and $n_H^C(0) = 0$.

$\alpha \in [\frac{v}{c}, 1]$

DH is still a stable DS as for the case $\alpha \in [\frac{1}{2}, \frac{v}{c})$, with the same argument.

If $\alpha > \frac{v}{c}$, when $n_H(\tau) = \alpha$, $U_C(H, n_H(\tau)) > U_C(D, n_H(\tau))$ and $U_M(D, n_H(\tau)) > U_M(H, n_H(\tau))$ (again, see Equation (3.1)). Therefore, whenever the system ends up in HD it will stay there. Intuitively if $\alpha > \frac{v}{c}$, $\frac{v-\alpha c}{(1-\alpha)c}$ is a negative number, that means that if $\alpha > \frac{v}{c}$, and if all conformists choose H, myopic best repliers best respond by playing a pure strategy (D) and not a mixed strategy. To show that HD is also absorbing, the argument is the same used for DH thanks to symmetry in payoffs.

□

A small note on type monomorphic DS and mixed DS. When the fraction of C types is exactly $\frac{v}{c}$, then HD coincides with HM. This is due to the fact that $\frac{v-\alpha(t)c}{c} = \frac{v}{c}$ if $\alpha(t) = \frac{v}{c}$. A similar reasoning applies for DM and DH. I formalize the result in the corollary below.

Corollary 9. *If $\alpha(t) = \frac{v}{c}$, $HD = HM$, and if $\alpha(t) = \frac{c-v}{c}$, $DH = DM$.*

The proof of Corollary 9 is straightforward and therefore is omitted.

Proof of Lemma 7.

I prove the results for $\frac{v}{c} > \frac{1}{2}$, and the results for $\frac{v}{c} < \frac{1}{2}$ hold for symmetry in payoffs. I prove the results in three steps. Firstly, I show the stochastically stable DS for $\alpha \in [0, \frac{1}{2})$, secondly, I show the stochastically stable DS in $\alpha \in [\frac{1}{2}, \frac{v}{c})$, and thirdly, I show the stochastically stable DS for $\alpha \in [\frac{v}{c}, 1]$.

$\alpha \in [0, \frac{1}{2})$

To show the stochastic stability for this range of parameters, it is enough to notice that whenever $\alpha \in [0, \frac{1}{2})$, there is only one absorbing DS, that is HM. Since it is the unique absorbing DS, it is automatically the stochastically stable DS.

$\alpha \in [\frac{1}{2}, \frac{v}{c})$

If $\alpha \in [\frac{1}{2}, \frac{v}{c})$, there are two absorbing DS. To show the stochastic stability of HM, consider the two paths of minimum resistance from HM to DH and from DH to HM.

Starting from DH, it would take $\alpha - \frac{c-v}{c}$ errors to make $U_M(H,\cdot) = U_M(D,\cdot)$. Indeed, consider to be in the state $n_H(\tau) = 1-\alpha+\varepsilon_M$,

$$U_M(H, 1-\alpha+\varepsilon_M) = (1-\alpha+\varepsilon_M)\left(\frac{v-c}{2}\right) + (\alpha-\varepsilon_M)\, v,$$

$$U_M(D, 1-\alpha+\varepsilon_M) = (\alpha-\varepsilon_M)\,\frac{v}{2}.$$

$U_M(H, 1-\alpha+\varepsilon_M) = U_M(D, 1-\alpha+\varepsilon_M) \leftrightarrow \varepsilon_M = \alpha - \frac{c-v}{c}$. With a similar reasoning, it takes $\varepsilon_C \approx \alpha - \frac{1}{2}$ errors to make $U_C(H, 1-\alpha+\varepsilon_C) = U_C(D, 1-\alpha+\varepsilon_C)$ starting from DH (i.e. to convince conformists to play H). Therefore, if this fraction of mistakes is enough to exit from the basin of attraction of DH, the minimum fraction of errors to go from DH to HM must pass through convincing C types since $\alpha - \frac{1}{2} < \alpha - \frac{c-v}{c}$. To show that $\alpha - \frac{1}{2}$ errors is enough to exit from DH's basin of attraction, consider to be in the state such that $n_H(\tau) = \alpha - \varepsilon_C = \frac{1}{2}$. In this DS, with positive probability all conformists are selected to revise the strategy and they choose H with probability $\frac{1}{2}$. Hence, if $n_H(\tau) = \frac{1}{2}$, $n_H(\tau+m) = 1$ with positive probability. Such a DS is out of the basin of attraction of DH since $U_C(H,1) > U_C(D,1)$ and $U_C(D,1) > U_C(H,1)$, i.e. if $n_H(\tau+m) = 1$, with positive probability $\boldsymbol{\sigma}(\tau+m') = HM$, for $m' > m$.

Now consider to start in HM. With a similar reasoning to the one above, it takes $\varepsilon'_C \approx \frac{v}{c} - \frac{1}{2}$ errors to convince C types to play D. With a similar reasoning of the one above, if $n_H(\tau) = \frac{1}{2}$, with positive probability $n_H(\tau+m) = 0$. However, if $n_H(\tau+m) = 0$, $\boldsymbol{\sigma}(\tau+m') = DH$, for $m' > m$ with positive probability since $U_M(H,0) > U_M(D,0)$. In other words, $\frac{v}{c} - \frac{1}{2}$ mistakes are enough to go from HM to DH. To show that this fraction of errors is also the minimum, consider the fraction of mistakes to make $U_M(H,\cdot) > U_M(D,\cdot)$ starting from HM. Such a fraction is $\varepsilon'_M \approx 0$, since it takes only a small fraction of agents playing D to make $U_M(H,\cdot) > U_M(D,\cdot)$. However, let us consider the agents that might deviate. Imagine that a small fraction ε' of M types that was playing H chose D by mistake. In this case, $n_H(\tau) = \frac{v}{c} + \varepsilon'$. If $n_H(\tau) = \frac{v}{c} + \varepsilon'$, $U_M(D, \frac{v}{c} + \varepsilon') > U_M(H, \frac{v}{c} + \varepsilon')$ and $U_C(H, \frac{v}{c} + \varepsilon') > U_C(D, \frac{v}{c} + \varepsilon')$. This means that if myopic best repliers who are playing H are given the revision opportunity, they will choose D. If conformists are given the revision opportunity they will still play H. Uniting the two above arguments, we can conclude that $\frac{v}{c} + \varepsilon'$ is still in the basin of attraction of HM. Note that ε' can be as big as we like and it will not change the result. With a similar argument, if a fraction of myopic best repliers who were playing H and who are given the revision opportunity play D by mistake, a fraction in the next time will choose H as best reply re-balancing the fraction of hawks to $\frac{v}{c}$. Therefore, the only path to go out from the basin of attraction of HM is the one that convinces conformists to play D, i.e. the one with $\frac{v}{c} - \frac{1}{2}$ errors.

To conclude this part of the proof, the minimum fraction of errors to go from DH to HM is $\alpha - \frac{1}{2}$, and the minimum fraction of errors to go from HM to DH is $\frac{v}{c} - \frac{1}{2}$. Since $\frac{v}{c} > \alpha$, this means that the path of minimum resistance to reach DH is bigger than the path of minimum resistance to reach HM. Therefore, HM is the stochastically stable DS if $\alpha \in [\frac{1}{2}, \frac{v}{c})$.

$\alpha \in [\frac{v}{c}, 1]$

By Lemma 12, in this range of values of α, there are two absorbing DS: HD and DH. To show the minimum fraction of errors to go from one DS to the other, consider first the case of myopic best repliers. With a similar argument to the one for $\alpha \in [\frac{1}{2}, \frac{v}{c})$, it is easy to show that whenever we start in DH, if more myopic best repliers play D, they will play H in the following round, as well as if we start in HD and more myopic best repliers play H, they will play D in the following round. Therefore, in both cases, whatever fraction of myopic best repliers play by mistake, it is not sufficient to exit form the basins of attraction of the two DS.

Therefore, we have to consider errors played by conformists in both cases. Let us start from DH. In this case, as I showed for the previous case, the minimum fraction of errors to go from DH to HD is approximately $\alpha - \frac{1}{2}$. Starting from HD, the minimum fraction of mistakes to convince conformists playing D is again approximately $\alpha - \frac{1}{2}$. However, the minimum fraction of mistakes necessary to make myopic best repliers play H starting from HD is $\alpha - \frac{v}{c}$.

To prove this result is enough to equate the two following equations and solve for x.

$$U_M(H, \alpha - x) = (\alpha - x)\left(\frac{v-c}{2}\right) + (1 - \alpha + x)\, v,$$

$$U_M(D, \alpha - x) = (1 - \alpha + x)\, \frac{v}{2}.$$

To show if this fraction is enough to exit from the basin of attraction of HD, let us consider the DS such that $n_H(\tau) = \alpha - (\alpha - \frac{v}{c}) = \frac{v}{c}$. In this DS, all conformists that are playing D (by mistake) would change strategy and play H if given the revision opportunity. While all myopic best repliers are indifferent between H and D. At this point, if another small fraction of conformists selects D, then $U_M(H, \cdot) > U_M(D, \cdot)$. This means that if all myopic best repliers could change the strategy, they would choose H. However, if that is the case, we will end up in a DS where $\frac{v}{c}-\epsilon+1-\alpha$ agents are playing H. In such a DS, all conformists who are not playing H would choose H if given the revision opportunity, and all myopic best repliers will choose D if given the revision opportunity. In other words we will return to HD with positive probability.

Hence, the only way to go out of the basin of attraction of HD and DH is through conformists. The minimum fraction of errors to leave the basin of attraction of both HD and DH is approximately $\alpha - \frac{1}{2}$. Now, I show in detail how to compute the exact number.

If we start from DH and a fraction of conformists $\varepsilon_C^{''}$ plays H by mistake the utilities from playing D and H for C types are defined by the following equations:

$$U_C(H, 1-\alpha+\varepsilon_C^{''}) = \lambda_C \left((1-\alpha+\varepsilon_C^{''}) + \frac{v-c}{2}(\alpha-\varepsilon_C^{''})v \right) + (1-\lambda_C)(1-\alpha+\varepsilon_C^{''}),$$

$$U_C(D, 1-\alpha+\varepsilon_C^{''}) = \lambda_C \left((\alpha-\varepsilon_C^{''})\frac{v}{2} \right) + (1-\lambda_C)(\alpha-\varepsilon_C^{''}).$$

From the above equations it can be retrieved that $U_C(H, 1-\alpha+\varepsilon_C^{''}) > U_C(D, 1-\alpha+\varepsilon_C^{''})$ iff

$$\varepsilon_C^{''} > \frac{\alpha c\lambda_C + 4\alpha\lambda_C - 4\alpha - c\lambda_C + \lambda_C v - 2\lambda_C + 2}{c\lambda_C + 4\lambda_C - 4} \tag{B.1}$$

Similarly, starting from HD, if a fration fo conformists $\varepsilon_C^{'''}$ plays D by mistake,

$$U_C(H, \alpha-\varepsilon_C^{'''}) = \lambda_C \left((\alpha-\varepsilon_C^{'''})\frac{v-c}{2} + (1-\alpha-\varepsilon_C^{'''})v \right) + (1-\lambda_C)(\alpha-\varepsilon_C^{'''}),$$

$$U_C(H, \alpha-\varepsilon_C^{'''}) = \lambda_C \left((1-\alpha-\varepsilon_C^{'''})\frac{v}{2} \right) + (1-\lambda_C)(1-\alpha+\varepsilon_C^{'''}),$$

$U_C(H, 1-\alpha+\varepsilon_C^{''}) > U_C(D, 1-\alpha+\varepsilon_C^{''})$ iff

$$\varepsilon_C^{''} > \frac{\alpha c\lambda_C + 4\alpha\lambda_C - 4\alpha - \lambda_C v - 2\lambda_C + 2}{c\lambda_C + 4\lambda_C - 4}. \tag{B.2}$$

Comparing (B.1) with (B.2), it emerges that the RHT of (B.1) is greater than the RHT of (B.2) if and only if $\frac{v}{c} > \frac{1}{2}$. This passage concludes the proof, as it shows that the fraction of mistakes to go from DH to HD is higher than the one from HD to DH whenever conflict is mild. In other words, DH is stochastically stable for $\alpha \in [\frac{v}{c}, 1]$ if and only if conflict is mild ($\frac{v}{c} > \frac{1}{2}$).

□

Proof of Lemma 8.

I prove the results for $\frac{v}{c} > \frac{1}{2}$, and the results for $\frac{v}{c} < \frac{1}{2}$ hold for symmetry in payoffs. I prove the result in two steps. Firstly, I show the the evolution of α for $\alpha \in [0, \frac{v}{c})$, and secondly, I show the the evolution of α for $\alpha \in [\frac{v}{c}, 1]$.

$\alpha \in [0, \frac{v}{c})$

By Lemma 7, we know that whenever $\alpha \in [0, \frac{v}{c})$, HM is the stochastically stable DS. If that is the case,

$$\bar{\Pi}_C\,(HM) = \frac{v}{c}\frac{v-c}{2} + \frac{c-v}{c}v,$$

$$\bar{\Pi}_M\,(HM) = \frac{v-\alpha c}{(1-\alpha)c}\left(\frac{v}{c}\frac{v-c}{2} + \frac{c-v}{c}v\right) + \frac{c-v}{(1-\alpha)c}\left(\frac{c-v}{c}\frac{v}{2}\right) - \kappa.$$

With some calculus it can be shown that $\bar{\Pi}_C - \bar{\Pi}_M = \kappa$. Therefore, referring to Equation (3.3),

$$\Delta\alpha\,(t) = \alpha(t)\,(1-\alpha(t))\,\kappa.$$

In other words, if $\alpha \in [0, \frac{v}{c})$, $\Delta\alpha\,(t) > 0$, $\forall \kappa > 0$.

$\alpha \in [\frac{v}{c}, 1]$

By Lemma 7, we know that DH is the stochastically stable DS in this case. Therefore,

$$\bar{\Pi}_C\,(DH) = \alpha(\frac{v}{2})$$

$$\bar{\Pi}_M\,(DH) = \alpha v + (1-\alpha)\left(\frac{v-c}{2}\right) - \kappa.$$

With some calculus it can be shown that $\bar{\Pi}_C\,(DH) - \bar{\Pi}_M\,(DH) = \frac{1}{2}\,((1-\alpha)c + 2\kappa - v)$. Hence,

$$\Delta\alpha\left(t\right)=\alpha(t)\left(1-\alpha(t)\right)\frac{1}{2}\left((1-\alpha(t))c+2\kappa-v\right),$$

and $\Delta\alpha\left(t\right)=0$ if and only if $\alpha(t)=\frac{c+2\kappa-v}{c}=\overline{\alpha}$. This last passage completes the proof.

□

Proof of Corollary 3.

The proof is straightforward from the proof of Lemma 8. Indeed, consider $\overline{\alpha}=\frac{c+2\kappa-v}{c}$. $\frac{c+2\kappa-v}{c}=\frac{v}{c}\leftrightarrow\kappa=\frac{2v-c}{2}$. Since $\Delta\alpha\left(t\right)>0,\forall\alpha(t)\in[0,\frac{v}{c})$, if $\kappa\leq\frac{2v-c}{2}$, $\overline{\alpha}=\frac{v}{c}$. Moreover, $\frac{c+2\kappa-v}{c}>\frac{v}{c}$, iff $\kappa>\frac{2v-c}{2}$, and therefore, it must be that $\overline{\alpha}>\frac{v}{c}$ iff $\kappa>\frac{2v-c}{2}$. The same argument applies for the case of $\frac{v}{c}<\frac{1}{2}$ for symmetry in payoffs. This completes the proof.

□

Proof of Corollary 7.

I prove the result for $\frac{v}{c}>\frac{1}{2}$, and the argument stands for $\frac{v}{c}>\frac{1}{2}$ for symmetry in payoffs.

$\alpha\in[\frac{1}{2},\frac{v}{c})$

From Lemma 12, we know that if conflict is mild and $\alpha\in[\frac{1}{2},\frac{v}{c})$ there are two possible absorbing distribution of strategies DH and HM. Therefore, I compute the fitness of each type assuming that they will be half of the times in DH and half of the times in HM.

$$\Pi_C\left(\frac{1}{2}HM+\frac{1}{2}DH\right)=\frac{1}{2}\left(\frac{v}{c}\frac{v-c}{2}+\frac{c-v}{c}v\right)+\frac{1}{2}\left(\alpha\frac{v}{2}\right),$$

$$\Pi_M\left(\frac{1}{2}HM+\frac{1}{2}DH\right)=\frac{1}{2}\left(\frac{v-\alpha c}{(1-\alpha)c}\left(\frac{v}{c}\frac{v-c}{2}+\frac{c-v}{c}v\right)+\frac{c-v}{(1-\alpha)c}\left(\frac{c-v}{c}\frac{v}{2}\right)\right)$$
$$+\frac{1}{2}\left(\alpha v+(1-\alpha)\frac{v-c}{2}\right)-\kappa.$$

From the two equations above it emerges that $\overline{\alpha}=\frac{c-v}{c}+\frac{4}{c}\kappa$. The corresponding level of κ such that $\alpha>\frac{1}{2}$ and $\overline{\alpha}=\frac{v}{c}$ are easily calculated, and they are $k_1=\frac{2v-c}{8}$ and $k_2=\frac{2v-c}{4}$.

$\alpha\in[\frac{v}{c},1]$

From Lemma 12, we know that if conflict is mild and $\alpha\in[\frac{v}{c},1]$ there are two possible absorbing distribution of strategies DH and HD. Therefore, I compute the fitness of each type assuming that they will be half of the times in DH and half of the times in HD.

$$\Pi_C\left(\frac{1}{2}HD+\frac{1}{2}DH\right)=\frac{1}{2}\left(\alpha\frac{v-c}{2}+(1-\alpha)v\right)+\frac{1}{2}\left(\alpha\frac{v}{2}\right),$$

$$\Pi_M\left(\frac{1}{2}HD+\frac{1}{2}DH\right)=\frac{1}{2}\left((1-\alpha)\frac{v}{2}\right)+\frac{1}{2}\left(\alpha v+(1-\alpha)\frac{v-c}{2}\right)-\kappa,$$

The above equations equated bring out an $\overline{\alpha}=\frac{1}{2}+\frac{2}{c}\kappa$. Therefore, κ s.t. $\overline{\alpha}>\frac{v}{c}$ is $\kappa_2=\frac{2v-c}{4}$. Comparing results for $\alpha\in[\frac{1}{2},\frac{v}{c})$ and $\alpha\in[\frac{v}{c},1]$, we obtain the thresholds in the corollary. This last passage completes the proof.

□

Proof of Theorem 6.

The proof of this theorem is built upon the arguments of previous proofs. Let us consider the matrix $G_3=pG_1+(1-p)G_2$. The payoffs of this matrix are as follows: $\pi(H,H)=\frac{pv_1+(1-p)v_2}{c}$, $\pi(H,D)=pv_1+(1-p)v_2$, $\pi(D,H)=0$, and $\pi(H,D)=\frac{pv_1+(1-p)v_2}{2}$. Given these payoffs, we can call $v_3=pv_1+(1-p)v_2$, and the payoffs of G_3 become $\pi(H,H)=\frac{v_3}{c}$, $\pi(H,D)=v_3$, $\pi(D,H)=0$, and $\pi(H,D)=\frac{v_3}{2}$. Given these payoffs, the proof is straightforward and it follows from Lemma 12, 7 and 8 and Theorem 5. The results depend on $\frac{v_3}{c}$, namely on $\frac{pv_1+(1-p)v_2}{c}$.

□

Before I move forward to the proof of Theorem 7, I introduce some useful DS. I name $(1,0,\frac{v-\alpha c}{(1-\alpha)c},\frac{c-v}{(1-\alpha)c})$ HM, $(0,1,1,0)$ DH, $(1,0,0,0)$ HB, and $(1,0,0,1)$ HD. Intuitively, HM is the mixed DS where conformists play H and myopic best repliers play a mixed strategy, DH is a type monomorphic DS such that conformists play D and myopic best repliers play H, HB is such that conformists play H and myopic best repliers play B, and finally HD is such that conformists play H and myopic best repliers play D.

Note that for simplicity of exposition I will use $n_B(\tau)$ instead of $1-n_H(\tau)-n_D(\tau)$.

Proof of Theorem 7.

I prove this theorem in a similar way of Theorem 5. I prove the results for $\frac{v}{c}>\frac{1}{2}$, and the ones for $\frac{v}{c}<\frac{1}{2}$ stands for symmetry in payoffs. I prove the results by lemmas, to be consistent with the main part of the paper and to (hopefully) ease the understanding of the proof to the reader.

Lemma 13.

For $\frac{v}{c} > \frac{1}{2}$,

- *if* $\alpha \in [0, \frac{1}{2})$, BB *and* HM *are the unique stable DS.*
- *If* $\alpha \in [\frac{1}{2}, \frac{v}{c})$, BB, HM, DH *and* HB *are the unique stable DS.*
- *If* $\alpha \in [\frac{v}{c}, 1]$, BB, DH *and* HD *are the unique stable DS.*

The proof for this lemma resembles the arguments of the proof of Lemma 12. I divide the proof in three parts: $\alpha \in [0, \frac{1}{2})$, $\alpha \in [\frac{1}{2}, \frac{v}{c})$, and $\alpha \in [\frac{v}{c}, 1]$.

$\alpha \in [0, \frac{1}{2})$

Firstly, it can be shown from Equation (3.7), that whenever $n_H(\tau) = \frac{v}{c}$, and $n_D(\tau) = \frac{c-v}{c}$ $U_M(H, \frac{v}{c}, \frac{c-v}{c}) = U_M(D, \frac{v}{c}, \frac{c-v}{c}) > U_M(B, \frac{v}{c}, \frac{c-v}{c})$, and $U_C(H, \frac{v}{c}, \frac{c-v}{c}) > U_C(D, \frac{v}{c}, \frac{c-v}{c}) > U_C(B, \frac{v}{c}, \frac{c-v}{c})$. Hence, whenever the system end up in HM, it stays there. Similarly, whenever $1 - n_H(\tau) - n_D(\tau) = 1$, $U_M(B, 0, 0) > U_M(H, 0, 0) > U_M(D, 0, 0)$, and $U_C(B, 0, 0) > U_C(H, 0, 0) \approx U_C(D, 0, 0)$. Therefore, in BB all agents are playing the best reply to the current DS, that is BB is a fixed point.

Let us start from proving that BB is absorbing. Consider to be in a DS such that $n_H(\tau) = \epsilon$, and $n_D(\tau) = 0$. In this case, $U_C(B, \epsilon, 0) > U_C(H, \epsilon, 0) > U_C(D, \epsilon, 0)$, for all $\epsilon < \frac{1}{2}$. Moreover, B is the myopic best reply strategy for M types, as long as $\epsilon < \frac{v}{c}$. A similar reasoning applies for $n_H(\tau) = 0$ and $n_D(\tau) = \epsilon$. From the above reasoning, we can conclude that there is a set of DS such that converges to BB with probability 1.

For HM, think about a state $n_H(\tau) = \frac{v}{c} - \epsilon$ and $n_D(\tau) = \frac{c-v}{c} + \epsilon$. Again H is the best reply for C types as long as $\epsilon < \frac{2v-c}{2c}$. For M types, H is the best reply, $\forall \epsilon > 0$. This means that the system will converge to HM with probability 1, from all DS of the same kind as described above. A similar reasoning can be applied to the DS such that $n_H(\tau) = \frac{v}{c} + \epsilon$ and $n_D(\tau) = \frac{c-v}{c} - \epsilon$, where H is still the best reply for all C types, and D is the best reply for all M types for all $\epsilon > 0$.

$\alpha \in [\frac{1}{2}, \frac{v}{c})$

The argument such that BB and HM are still stable is the same as for $\alpha \in [0, \frac{1}{2})$, and therefore, is omitted.

To see that DH is a fixed point, note that when $n_H(\tau) = 1 - \alpha$, and $n_D(\tau) = \alpha$, $U_C(D, 1 - \alpha, \alpha) > U_C(H, 1 - \alpha, \alpha) > U_C(B, 1 - \alpha, \alpha)$, and $U_M(H, 1 - \alpha, \alpha) > U_M(B, 1 - \alpha, \alpha) > U_M(D, 1 - \alpha, \alpha)$, referring to Equation (3.7). Therefore, whenever the system ends up in

DH it will stay there. Similarly, referring to Equation (3.7), if $n_H(\tau) = \alpha$ and $n_D(\tau) = 0$, $U_C(H, \alpha, 0) > U_C(B, \alpha, 0) > U_C(D, \alpha, 0)$, and $U_M(B, \alpha, 0) > U_M(D, \alpha, 0) > U_M(H, \alpha, 0)$. Again, this means that in HB every agent is playing the myopic best reply to the current DS, and hence, if the system ends up in HB, it will stay there.

Now, I show that these DS are also absorbing. First of all, let us consider DH. If $n_D(\tau) = \alpha + \epsilon$, and $n_H(\tau) = 1 - \alpha - \epsilon$, surely D is still the best reply for C types $\forall \epsilon > 0$. Moreover, also H is still the best reply for M types, $\forall \epsilon > 0$. Therefore, whenever the system is near DH as described above, it converges to DH with probability 1. Secondly, let us consider HB. If $n_H(\tau) = \alpha + \epsilon$, and $n_D(\tau) = 0$, H is the best reply for all C types $\forall \epsilon > 0$. Furthermore, B is still the best reply for all M types $\forall \epsilon < \frac{v}{c} - \alpha$. This last passage proves that there is a set of DS that converges with probability 1 to HB, i.e. HB is an absorbing DS.

$\alpha \in [\frac{v}{c}, 1]$

The argument such that BB and DH are still stable is the same as for $\alpha \in [0, \frac{v}{c})$, and therefore, is omitted.

Let us now consider HD, similar as for the other DS, it can be shown using Equation (3.7), that if $n_H(\tau) = \alpha$ and $n_D(\tau) = 1 - \alpha$, $U_C(H, \alpha, 1 - \alpha) > U_C(D, \alpha, 1 - \alpha) > U_C(B, \alpha, 1 - \alpha)$, and $U_M(D, \alpha, 1 - \alpha) > U_M(B, \alpha, 1 - \alpha) > U_M(H, \alpha, 1 - \alpha)$. Hence, all agents are playing the best reply at HD, that is, HD is a fixed point. Similar to what was done above, consider to be in a DS such that $n_H(\tau) = \alpha + \epsilon$. Conformists would still play H, and for myopic best repliers, D is still the best reply $\forall \epsilon > 0$. In other words, there is a set of DS that converges to HD with probability 1, i.e. HD is absorbing.

Lemma 14. *If* $\frac{v}{c} > \frac{1}{2}$, $\forall \alpha \in [0, 1]$, *$BB$ is uniquely stochastically stable.*

I divide the proof in three parts: $\alpha \in [0, \frac{1}{2})$, $\alpha \in [\frac{1}{2}, \frac{v}{c})$, and $\alpha \in [\frac{v}{c}, 1]$.

$\alpha \in [0, \frac{1}{2})$

For this range of values, only BB and HM are absorbing. However, it is easy to show that HM is hardly stable (by means of stochastic stability). Indeed, it is sufficient that a small fraction of agents that was playing H play B by mistake, to make M types choose B with positive probability. Indeed, consider the DS $n_H(\tau) = \frac{v}{c} - \varepsilon$, $n_D(\tau) = \frac{c-v}{c}$, with $\varepsilon \approx 0$. Using Equation (3.7), it can be shown that $U_M(B, \frac{v}{c} - \varepsilon, \frac{c-v}{c}) = U_M(H, \frac{v}{c} - \varepsilon, \frac{c-v}{c}) > U_M(D, \frac{v}{c} - \varepsilon, \frac{c-v}{c})$. If this is the case, it means that if $n_H(\tau) = \frac{v}{c} - \epsilon$, $n_D(\tau) = \frac{c-v}{c}$, we can reach a DS such that $n_H(\tau + m) = \alpha$, $n_D(\tau + m) = 0$, $m > 0$, with zero resistance. In such a DS, B is the best reply for all M types and also for all C types, that is, we can reach a DS such that $n_B(\tau + m') = 1$, for $m' > m$ with zero resistance. In other words, it is sufficient a small deviation of size ε to exit from the basin of attraction of HM and enter in the basin of attraction of BB.

Furthermore, it can be shown using Equation (3.7), that there exists no small deviation of size $\varepsilon \approx 0$ sufficient to leave the basin of attraction of BB. Therefore, we can conclude that BB is the uniquely stochastically stable DS $\forall \alpha \in [0, \frac{1}{2})$.

$\alpha \in [\frac{1}{2}, \frac{v}{c})$

For this range of values of α, there are four absorbing DS. However, as it was shown before, it is sufficient a small deviation of size $\varepsilon \approx 0$ to exit from the basin of attraction of HM. This means that HM can never be the stochastically stable DS. However it should not be removed from the analysis. Indeed, the path of minimum resistance from one absorbing DS to another one, may pass through HM as well. I will call an indirect path from one absorbing DS to another, a path that passes through HM. I will call a direct path from one absorbing DS to another, a path that does not passes through HM. For each absorbing DS I will provide the path of minimum resistance from that absorbing DS to each of the other absorbing DS. In the end, I calculate the stochastic potential of each absorbing DS in the way of Young (1993a). From now on, I call $r(\boldsymbol{\sigma'}, \boldsymbol{\sigma''})$ the path of minimum resistance from $\boldsymbol{\sigma'}$ to $\boldsymbol{\sigma''}$. Moreover, I call $r_d(\boldsymbol{\sigma'}, \boldsymbol{\sigma''})$ the direct path of minimum resistance and $r_{id}(\boldsymbol{\sigma'}, \boldsymbol{\sigma''})$ the indirect path of minimum resistance.

Let us start from the paths from BB to DH and HB.

$r(BB, DH)$

Firstly, consider the direct path. The minimum fraction of mistakes, to convince C types to play D starting from BB is $\frac{1}{2}$. Indeed, when 100% of the population plays B, it takes half of the population playing D by mistake to make D the best reply for conformists. The minimum fraction of mistakes necessary to convince M types to play H starting from BB is $\frac{c-v}{c}$. This fraction of mistakes is sufficient to exit the basin of attraction of BB under one condition. Indeed, suppose $n_B(0) = 1$, and that after m steps $n_B(m) = \frac{v}{c}$, and $n_D(m) = \frac{c-v}{c}$. This DS is reached with a resistance of $\frac{c-v}{c}$. Suppose that all the agents playing D at this DS are conformists. In this DS, H is the best reply for myopic best repliers; hence, with positive probability after other m' steps, $n_H(m + m') = 1 - \alpha\ n_D(m) = \frac{c-v}{c}$. In this DS, $n_H > n_D$, but $n_H > n_B$ iff $\alpha < \frac{2c-v}{2c}$. Furthermore, $n_D > n_B$ iff $\alpha < \frac{2(c-v)}{c}$. From the above statements, it emerges that the direct path of minimum resistance from BB to DH depends on three situations. The first is when $n_H > n_D > n_B$, the second is when $n_H > n_B > n_D$, and the third is when $n_B > n_H > n_D$. In the first case, it is sufficient that $\frac{v}{c} - \alpha$ conformists that were choosing B choose D by mistake to make D the best reply for all conformists (still H the best reply for all myopic best repliers). In other words, an additional resistance of $\frac{v}{c} - \alpha$ is needed to reach DH. This means that in the first case, $r(BB, DH) = 1 - \alpha$. In the second case, the reasoning is similar, and the minimum path from BB to DH is still $1 - \alpha$. In the third case, the minimum fraction of errors needed to reach DH is such that D becomes the best reply for C types instead of B (since $n_B > n_H > n_D$). This happens with a resistance of $\frac{\alpha}{2} - \frac{c-v}{c}$ in case $\alpha > \frac{2}{3}$ and of $\frac{2-\alpha}{4} - \frac{c-v}{c}$ otherwise. Consequently, in the third case, the path of minimum resistance from BB to DH is $\frac{\alpha}{2}$ or $\frac{2-\alpha}{4}$. Wrapping up,

$$r_d(BB, DH) \approx \begin{cases} 1 - \alpha & \text{if } \frac{v}{c} < \frac{3}{4} \\ \frac{\alpha}{2} & \text{if } \frac{v}{c} > \frac{3}{4} \wedge \alpha > \frac{2}{3} \\ \frac{2-\alpha}{4} & \text{if } \frac{v}{c} > \frac{3}{4} \wedge \alpha < \frac{2}{3} \end{cases} . \tag{B.3}$$

To show that the direct path has a lower resistance than the indirect one, let us consider the path of minimum resistance from BB to HM. Such a path is similar to the one from BB to DH; indeed, it passes through mistakes of conformists to convince M types to play H. We know from the above part of the proof that the minimum fraction of mistakes to make H the best reply for M types is $\frac{c-v}{c}$, and that with zero resistance, then we can reach a DS s.t. $n_B = \alpha - \frac{c-v}{c}$, $n_H = 1 - \alpha$, $n_D = \frac{c-v}{c}$. In this DS, $n_H > n_B$ iff $\alpha < \frac{2c-v}{c}$. Therefore, if $\alpha < \frac{2c-v}{c}$, $\frac{c-v}{c}$ is enough to reach HM, while if $\alpha > \frac{2c-v}{c}$, a fraction of $\alpha - \frac{2c-v}{2c}$ mistakes is necessary to reach HM. Consequently, $r(BB, HM) = \begin{cases} \frac{c-v}{c} & \text{if } \frac{v}{c} < \frac{2}{3} \vee \left(\frac{v}{c} > \frac{2}{3} \wedge \alpha < \frac{2c-v}{2c}\right) \\ \alpha - \frac{v}{2c} & \text{if } \frac{v}{c} > \frac{2}{3} \wedge \alpha > \frac{2c-v}{2c} \end{cases}$.

Moreover, similarly to the proof of Lemma 7, $r(HM, DH) = \frac{v}{c} - \frac{1}{2}$. For this reason, the indirect path from BB to DH is always at least $\frac{1}{2}$. Since the RHT of (B.3) is always lower than $\frac{1}{2}$, we conclude that the direct path from BB to DH must be the one defined in (B.3). Therefore, $r(BB, DH) = r_d(BB, DH)$.

$r(BB, HB)$

The direct path of minimum resistance from BB to HB always passes through convincing conformists to play H, since myopic best repliers still play B in this DS, and hence, it would be inconvenient to induce them to change the strategy to pass from BB to HB. Similar to the case above, the minimum fraction of mistakes to induce C types play H starting from $n_B = 1$ is $\frac{1}{2}$. However, the indirect path of minimum resistance is lower than $\frac{1}{2}$. Indeed, as I showed above, $r(BB, HM) = \begin{cases} \frac{c-v}{c} & \text{if } \frac{v}{c} < \frac{2}{3} \vee \left(\frac{v}{c} > \frac{2}{3} \wedge \alpha < \frac{2c-v}{2c}\right) \\ \alpha - \frac{v}{2c} & \text{if } \frac{v}{c} > \frac{2}{3} \wedge \alpha > \frac{2c-v}{2c} \end{cases}$, and $r(HM, HB) = \varepsilon \approx 0$ as shown earlier in this proof. Therefore,

$$r(BB,HB) = r_{id}(BB,HB) \approx \begin{cases} \frac{c-v}{c} & \text{if } \frac{v}{c} < \frac{2}{3} \vee \left(\frac{v}{c} > \frac{2}{3} \wedge \alpha < \frac{2c-v}{2c}\right) \\ \alpha - \frac{v}{2c} & \text{if } \frac{v}{c} > \frac{2}{3} \wedge \alpha > \frac{2c-v}{2c} \end{cases} . \tag{B.4}$$

$r(DH, HB)$

Concerning the direct path, the minimum fraction of mistakes to go from DH to HB convincing conformists is $\alpha - \frac{1}{2}$. Such resistance is lower than the one which involves myopic best repliers changing their strategy. Indeed, consider being in a DS such that $n_B = x$, $n_H = 1 - \alpha$, and $n_D = \alpha - x$. B becomes best reply for M types when $x > \alpha - \frac{c-v}{c}$, that is the resistance to make D best reply for M types is $\alpha - \frac{c-v}{c} > \alpha - \frac{1}{2}$. Therefore, $r_d(DH, HB) = \alpha - \frac{1}{2}$.

Concerning the indirect path, we know from the proof of Lemma 7 that $r(DH, HM) = \alpha - \frac{1}{2}$, and that $r(HM, HB) = \varepsilon$. therefore, the path of minimum resistance from DH to HB must be the direct one. That is,

$$r(DH, HB) = \alpha - \frac{1}{2}. \tag{B.5}$$

$r(DH, BB)$

Concerning the direct path, with similar calculus as the ones for the last two cases of (B.3), it can be shown that the minimum resistance to make B best reply for conformists is $\frac{\alpha}{2}$ if $\alpha > \frac{2}{3}$ and $\frac{2-\alpha}{4}$ if $\alpha < \frac{2}{3}$. However, as I showed above, the minimum resistance necessary to make B best reply for M types is $\alpha - \frac{c-v}{c}$. Such resistance is also enough to reach BB starting from DH. Indeed, consider to start in a DS s.t. $n_D(0) = \frac{c-v}{c}$ and $n_H(0) = 1 - \alpha$. As stated above, B is the best reply for M types; therefore, after m periods, with positive probability, we will end up in a DS s.t. $n_D(m) = \frac{c-v}{c}$ and $n_H(m) = 0$. However, in this DS, B is also the best reply for C types. This means that with positive probability $n_B(m+m') = 1$, with m'.

Consequently, the direct path of minimum resistance from DH to BB is either $\alpha - \frac{c-v}{c}$ or $\frac{\alpha}{2}$ if $\alpha > \frac{2}{3}$ or $\alpha - \frac{c-v}{c}$ or $\frac{2-\alpha}{4}$ if $\alpha < \frac{2}{3}$.

Concerning the indirect path, we know that $r(DH, HM) = \alpha - \frac{1}{2}$. Moreover, the minimum resistance from HM to BB, always involves HB and myopic best repliers. Indeed, the minimum resistance to make conformists play B is $\frac{v}{c} - \frac{1}{2}$. While considering M types, the minimum resistance passes through convincing M types playing B, which has resistance ε (as for $r(HM, HB)$), and then from HB convincing C types playing B, which has resistance

$\alpha - \frac{1}{2}$. This means that the indirect path of minimum resistance from DH to BB is $2\alpha - 1$, which is greater than all the cases of the direct path. Hence,

$$r(DH, BB) = \begin{cases} \frac{\alpha}{2} & \text{if } \left(\frac{v}{c} > \frac{3}{4} \wedge \alpha > \frac{2(c-v)}{c}\right) \vee \frac{2}{3} < \frac{v}{c} < \frac{3}{4} \\ \frac{2-\alpha}{4} & \text{if } \frac{v}{c} > \frac{7}{8} \vee \left(\frac{2}{3} < \frac{v}{c} < \frac{7}{8} \wedge \alpha > \frac{6c-4v}{3c}\right) \\ \alpha - \frac{c-v}{c} & \text{otherwise} \end{cases} . \tag{B.6}$$

$r(HB, DH)$

Concerning the direct path, with similar calculus as for (B.3), it can be shown that this path is either $\frac{\alpha}{2}$ or $\frac{2-\alpha}{4}$, depending on whether $\alpha > \frac{2}{3}$ or not. Concerning the indirect path, we already know that $r(HM, DH) = \frac{v}{c} - \frac{1}{2}$, while $r(HB, HM)$, can be calculated as the minimum resistance to make H the best reply for M types. This resistance is equal to $\frac{c-v}{c}$ (see $r(BB, HM)$). The sum of these two resistance is equal to $\frac{1}{2}$, which is always greater than the two direct cases. Hence,

$$r(HB, DH) = \begin{cases} \frac{\alpha}{2} & \text{if } \frac{v}{c} > \frac{2}{3} \wedge \alpha > \frac{2}{3} \\ \frac{2-\alpha}{4} & \text{if } \frac{v}{c} < \frac{2}{3} \vee \left(\frac{v}{c} > \frac{2}{3} \wedge \alpha < \frac{2}{3}\right) \end{cases} . \tag{B.7}$$

$r(HB, BB)$

The direct path of minimum resistance from HB to BB, is $\alpha - \frac{1}{2}$. The calculus are the same as for $r(DH, HB)$. The indirect path of minimum resistance cannot be lower since we already know from previous calculus that $r(HM, BB)$; therefore, it must be that

$$r(HB, BB) = \alpha - \frac{1}{2}. \tag{B.8}$$

Once we know these paths, we need to calculate the stochastic potential of each absorbing DS. The stochastic potential is the minimum path to reach each absorbing DS starting from all the other DS (Young, 1993a). Let us call $\rho_{\boldsymbol{\sigma'}}$ the stochastic potential of $\boldsymbol{\sigma'}$. In this case, $\rho_{BB} = \min\{(B.6)+(B.8), (B.5)+(B.8), (B.7)+(B.6)\}$, $\rho_{HB} = \{(B.5) + (B.4), (B.6) + (B.4), (B.3) + (B.5)\}$, $\rho_{DH} = \min\{(B.3) + (B.7), (B.4) + (B.7), (B.8) + (B.3)\}$. With calculations it can be shown that ρ_{BB} is the minimum $\forall \alpha \in [\frac{1}{2}, \frac{v}{c})$.

$\alpha \in [\frac{v}{c}, 1]$

For this range of values of α, HM is no longer stable. Therefore, we only need to calculate the direct paths from one stable DS to the other.

Similar reasoning as for the case of $\alpha \in [\frac{1}{2}, \frac{v}{c})$ lead to

$$r(BB, DH) = \begin{cases} \frac{c-v}{c} & \text{if } \frac{v}{c} < \frac{2}{3} \wedge \alpha < \frac{2(c-v)}{c} \\ \frac{\alpha}{2} & \text{if } \frac{v}{c} > \frac{2}{3} \vee \left(\frac{v}{c} < \frac{2}{3} \wedge \alpha > \frac{2(c-v)}{c}\right) \end{cases}, \tag{B.9}$$

$$r(DH, BB) = \begin{cases} \alpha - \frac{c-v}{c} & \text{if } \frac{v}{c} < \frac{2}{3} \wedge \alpha < \frac{2(c-v)}{c} \\ \frac{\alpha}{2} & \text{if } \frac{v}{c} > \frac{2}{3} \vee \left(\frac{v}{c} < \frac{2}{3} \wedge \alpha > \frac{2(c-v)}{c}\right) \end{cases}, \tag{B.10}$$

$$r(HD, DH) = r(DH, HD) = r(HD, BB) = \alpha - \frac{1}{2}. \tag{B.11}$$

Concerning calculus and cases, $r(BB, HD)$ is the longest to calculate for this part of the proof. Intuitively, the path of minimum resistance cannot path through inducing only conformists to play H since the cost would be $\frac{1}{2}$. Such a path is obtained by passing through inducing M types playing D and then inducing C types playing H. Firstly, with a resistance of $\frac{c-v}{c}$, we can reach one of these two DS such that $n_B = \frac{v}{c}$, $n_H = 1-\alpha$ and $n_D = \frac{c-v}{c} - (1-\alpha)$ or $n_B = \frac{v}{c} - (1-\alpha)$, $n_H = 1-\alpha$ and $n_D = \frac{c-v}{c}$. From these DS, the path of minimum resistance depends on which between n_B, n_H, and n_D at these DS. With some calculations, we can retrieve

$$r(BB, HD) = \begin{cases} \alpha - \frac{v}{2c} & \text{if } \frac{v}{c} > \frac{2}{3} \\ \min\{\alpha - \frac{5}{4}\frac{v}{c} + \frac{1}{2}, \frac{\alpha c + c - v}{2c}\} & \text{if } \frac{v}{c} < \frac{2}{3} \wedge \alpha > \frac{2(c-v)}{c} \\ \frac{\alpha c + 2c - 3v}{2c} & \text{if } \frac{v}{c} < \frac{2}{3} \wedge \left(\frac{4c-3v}{4c} < \alpha < \frac{2(c-v)}{c} \vee \alpha < \frac{2c-v}{2c}\right) \\ \min\{\frac{\alpha + 4\alpha c - 3v}{4c}, \frac{\alpha c + c - v}{2c}\} & \text{if } \frac{2c-v}{2c} < \alpha < \frac{4c-3v}{4c} \end{cases} \tag{B.12}$$

Similar to the previous calculus, we obtain $\rho_{BB} = \min\{$(B.10)+(B.11), (B.11)+(B.11), (B.11)+(B.10)$\}$, $\rho_{HD} = \{$(B.11) + (B.12), (B.10) + (B.12), (B.9) + (B.11)$\}$, $\rho_{DH} = \min\{$(B.9) + (B.11), (B.12) + (B.11), (B.11) + (B.9)$\}$. Again, it can be shown that ρ_{BB} is always the minimum of the three, and therefore, BB is the stochastically stable DS.

Lemma 15. *If* $\frac{v}{c} > \frac{1}{2}$, $\forall \kappa > 0$, $\Delta\alpha\left(t, \boldsymbol{\sigma}^{*}(\boldsymbol{\alpha}(\boldsymbol{t}))\right) > 0$ $\forall \alpha \in (0, 1)$.

The proof of Lemma 15 is straightforward. $\forall \alpha \in [0, 1]$, BB is the stochastically stable DS.

Therefore, $\bar{\Pi}_C(BB) = \frac{v}{2}$, $\bar{\Pi}_M(BB) = \frac{v}{2} - \kappa$, and

$$\Delta\alpha\left(t\right) = \alpha(t)\left(1 - \alpha(t)\right)\kappa.$$

In other words, $\Delta\alpha\left(t\right) > 0$, $\forall\alpha(t) \in (0,1)$. From Lemma 15, we can reach the conclusions of Theorem 7, that is α continues to grow up to $\alpha = 1$, and the stochastically stable DS is BB.

$\square$

B.2 Simulations

B.2.1 Simulations from relaxations of Assumptions 1 and 2

In this appendix, I give further details about the simulations used to compute results in Section 3.3.5. First of all, in Table D1 I present the parameters to calibrate the simulations.

Parameter	Value
Periods	600
Number of players	1000
Mistake probability	0.01
Probability to revise strategy	0.05
v	$0.2, 0.3, 0.4, 0.6, 0.7, 0.8$
κ	$0.01, 0.05, 0.31, 0.21, 0.11$
Generation length	20 or 50
c	1
λ	1×10^{-16}

Table D1: Model calibration for simulations from relaxations of Assumptions 1 and 2.

I choose 600 (1000) as a limit period because it is a number large enough, which allows the system to reach all the possible states (given the initial condition) and allows for the system to stabilize around a cycle when possible. I chose a number of players large enough to justify the continuum considered in previous sections. I picked a mistake probability small enough (0.01). I assume that at each step, only the 5% of the agents can revise their strategy: this assumption rules out continuous cycling and is in line with usual works in the evolutionary game theory literature. I normalize the payoffs so that $c = 1$, and v varies from 0.2 to 0.8 depending on the harshness of conflict. I choose κ to be either 0.01, 0.05, 0.31, 0.21, or 0.11. The last three values are respectively the thresholds $\frac{2v-c}{2}$ ($\frac{c-2v}{2}$) for the cases when conflict is mild (harsh) and v takes different values. Each generation lasts either 20 or 50 periods: types evolve according to their fitness every 20 (50) rounds. I evaluate the last 4 or 30 periods depending on the specification. I chose λ as the smallest possible number that NetLogo distinguishes from 0.

In the following tables, I present the results of the simulations as presented in Section 3.3.5.

Harsh conflict	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$v = 0.2, \kappa = 0.01$	841	49.183	0.009	0.005	0.969	0.024	0.161	0.042
$v = 0.2, \kappa = 0.05$	1000	0.000	0.005	0.002	0.000	0.000	0.005	0.002
$v = 0.2, \kappa = 0.31$	967	67.904	0.291	0.449	0.025	0.019	0.260	0.401
$v = 0.3, \kappa = 0.01$	749	49.990	0.010	0.006	0.980	0.018	0.253	0.042
$v = 0.3, \kappa = 0.05$	833	60.092	0.009	0.005	0.984	0.018	0.172	0.054
$v = 0.3, \kappa = 0.21$	960	104.283	0.173	0.448	0.016	0.011	0.134	0.349
$v = 0.4, \kappa = 0.01$	640	50.990	0.010	0.007	0.979	0.025	0.358	0.042
$v = 0.4, \kappa = 0.05$	715	71.239	0.011	0.006	0.986	0.012	0.288	0.065
$v = 0.4, \kappa = 0.11$	800	98.995	0.285	0.441	0.710	0.439	0.296	0.233

Table D2: Harsh Conflict, with generation length 50 and average payoff of last 30 rounds.

Mild conflict	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$v = 0.8, \kappa = 0.01$	846	49.840	0.992	0.004	0.028	0.024	0.845	0.043
$v = 0.8, \kappa = 0.05$	1000	0.000	0.995	0.003	0.000	0.000	0.995	0.003
$v = 0.8, \kappa = 0.31$	968	61.449	0.719	0.443	0.983	0.000	0.750	0.443
$v = 0.7, \kappa = 0.01$	755	49.749	0.992	0.004	0.017	0.014	0.754	0.044
$v = 0.7, \kappa = 0.05$	794	64.529	0.993	0.004	0.022	0.017	0.793	0.061
$v = 0.7, \kappa = 0.21$	933	113.186	0.729	0.437	0.987	0.013	0.795	0.330
$v = 0.6, \kappa = 0.01$	636	52.000	0.992	0.005	0.021	0.030	0.639	0.043
$v = 0.6, \kappa = 0.05$	680	74.833	0.993	0.004	0.018	0.011	0.682	0.070
$v = 0.6, \kappa = 0.11$	783	103.010	0.698	0.450	0.318	0.443	0.695	0.227

Table D3: Mild Conflict, with generation length 50 and average payoff of last 30 rounds.

Harsh conflict	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
v = 0.2, κ = 0.01	837	48.280	0.008	0.004	0.964	0.027	0.163	0.040
v = 0.2, κ = 0.05	978	46.000	0.005	0.002	0.986	0.026	0.026	0.044
v = 0.2, κ = 0.31	956	77.872	0.272	0.439	0.023	0.015	0.229	0.370
v = 0.3, κ = 0.01	749	49.990	0.010	0.006	0.979	0.016	0.253	0.042
v = 0.3, κ = 0.05	828	58.447	0.010	0.005	0.980	0.019	0.176	0.053
v = 0.3, κ = 0.21	960	90.554	0.173	0.372	0.016	0.014	0.134	0.286
v = 0.4, κ = 0.01	646	51.807	0.010	0.007	0.980	0.027	0.353	0.042
v = 0.4, κ = 0.05	726	68.731	0.011	0.006	0.987	0.012	0.278	0.063
v = 0.4, κ = 0.11	781	105.541	0.334	0.462	0.663	0.455	0.320	0.235

Table D4: Harsh Conflict, with generation length 50 and average payoff of last 4 rounds.

Mild conflict	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
v = 0.8, $\kappa = 0.01$	851	49.990	0.983	0.098	0.039	0.099	0.843	0.086
v = 0.8, $\kappa = 0.05$	967	49.102	0.994	0.007	0.006	0.011	0.961	0.050
v = 0.8, $\kappa = 0.31$	971	65.261	0.768	0.416	0.000	0.000	0.796	0.366
v = 0.7, $\kappa = 0.01$	763	48.280	0.992	0.005	0.016	0.013	0.761	0.043
v = 0.7, $\kappa = 0.05$	820	66.332	0.992	0.004	0.020	0.017	0.818	0.062
v = 0.7, $\kappa = 0.21$	913	127.008	0.669	0.464	0.986	0.011	0.754	0.343
v = 0.6, $\kappa = 0.01$	646	49.840	0.992	0.005	0.017	0.012	0.647	0.044
v = 0.6, $\kappa = 0.05$	713	71.631	0.993	0.005	0.016	0.014	0.713	0.067
v = 0.6, $\kappa = 0.11$	816	101.705	0.806	0.385	0.203	0.382	0.764	0.195

Table D5: Mild Conflict, with generation length 50 and average payoff of last 4 rounds.

Harsh conflict	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
v = 0.2, $\kappa = 0.01$	839	106.672	0.031	0.016	0.877	0.095	0.159	0.049
v = 0.2, $\kappa = 0.05$	1000	99.499	0.005	0.002	0.000	0.000	0.007	0.020
v = 0.2, $\kappa = 0.31$	1000	0.000	0.292	0.450	0.000	0.000	0.292	0.450
v = 0.3, $\kappa = 0.01$	749	68.549	0.044	0.019	0.911	0.072	0.258	0.038
v = 0.3, $\kappa = 0.05$	849	69.993	0.029	0.016	0.926	0.051	0.163	0.060
v = 0.3, $\kappa = 0.21$	955	93.140	0.250	0.424	0.099	0.043	0.210	0.359
v = 0.4, $\kappa = 0.01$	626	86.741	0.042	0.022	0.924	0.072	0.367	0.028
v = 0.4, $\kappa = 0.05$	719	68.840	0.035	0.020	0.945	0.036	0.290	0.045
v = 0.4, $\kappa = 0.11$	928	141.478	0.219	0.403	0.064	0.036	0.154	0.282

Table D6: Harsh Conflict, with generation length 20 and average payoff of last 12 rounds.

Mild conflict	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
v = 0.8, $\kappa = 0.01$	805	112.583	0.979	0.015	0.189	0.123	0.836	0.047
v = 0.8, $\kappa = 0.05$	1000	0.000	0.995	0.002	0.000	0.000	0.995	0.002
v = 0.8, $\kappa = 0.31$	1000	0.000	0.777	0.410	0.000	0.000	0.777	0.410
v = 0.7, $\kappa = 0.01$	699	76.805	0.977	0.018	0.143	0.088	0.731	0.039
v = 0.7, $\kappa = 0.05$	828	78.842	0.981	0.014	0.100	0.059	0.831	0.065
v = 0.7, $\kappa = 0.21$	969	71.687	0.774	0.404	0.936	0.063	0.802	0.356
v = 0.6, $\kappa = 0.01$	591	66.476	0.975	0.018	0.118	0.068	0.628	0.029
v = 0.6, $\kappa = 0.05$	711	94.758	0.978	0.017	0.070	0.042	0.717	0.077
v = 0.6, $\kappa = 0.11$	871	143.035	0.724	0.425	0.510	0.435	0.768	0.285

Table D7: Mild Conflict, with generation length 20 and average payoff of last 12 rounds.

Harsh conflict	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
v = 0.2, $\kappa = 0.01$	837	128.293	0.033	0.016	0.874	0.111	0.171	0.031
v = 0.2, $\kappa = 0.05$	1000	0.000	0.005	0.002	0.000	0.000	0.005	0.002
v = 0.2, $\kappa = 0.31$	1000	0.000	0.272	0.440	0.000	0.000	0.272	0.440
v = 0.3, $\kappa = 0.01$	746	59.025	0.042	0.019	0.920	0.060	0.263	0.027
v = 0.3, $\kappa = 0.05$	842	51.342	0.032	0.015	0.919	0.059	0.172	0.034
v = 0.3, $\kappa = 0.21$	931	113.750	0.288	0.443	0.089	0.046	0.227	0.348
v = 0.4, $\kappa = 0.01$	634	62.000	0.042	0.022	0.930	0.058	0.365	0.027
v = 0.4, $\kappa = 0.05$	731	68.840	0.037	0.019	0.957	0.034	0.284	0.048
v = 0.4, $\kappa = 0.11$	866	156.346	0.301	0.440	0.431	0.429	0.236	0.289

Table D8: Harsh Conflict, with generation length 20 and average payoff of last 4 rounds.

Mild conflict	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
v = 0.8, $\kappa = 0.01$	810	132.288	0.978	0.014	0.162	0.128	0.838	0.040
v = 0.8, $\kappa = 0.05$	1000	0.000	0.995	0.002	0.000	0.000	0.984	0.002
v = 0.8, $\kappa = 0.31$	1000	0.000	0.728	0.440	0.000	0.000	0.728	0.440
v = 0.7, $\kappa = 0.01$	716	78.384	0.977	0.018	0.127	0.088	0.741	0.041
v = 0.7, $\kappa = 0.05$	816	77.097	0.981	0.012	0.097	0.060	0.821	0.060
v = 0.7, $\kappa = 0.21$	942	106.939	0.705	0.443	0.944	0.043	0.758	0.366
v = 0.6, $\kappa = 0.01$	606	69.022	0.969	0.021	0.107	0.069	0.633	0.029
v = 0.6, $\kappa = 0.05$	696	77.356	0.975	0.017	0.075	0.046	0.703	0.053
v = 0.6, $\kappa = 0.11$	824	118.423	0.744	0.407	0.333	0.392	0.748	0.228

Table D9: Mild Conflict, with generation length 20 and average payoff of last 4 rounds.

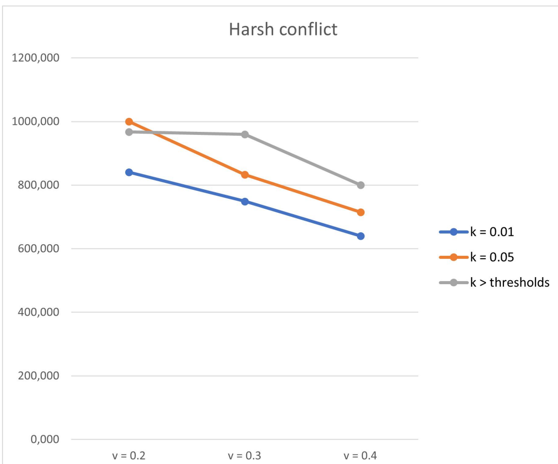

(a) $\overline{\alpha}$ *for harsh conflict*

(b) $\overline{\alpha}$ *for mild conflict*

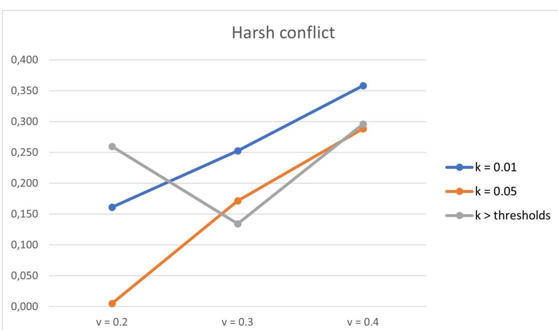

(c) n_H *for harsh conflict*

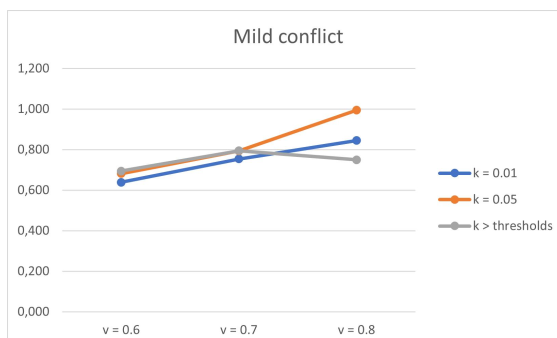

(d) n_H *for mild conflict*

Figure D1: Results for generation length 50 and average payoff of last 30 rounds

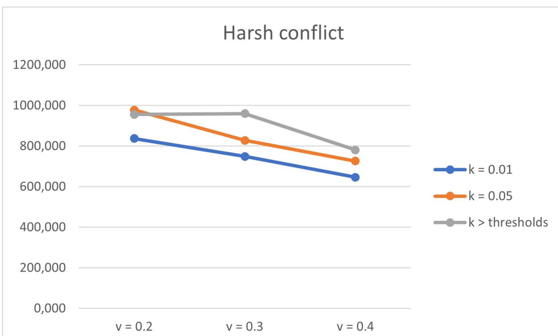

(a) $\overline{\alpha}$ *for harsh conflict*

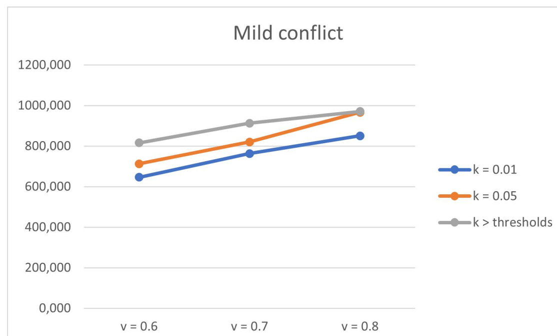

(b) $\overline{\alpha}$ *for mild conflict*

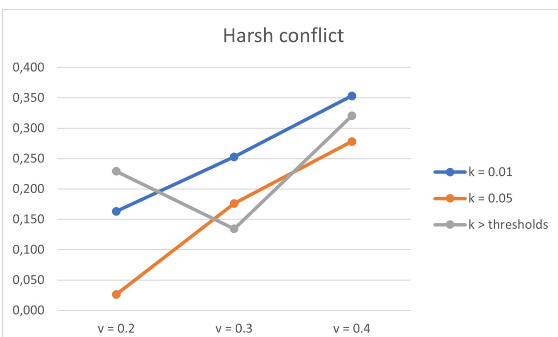

(c) n_H *for harsh conflict*

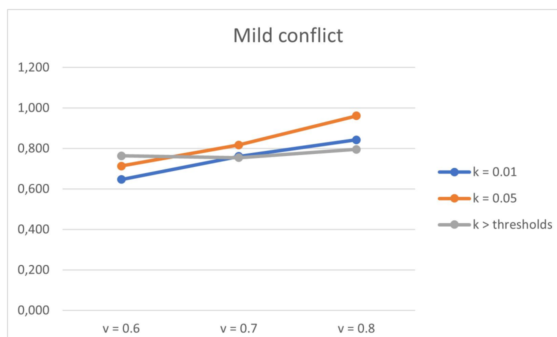

(d) n_H *for mild conflict*

Figure D2: Results for generation length 50 and average payoff of last 4 rounds

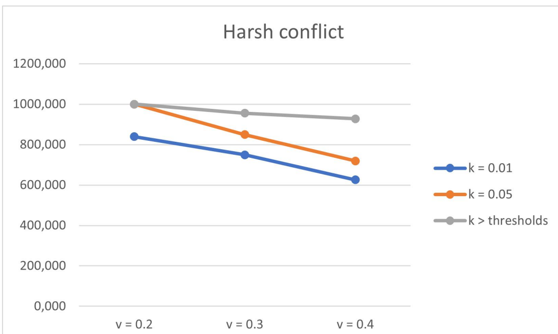

(a) $\overline{\alpha}$ *for harsh conflict*

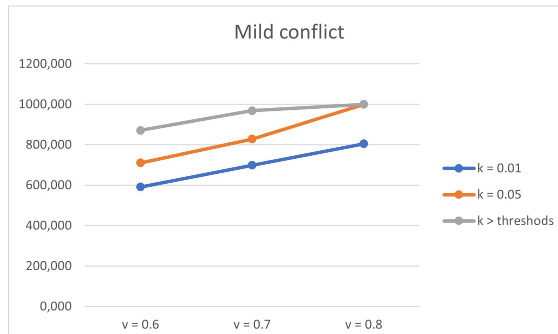

(b) $\overline{\alpha}$ *for mild conflict*

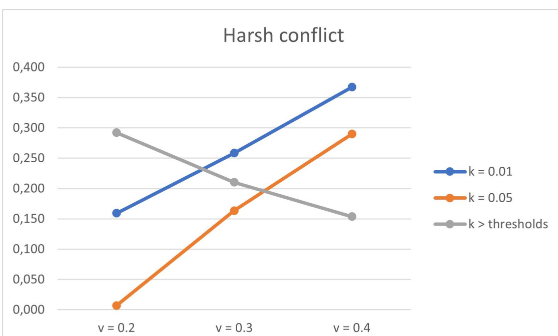

(c) n_H *for harsh conflict*

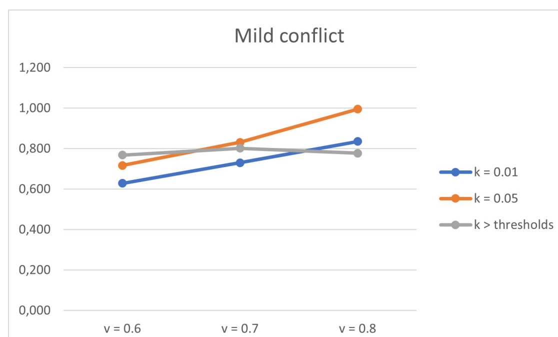

(d) n_H *for mild conflict*

Figure D3: Results for generation length 20 and average payoff of last 12 rounds

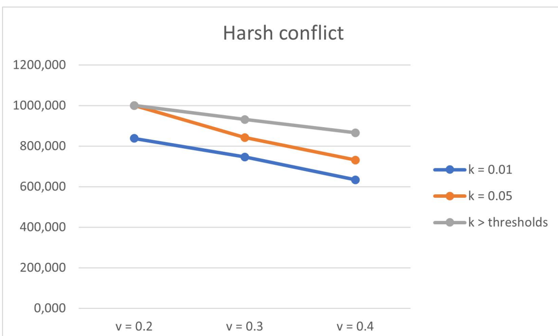

(a) $\overline{\alpha}$ *for harsh conflict*

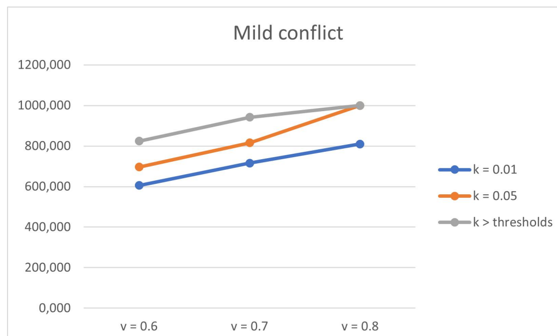

(b) $\overline{\alpha}$ *for mild conflict*

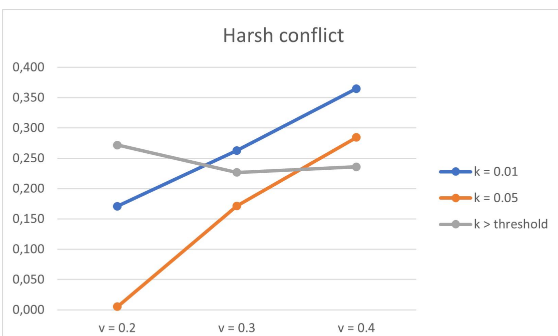

(c) n_H *for harsh conflict*

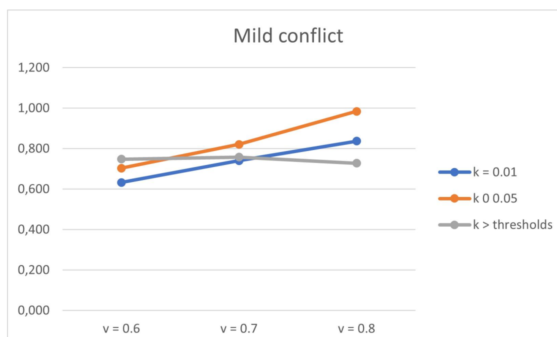

(d) n_H *for mild conflict*

Figure D4: Results for generation length 20 and average payoff of last 4 rounds

B.2.2 Simulations for uncertainty between steps

I summarize the parameters used for the calibrations of the simulations of Section 3.4.2 in the following table.

Parameter	Value
Periods	1000
Number of players	1000
Mistake probability	0.01
Probability to revise strategy	0.05
p	0.1 to 0.9 with 0.1 increment
v_1	$0.6, 0.8$
v_2	$0.2, 0.4$
κ	$0.01, 0.05, 0.31, 0.11, p0.31 + (1-p)0.11, p0.11 + (1-p)0.31$
Generation length	100
c	1
λ	1×10^{-16}

Table D10: Model calibration for simulations with uncertainty between steps.

Much of these parameters resemble the ones in Table D1. p is the probability that G_1 is played at each step, and it is varied between each simulation with 0.1 increment. v is now split into two values v_1 and v_2, for simplicity, across treatments, I choose to set the following couples: $v_1 = 0.8$, $v_2 = 0.2$, $v_1 = 0.6$, $v_2 = 0.2$, $v_1 = 0.8$, $v_2 = 0.4$, and $v_1 = 0.6$, $v_2 = 0.4$. Values of κ are 0.01, 0.05 and such that in the non uncertainty case α should reach the critical thresholds. Such a test is meant to prove the robustness of the statements in Theorems 5 and 6. The generations length is 100 times, to allow for convergence within each generation (and recreate somehow the conditions in Theorem 5). According to this length, I also increased the total number of periods to 1000.

$\kappa = 0.01$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	397	177.457	0.006	0.004	0.376	0.126	0.210	0.016
$p = 0.2$	96	130.323	0.005	0.005	0.279	0.054	0.249	0.045
$p = 0.3$	25	65.383	0.008	0.005	0.316	0.073	0.307	0.072
$p = 0.4$	18	51.730	0.007	0.004	0.389	0.082	0.382	0.085
$p = 0.5$	26	62.642	0.518	0.428	0.504	0.083	0.504	0.093
$p = 0.6$	35	80.467	0.982	0.029	0.597	0.080	0.611	0.081
$p = 0.7$	21	65.261	0.992	0.005	0.683	0.065	0.690	0.067
$p = 0.8$	110	132.288	0.994	0.005	0.727	0.056	0.761	0.044
$p = 0.9$	494	148.876	0.995	0.004	0.565	0.132	0.795	0.012

Table D11: Results for the case with uncertainty between steps, when $v_1 = 0.8, v_2 = 0.2$ and $\kappa = 0.01$.

$\kappa = 0.05$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	858	91.848	0.005	0.002	0.939	0.106	0.133	0.073
$p = 0.2$	759	90.659	0.005	0.003	0.808	0.181	0.184	0.034
$p = 0.3$	628	107.778	0.005	0.003	0.581	0.166	0.203	0.021
$p = 0.4$	509	109.631	0.089	0.270	0.486	0.100	0.266	0.126
$p = 0.5$	407	104.168	0.520	0.487	0.495	0.067	0.506	0.208
$p = 0.6$	488	119.398	0.971	0.143	0.525	0.092	0.754	0.062
$p = 0.7$	612	109.800	0.995	0.003	0.439	0.159	0.795	0.022
$p = 0.8$	760	92.736	0.994	0.002	0.199	0.181	0.817	0.041
$p = 0.9$	871	87.516	0.995	0.002	0.050	0.085	0.876	0.076

Table D12: Results for the case with uncertainty between steps, when $v_1 = 0.8, v_2 = 0.2$ and $\kappa = 0.05$.

$\kappa = 0.31$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	994	31.048	0.292	0.449	0.014	0.010	0.287	0.441
$p = 0.2$	998	19.900	0.263	0.434	0.020	0.000	0.261	0.431
$p = 0.3$	1000	0.000	0.391	0.483	0.000	0.000	0.391	0.483
$p = 0.4$	1000	0.000	0.272	0.440	0.000	0.000	0.272	0.440
$p = 0.5$	1000	0.000	0.530	0.494	0.000	0.000	0.530	0.494
$p = 0.6$	1000	0.000	0.649	0.472	0.000	0.000	0.649	0.472
$p = 0.7$	1000	0.000	0.589	0.487	0.000	0.000	0.589	0.487
$p = 0.8$	1000	0.000	0.787	0.403	0.000	0.000	0.787	0.403
$p = 0.9$	991	34.914	0.688	0.458	0.982	0.024	0.697	0.446

Table D13: Results for the case with uncertainty between steps, when $v_1 = 0.8, v_2 = 0.2$ and $\kappa = 0.31$.

$\kappa = 0.01$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	628	114.961	0.005	0.002	0.583	0.186	0.200	0.014
$p = 0.2$	331	155.367	0.006	0.004	0.345	0.078	0.222	0.026
$p = 0.3$	199	122.061	0.007	0.006	0.340	0.067	0.271	0.057
$p = 0.4$	211	101.877	0.033	0.111	0.416	0.080	0.334	0.080
$p = 0.5$	181	135.422	0.201	0.333	0.497	0.069	0.436	0.097
$p = 0.6$	38	91.411	0.785	0.348	0.541	0.060	0.546	0.061
$p = 0.7$	8	33.705	0.972	0.033	0.560	0.044	0.564	0.040
$p = 0.8$	16	50.438	0.990	0.017	0.579	0.031	0.587	0.025
$p = 0.9$	230	157.162	0.994	0.005	0.458	0.113	0.599	0.014

Table D14: Results for the case with uncertainty between steps, when $v_1 = 0.6, v_2 = 0.2$ and $\kappa = 0.01$.

$\kappa = 0.05$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	868	73.321	0.005	0.002	0.971	0.027	0.132	0.069
$p = 0.2$	834	71.021	0.005	0.002	0.940	0.098	0.156	0.055
$p = 0.3$	755	82.916	0.005	0.003	0.809	0.182	0.188	0.028
$p = 0.4$	669	99.695	0.006	0.003	0.652	0.169	0.204	0.027
$p = 0.5$	526	148.068	0.102	0.282	0.544	0.128	0.280	0.112
$p = 0.6$	452	130.752	0.575	0.486	0.456	0.158	0.457	0.164
$p = 0.7$	522	87.841	0.785	0.403	0.287	0.251	0.547	0.120
$p = 0.8$	602	72.083	0.940	0.222	0.096	0.184	0.616	0.073
$p = 0.9$	651	65.567	0.965	0.168	0.044	0.155	0.653	0.064

Table D15: Results for the case with uncertainty between steps, when $v_1 = 0.6, v_2 = 0.2$ and $\kappa = 0.05$.

$\kappa = p0.11 + (1-p)0.31$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	966	73.783	0.222	0.410	0.006	0.006	0.189	0.349
$p = 0.2$	956	81.633	0.272	0.439	0.005	0.005	0.228	0.369
$p = 0.3$	966	72.415	0.253	0.428	0.005	0.005	0.219	0.373
$p = 0.4$	960	82.462	0.282	0.444	0.007	0.006	0.243	0.384
$p = 0.5$	925	97.340	0.480	0.494	0.004	0.005	0.406	0.421
$p = 0.6$	885	95.263	0.727	0.439	0.006	0.006	0.614	0.376
$p = 0.7$	872	93.894	0.619	0.480	0.192	0.380	0.539	0.378
$p = 0.8$	814	72.139	0.678	0.462	0.305	0.434	0.623	0.299
$p = 0.9$	803	87.698	0.718	0.444	0.268	0.418	0.676	0.261

Table D16: Results for the case with uncertainty between steps, when $v_1 = 0.6, v_2 = 0.2$ and $\kappa = p0.11 + (1-p)0.31$.

$\kappa = 0.01$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	254	170.540	0.005	0.004	0.563	0.127	0.400	0.017
$p = 0.2$	29	68.257	0.011	0.024	0.430	0.040	0.416	0.024
$p = 0.3$	12	51.536	0.005	0.007	0.437	0.044	0.430	0.036
$p = 0.4$	59	107.791	0.241	0.319	0.479	0.063	0.465	0.065
$p = 0.5$	182	126.791	0.752	0.361	0.507	0.084	0.566	0.093
$p = 0.6$	215	99.373	0.973	0.110	0.590	0.079	0.675	0.078
$p = 0.7$	188	131.362	0.995	0.005	0.662	0.061	0.727	0.059
$p = 0.8$	321	140.922	0.994	0.005	0.668	0.069	0.781	0.026
$p = 0.9$	645	126.787	0.994	0.003	0.387	0.205	0.803	0.018

Table D17: Results for the case with uncertainty between steps, when $v_1 = 0.8, v_2 = 0.4$ and $\kappa = 0.01$.

$\kappa = 0.05$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	653	71.351	0.025	0.139	0.961	0.131	0.342	0.064
$p = 0.2$	590	80.623	0.038	0.170	0.896	0.168	0.380	0.060
$p = 0.3$	494	100.817	0.184	0.374	0.692	0.237	0.442	0.107
$p = 0.4$	464	122.082	0.402	0.484	0.563	0.167	0.538	0.166
$p = 0.5$	575	117.792	0.956	0.186	0.427	0.124	0.749	0.082
$p = 0.6$	670	91.104	0.995	0.003	0.354	0.161	0.797	0.021
$p = 0.7$	763	91.274	0.995	0.003	0.171	0.184	0.815	0.033
$p = 0.8$	813	71.631	0.995	0.003	0.091	0.121	0.832	0.049
$p = 0.9$	875	75.333	0.995	0.002	0.029	0.045	0.876	0.069

Table D18: Results for the case with uncertainty between steps, when $v_1 = 0.8, v_2 = 0.4$ and $\kappa = 0.05$.

$\kappa = p0.31 + (1-p)0.11$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	798	87.155	0.273	0.439	0.754	0.398	0.323	0.257
$p = 0.2$	817	66.415	0.282	0.445	0.740	0.414	0.355	0.290
$p = 0.3$	866	87.430	0.263	0.434	0.885	0.308	0.375	0.351
$p = 0.4$	912	98.265	0.431	0.490	0.995	0.005	0.518	0.419
$p = 0.5$	937	90.172	0.579	0.489	0.994	0.006	0.642	0.419
$p = 0.6$	947	95.347	0.679	0.462	0.996	0.005	0.731	0.389
$p = 0.7$	947	89.950	0.678	0.462	0.994	0.006	0.731	0.388
$p = 0.8$	961	74.693	0.747	0.429	0.993	0.006	0.786	0.363
$p = 0.9$	968	67.646	0.797	0.396	0.994	0.006	0.829	0.333

Table D19: Results for the case with uncertainty between steps, when $v_1 = 0.8, v_2 = 0.4$ and $\kappa = p0.31 + (1-p)0.11$.

$\kappa = 0.01$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	481	114.625	0.005	0.003	0.774	0.150	0.388	0.026
$p = 0.2$	249	124.495	0.006	0.004	0.550	0.088	0.404	0.016
$p = 0.3$	106	108.462	0.029	0.103	0.471	0.057	0.418	0.031
$p = 0.4$	92	103.615	0.104	0.219	0.488	0.053	0.449	0.047
$p = 0.5$	135	108.972	0.534	0.392	0.499	0.069	0.509	0.068
$p = 0.6$	66	88.566	0.913	0.182	0.530	0.052	0.559	0.041
$p = 0.7$	118	116.086	0.991	0.013	0.516	0.062	0.578	0.035
$p = 0.8$	245	134.443	0.994	0.004	0.452	0.091	0.596	0.016
$p = 0.9$	514	128.078	0.995	0.004	0.178	0.163	0.618	0.035

Table D20: Results for the case with uncertainty between steps, when $v_1 = 0.6, v_2 = 0.4$ and $\kappa = 0.01$.

$\kappa = 0.05$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
$p = 0.1$	654	57.306	0.005	0.002	0.985	0.014	0.344	0.053
$p = 0.2$	655	49.749	0.005	0.003	0.984	0.015	0.342	0.044
$p = 0.3$	634	63.592	0.075	0.253	0.931	0.194	0.370	0.077
$p = 0.4$	595	81.701	0.322	0.462	0.701	0.390	0.435	0.118
$p = 0.5$	580	77.460	0.540	0.493	0.454	0.420	0.511	0.127
$p = 0.6$	574	74.324	0.619	0.479	0.364	0.399	0.538	0.114
$p = 0.7$	628	67.941	0.966	0.168	0.051	0.138	0.637	0.056
$p = 0.8$	647	51.875	0.995	0.003	0.017	0.015	0.650	0.048
$p = 0.9$	654	59.025	0.995	0.003	0.016	0.015	0.657	0.055

Table D21: Results for the case with uncertainty between steps, when $v_1 = 0.6, v_2 = 0.4$ and $\kappa = 0.05$.

$\kappa = 0.11$	Avg $\overline{\alpha}$	Std $\overline{\alpha}$	Avg n_H^C	Std n_H^C	Avg n_H^M	Std n_H^M	Avg n_H	Std n_H
p = 0.1	792	86.810	0.163	0.363	0.837	0.361	0.263	0.186
p = 0.2	740	78.740	0.243	0.423	0.758	0.421	0.334	0.187
p = 0.3	744	71.162	0.282	0.445	0.718	0.443	0.352	0.204
p = 0.4	740	64.807	0.351	0.472	0.650	0.470	0.406	0.227
p = 0.5	741	58.472	0.540	0.493	0.461	0.492	0.517	0.245
p = 0.6	750	60.828	0.678	0.462	0.320	0.458	0.612	0.229
p = 0.7	725	72.629	0.648	0.471	0.347	0.466	0.610	0.207
p = 0.8	740	78.740	0.757	0.423	0.241	0.418	0.666	0.188
p = 0.9	753	103.397	0.659	0.469	0.338	0.463	0.660	0.219

Table D22: Results for the case with uncertainty between steps, when $v_1 = 0.6, v_2 = 0.4$ and $\kappa = 0.11$.

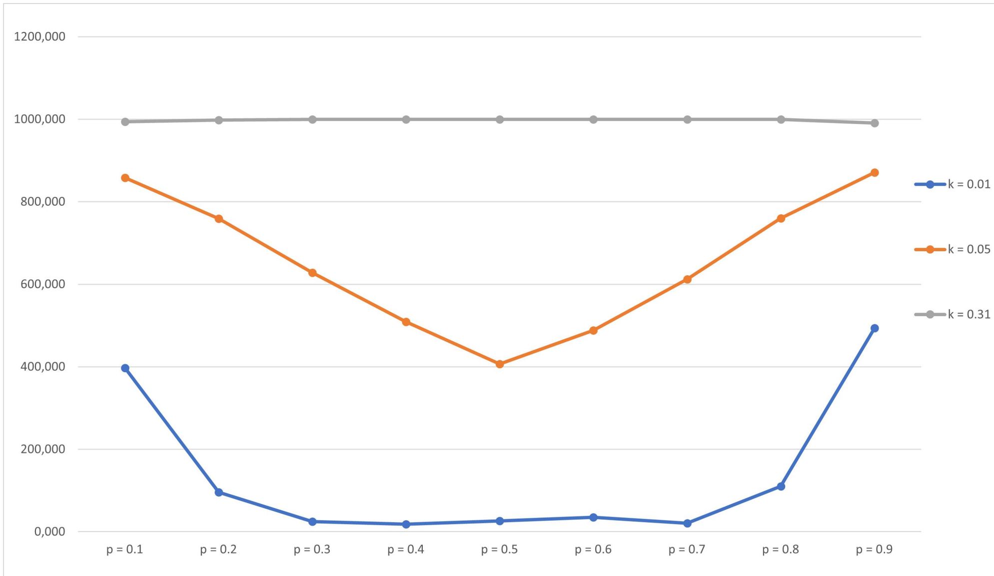

(a) $\overline{\alpha}$

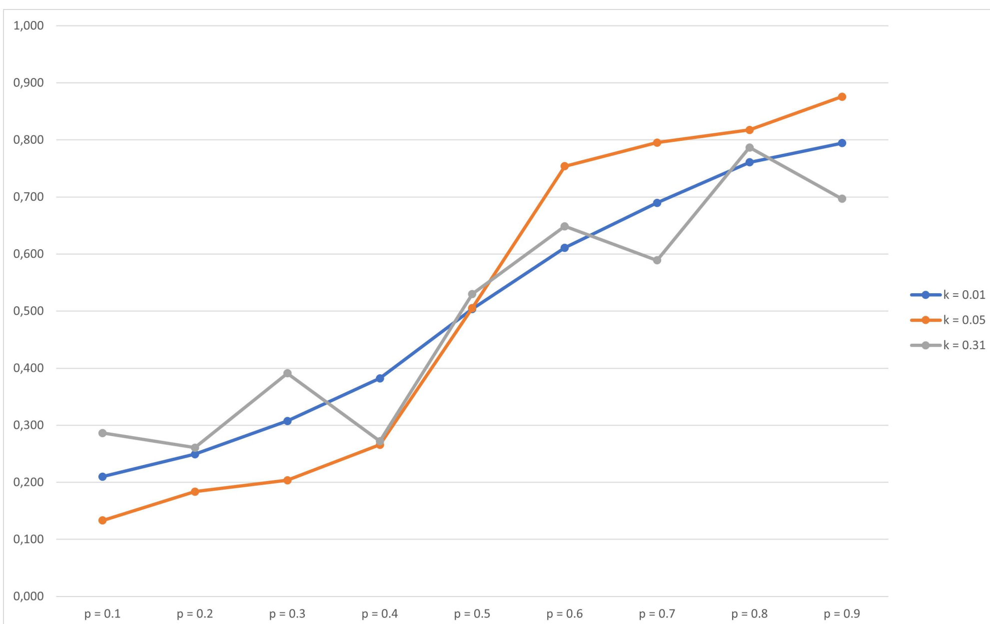

(b) n_H

Figure D5: Results for $v_1 = 0.8$ and $v_2 = 0.2$.

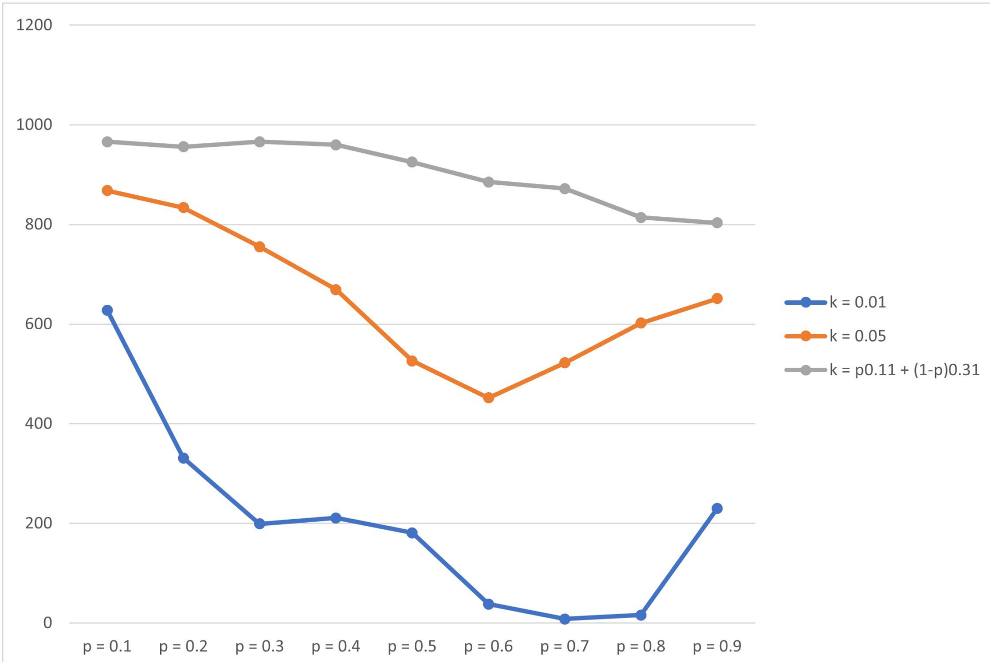

(a) $\overline{\alpha}$

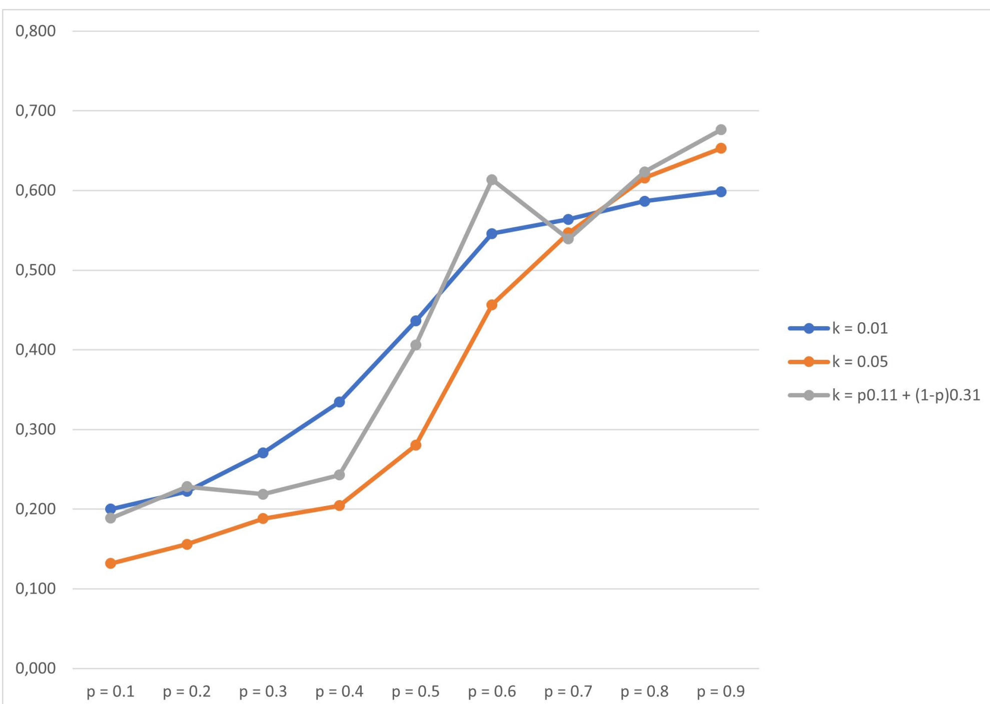

(b) n_H

Figure D6: Results for $v_1 = 0.6$ and $v_2 = 0.2$.

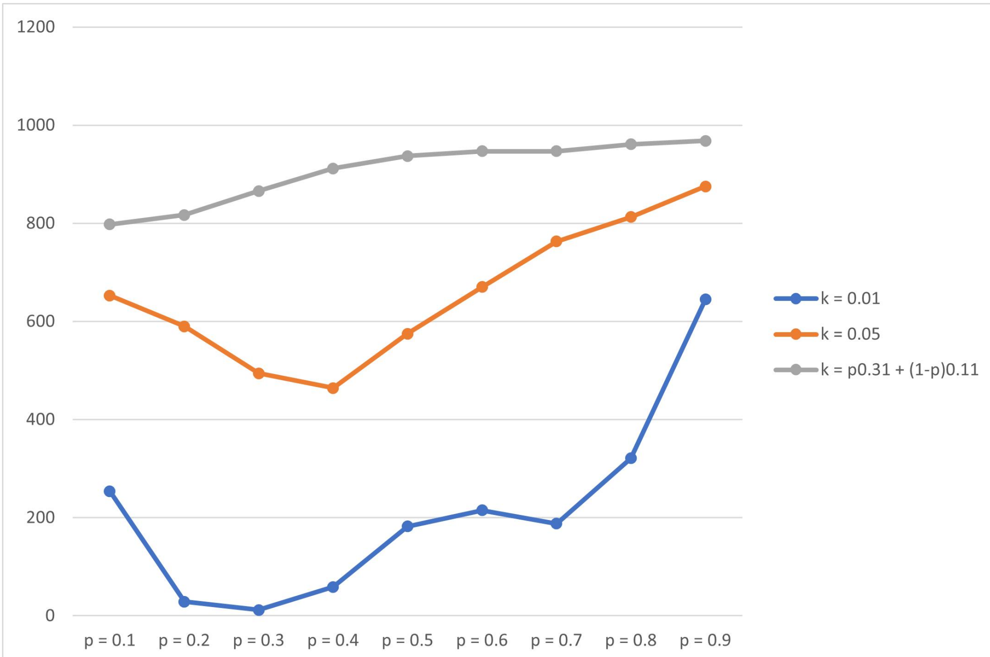

(a) $\overline{\alpha}$

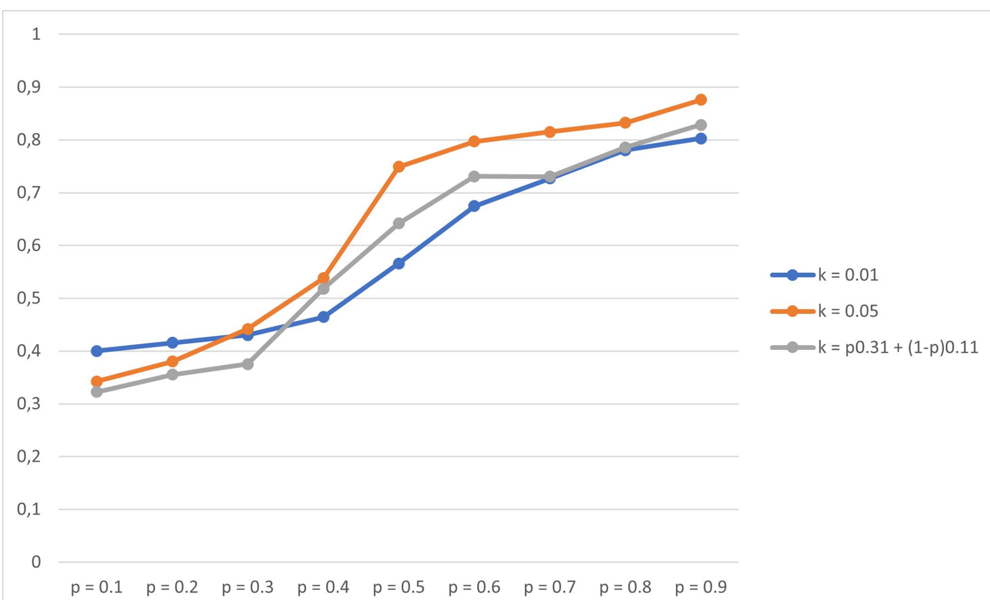

(b) n_H

Figure D7: Results for $v_1 = 0.8$ and $v_2 = 0.4$.

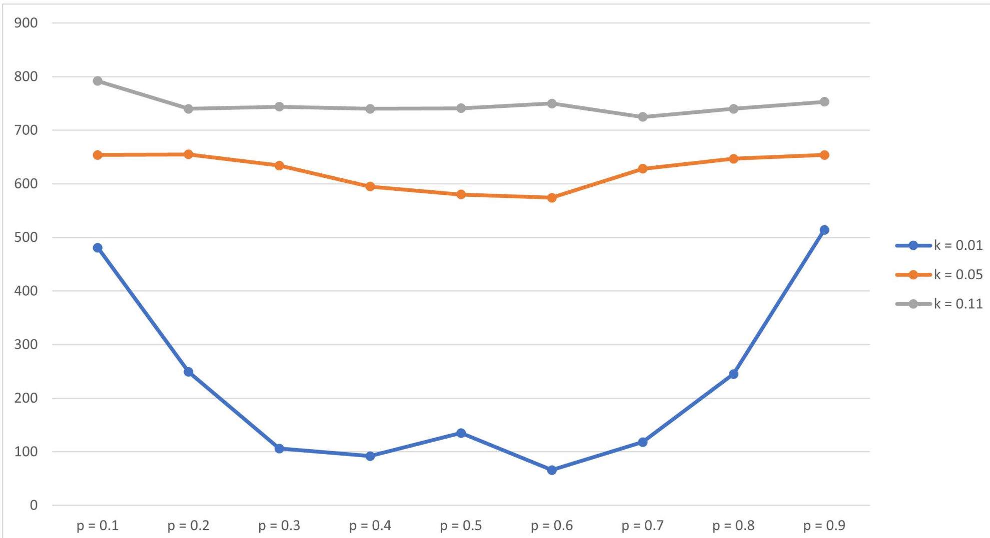

(a) $\overline{\alpha}$

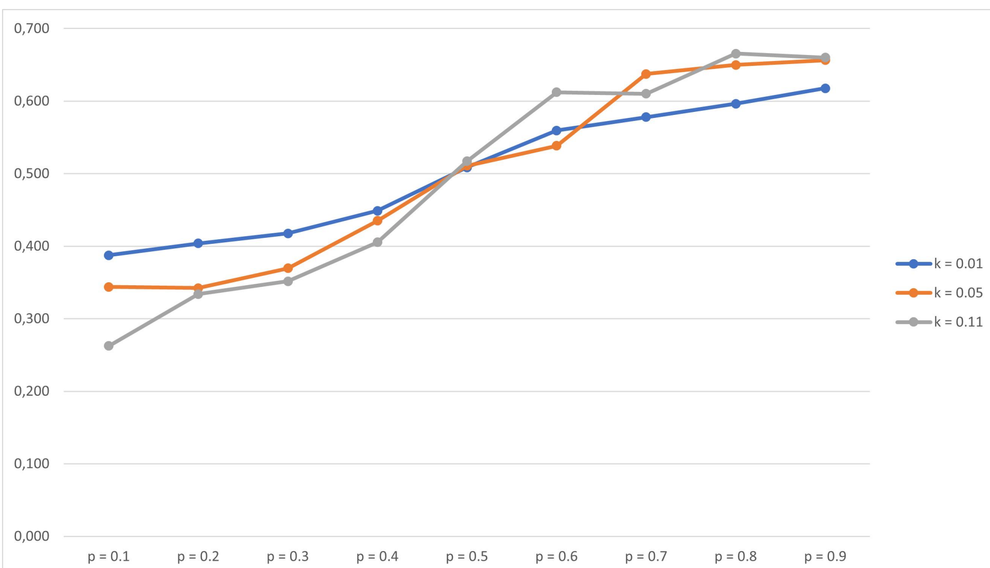

(b) n_H

Figure D8: Results for $v_1 = 0.6$ and $v_2 = 0.4$.

Appendix C

Chapter 3 Appendix

C.1 Proofs

Proof of Proposition 1. To prove the first statement, assume that f is first convex and then concave; and also that $s_0 \neq \hat{s}$. (The proof of the second statement is entirely analogous and therefore omitted). As a preliminary, observe that, if $s_0 = 0$, then $s_t = f^t(0) = 0$ for all $t \in \mathbb{N}$ (recall that 0 is a fixed point). Hence, if $s_0 = 0$, then $\lim_{t\to\infty} s_t = 0$. Similarly, $\lim_{t\to\infty} s_t = 1$ if $s_t = 1$. Therefore, the statement holds trivially in the cases of $s_0 = 0$ and $s_1 = 1$. It remains to consider the cases of $s_0 \in (0, \hat{s})$ and $s_0 \in (\hat{s}, 1)$.

Suppose then that $s_0 \in (0, \hat{s})$. (The argument when $s_0 \in (\hat{s}, 1)$ follows similar lines and is therefore omitted.) Given that f is convex on $[0, \hat{s}]$, and furthermore that $f(0) = 0$ and $f(\hat{s}) = \hat{s}$, one may check that $f(s) < s$ for all $s \in (0, \hat{s})$. In addition, given that f is increasing, $f(s) > f(0) = 0$ for all such s. So for all $s \in (0, s^*)$, $f(s) \in (0, s)$.

Given that this fact, and also that $s_0 \in (0, s^*)$, one can show by induction that $s_t \in (0, s_{t-1})$ for all $t \in \mathbb{N}$. Hence, the sequence $\{s_t\}_{t=0}^{\infty}$ is strictly decreasing and bounded from below by zero. By the monotone convergence theorem, it therefore has a limit s^*. Furthermore, since f is continuous, every limit s^* must be a fixed point:

$$s^* = \lim_{t\to\infty} s_t = \lim_{t\to\infty} s_{t+1} = \lim_{t\to\infty} f(s_t) = f(\lim_{t\to\infty} s_t) = f(s^*), \tag{C.1}$$

where the penultimate equality uses the continuity of f. By assumption, f only has three fixed points: 0, $\hat{s}$, and 1. Since $s^* < s_0 < \hat{s}$, we see that the only possible limit is 0.

The argument given above establishes that $\lim_{t\to\infty} s_t = 0$ if $s_0 \in (0, \hat{s})$. By an analogous argument, one may show that $\lim_{t\to\infty} s_t = 1$ if $s_0 \in (\hat{s}, 1)$. Furthermore, we have already observed that $\lim_{t\to\infty} s_t \in \{0,1\}$ if either $s_0 = 0$ or $s_0 = 1$. This establishes that $\lim_{t\to\infty} s_t \in \{0,1\}$ for any initial value $s_0 \neq \hat{s}$. □

Proof of Proposition 2. As noted earlier, given that f is continuous, we know that every limit s^* must be a fixed point. Moreover, since f is increasing, $f(s) \geq f(0)$ for all $s \in [0,1]$. In particular, then, $f(s^*) \geq f(0)$. However, since s^* is a fixed point (established earlier), $s^* = f(s^*)$. From this, we conclude that $s^* \geq f(0)$; and a symmetric argument establishes that $s^* \leq f(1)$. □

Proof of Proposition 3. Let us define the difference in utilities by

$$\Delta(\alpha_i, s) \equiv U(a_i = 1) - U(a_i = 0) = \alpha_i + m(s) - m(1-s) \tag{C.2}$$

To show that the individual has tipping point preferences, we consider three cases:

Case 1. For all $s \in [0,1]$, $\Delta(\alpha_i, s) \geq 0$. In that case, $a_i^* = 1$ for all s (where a_i^* denotes the optimal action). Equivalently, $a_i^* = 1$ if and only if $s \geq t_i$, for any tipping point that satisfies $t_i \leq 0$. Thus, the individual has tipping point preferences (for example, we may set $t_i = 0$).

Case 2. For all $s \in [0,1]$, $\Delta(\alpha_i, s) < 0$. Similarly to before, this means that $a_i^* = 0$ for all s. Equivalently, $a_i^* = 1$ if and only if $s \geq t_i$ for any tipping point that satisfies $t_i > 1$: for example, we can set $t_i = 2$. Thus, the individual again has tipping point preferences.

Case 3. $\Delta(\alpha_i, s) \geq 0$ for some $s \in [0,1]$; but also $\Delta(\alpha_i, s) < 0$ for some $s' \in [0,1]$. Differentiating with respect to s, we see that

$$\Delta'(s) = m'(s) + m'(1-s) > 0. \tag{C.3}$$

Thus, $\Delta(s)$ is strictly increasing (and continuous) in s. This means that there is a unique $s^* \in (0,1]$ such that $\Delta(s) > 0$ for $s > s^*$, $\Delta(s) = 0$ when $s = s^*$, and $\Delta(s) < 0$ for $s < s^*$. Thus, the individual has tipping point preferences for $t_i = s^*$.

To prove the second statement, observe that, if $t_i \in (0,1)$, then

$$\alpha_i + m(t_i) - m(1-t_i) = 0 \tag{C.4}$$

We can totally differentiate to obtain

$$\frac{\partial \alpha_i}{\partial \alpha_i} + \frac{\partial m(t_i)}{\partial t_i}\frac{\partial t_i}{\partial \alpha_i} + \frac{\partial m(1-t_i)}{\partial t_i}\frac{\partial t_i}{\partial \alpha_i} = 0 \tag{C.5}$$

which implies that

$$\frac{\partial t_i}{\partial \alpha_i} = -\frac{1}{m'(t_i) + m'(1-t_i)} < 0 \tag{C.6}$$

where the inequality holds since both derivatives are strictly positive. Thus, t_i is strictly decreasing in α_i for $t_i \in (0,1)$. $\square$

C.2 Local interaction

In this section, we extend the model described in Section 4.2 to allow for local interaction in overlapping networks. The model presented here shares some similarities to the model studied by Efferson et al. (2020). An important difference is that, while Efferson et al. (2020) assume that decision makers choose randomly, we instead assume that they choose deterministically but with heterogenous decision rules. In addition, our model assumes that individuals respond to the decisions of their 'neighbours' (in line with our experimental settings); whereas Efferson et al. (2020) assume that they best respond to the entire population.

In our baseline model, we assume the following:[1]

- There are l^2 agents, each located on a node of a grid with side length $l \in \mathbb{N}^+$. Let (r, c) denote the agent located at row r and column c; so the set of agents is the set $N = \{(r, c) \colon r \in \{1, ..., l\}, c \in \{1, ..., l\}\}$.
- As in our experiments, agents are faced with a binary choice: they must either take an action (denoted $a_{r,c} = 1$) or not take the action (denoted $a_{r,c} = 0$).
- Each agent (r, c) has a set of 'neighbours' $N_{r,c}$ whose actions they can see. For each agent, we assume that $N_{r,c} = \{(i, j) \colon (i, j) \in N, |i - r| \leq 1, |j - c| \leq 1, (i, j) \neq (r, c)\}$. Observe that agents in the interior have 4 neighbours, agents on the edge have 3 neighbours, and agents in the corners have 2 neighbours.
- We define $m^1_{r,c}$ as the share of individual (r, c)'s neighbours who have chosen to do the action. Formally, $m^1_{r,c} = \frac{1}{|N_{r,c}|} \sum_{(r,c) \in N_{r,c}} a_{r,c}$ where $|N_{r,c}|$ is the cardinality of $N_{r,c}$.
- Each agent is endowed with a (fixed) tipping point $t_{r,c} \in [0, 1]$. As in the main text, we assume that they choose $a_{r,c} = 1$ if and only if $t_{r,c} \geq m^1_{r,c}$.
- Agents interact over multiple periods. In each period, one agent is chosen to move at random; and updates their action (if necessary) by comparing their tipping point $t_{r,c}$ with the share of their neighbours who are taking the action $m^1_{r,c}$.

To assess the robustness of our results, we also study an alternative model that departs from the model sketched above in various ways.[2] In this model — which we label the *edgeless model* — each agent is linked with the same number of neighbours. In addition, each agent has a probability $\epsilon \in [0, 1]$ of making a 'mistake', i.e. choosing the opposite action as that required by their tipping point. Finally, a share $p \in [0, 1]$ of agents are selected in period to

[1] The corresponding code can be viewed here: https://github.com/Itzhak95/tipping_points

[2] The corresponding code can be viewed here: https://github.com/rrozzi/tipping_point-netlogo-

revise their action; so in principle multiple agents can update their action simultaneously.

To simulate the results of our models, we use the following procedure:

- We specify a distribution of tipping points in the population, and randomly scatter these tipping points across the agents.
- We also specify the share of agents who initially take the action; and we randomly scatter the agents who are taking the action on grid.
- We then allow the model to run for 1000 periods (or until it is 'stable' so no further changes can occur).

As discussed above, the results of the model could in principle depend on the way in which tipping points and initial actions are scattered. As a result, we conduct all simulations 1000 times and report the distribution of results across simulations.

Before turning to our main results, we provide a simple example to illustrate the mechanics of the model. To generate this example, we suppose that, initially, 40% of agents are taking the action; and we set $s = 5$. In addition, we assume (for expositional simplicity) that all agents have a tipping point $t_{r,c} = 0.5$, so choose $a_{r,c} = 1$ if and only if half or more of their neighbours are doing the action. While one would normally repeat the simulation many times, here we just report the outcome of one simulation.

After randomly scattering the initial actions, we obtain the initial state

$$\begin{array}{ccccc}
0 & 1 & 1 & 1 & 0 \\
{\color{blue}0} & 1 & 0 & 0 & 1 \\
{\color{red}1} & {\color{blue}0} & 1 & 1 & 0 \\
{\color{blue}0} & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 1 & 0
\end{array}$$

As can be seen, 10 of the $s^2 = 25$ agents initially take the action (indicated by a 1); the rest do not. Several rounds now progress in which the player chosen to move does not wish to update their action. Eventually, however, the player at row 3 and column 1 is chosen to move (they are coloured in red). Since none of their neighbours (coloured in blue) were taking the action, they choose to switch to action 0. This yields the new state

$$\begin{array}{ccccc} 0 & 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \end{array}$$

As the process continues, additional players are given the opportunity to also revise their action. After 13 such revisions, we finally obtain the state

$$\begin{array}{ccccc} 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array}$$

This state is *stable* in the sense that no agent has an incentive to change their behaviour. The agent in the top left is surrounded by neighbours who choose $a_{r,c} = 1$, so would also want to choose $a_{r,c} = 1$ if allowed to update their action. The agent at $(2, 2)$ is surrounded by 4 neighbours, half of whom are taking the action; so also chooses $a_{r,c} = 1$ (recall that all tipping points are set at $t_{r,c} = 0.5$). Meanwhile, the agents at $(2, 1)$ and $(1, 2)$ are each surrounded by 3 neighbours, 2 of whom are choosing the action; so they also wish to choose the action. Finally, one can verify that the agents choosing $a_{r,c} = 0$ are choosing optimally given their tipping point and the share of their neighbours who are taking the action.

We now calibrate our model using the tipping point distributions calculated in Section 4.5. We assume a population size of 100; and the edgeless model further assumes an error probability $\epsilon = 0.01$ and a probability of revision $p = 0.07$. As stated above, each simulation is run for 1000 periods (or until the obtained state is stable); and all simulations are conducted $1,000$ times. Tables C1 and C2 display the results for experiment 1 (face masks) and experiment 2 (Zoom calls) respectively. The first row specifies the initial share who are assumed to do the action. The rows 'mean (main)' and 'mean (edgeless)' display the average share who end up doing the activity in the main specification and edgeless model respectively. The rows 'Var (main)' and 'Var (edgeless)' specify the variance of outcomes across simulations.

Initial share	0.05	0.15	0.25	0.35	0.45	0.55	0.65	0.75	0.85	0.95
Mean (main)	.228	.230	.230	.231	.231	.233	.234	.235	.235	.236
Var (main)	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
Mean (edgeless)	.237	.242	.242	.242	.240	.242	.240	.239	.240	.239
Var (edgeless)	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000

Notes. This table shows the results of simulating our models using the distribution of tipping points obtained by experiment 1 (see Table 4.2). That is, we set $p_0 = .203$, $p_1 = .033$, $p_2 = .044$, $p_3 = .085$, $p_4 = .123$, $p_5 = .513$.

Table C1: Simulation results (experiment 1)

Initial share	0.05	0.15	0.25	0.35	0.45	0.55	0.65	0.75	0.85	0.95
Mean (main)	.311	.316	.321	.329	.333	.338	.341	.345	.349	.353
Var (main)	.001	.001	.001	.001	.001	.001	.001	.001	.001	.001
Mean (edgeless)	.353	.353	.347	.356	.355	.362	.354	.360	.356	.355
Var (edgeless)	.001	.001	.001	.001	.001	.001	.001	.001	.001	.001

Notes. This table shows the results of simulating our models using the distribution of tipping points obtained by experiment 2 (see Table 4.4). That is, we set $p_0 = .209$, $p_1 = .118$, $p_2 = .091$, $p_3 = .099$, $p_4 = .072$, $p_5 = .411$.

Table C2: Simulation results (experiment 2)

Three results are apparent. First, we see that the results of the simulations are relatively insensitive to the initial share who are assumed to do the activity. This is especially true in the edgeless model since this assumes that agents occasionally make errors, which weakens dependence on initial conditions in the usual way (Young, 1993a). Second, the variance in outcomes across simulations is very low, which is again points to the lack of importance of initial conditions (since different simulations generate different outcomes only due to variation in initial conditions). Finally, and most importantly, we see that our models generate convergence to interior equilibria that resemble those obtained from the simple model of Section 4.2. This should not come as a surprise given that some agents always do the action, that other agents never do the action, and that a final group of agents engage in copying behaviour.

C.3 Tables and figures

Treatment	Frequency	Percentage
0	127	19.7
1	134	20.7
2	128	19.8
3	124	19.2
4	133	20.6
Total	646	100.0

Notes. This table shows how many subjects were allocated into each of the five treatments in the first experiment.

Table C3: Sample allocation (experiment 1)

Variable	Mean	Std. Dev.
Age	20.8	3.90
Male	.497	.500
Humanities	.283	.451
MPLS	.240	.427
Medical Sciences	.127	.333
Social Sciences	.333	.471
Pre	.201	.401
n	646	

Notes. This table shows the descriptive statistics for experiment 1 (see Table 4.1 for a description of the variables).

Table C4: Descriptive statistics (experiment 1)

Variable	No controls	Main Specification	All Controls
Treatment 1	.044	.033	.029
	[.047]	[.030]	[.034]
Treatment 2	.171***	.073**	.079**
	[.053]	[.032]	[.035]
Treatment 3	.238***	.162***	.168***
	[.055]	[.040]	[.043]
Treatment 4	.331***	.283***	.304***
	[.054]	[.042]	[.046]
Pre		.504***	.498***
		[.030]	[.031]
Age		.003	.002
		[.005]	[.004]
Male		-.006	-.002
		[.026]	[.028]
n	646	646	620

Notes. This table reports the exact same specifications reported on in Table 4.2, except using logistic instead of linear regressions. Robust standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table C5: Logit regressions (experiment 1)

Variable	No controls	Main Specification	All Controls
Treatment 1	.044	.036	.029
	[.047]	[.031]	[.034]
Treatment 2	.171***	.078**	.078**
	[.053]	[.033]	[.035]
Treatment 3	.238***	.163***	.162***
	[.055]	[.040]	[.043]
Treatment 4	.331***	.284***	.298***
	[.054]	[.043]	[.046]
Pre		.518***	.512***
		[.024]	[.027]
Age		.002	.001
		[.004]	[.004]
Male		-.007	-.004
		[.026]	[.028]
n	646	646	620

Notes. This table reports the exact same specifications reported on in Table 4.2, except using probit instead of linear regressions. Robust standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table C6: Probit regressions (experiment 1)

Comparison	No controls	Main specification	All controls
T0 vs T1	.355	.278	.536
T1 vs T2	.019	.205	.163
T2 vs T3	.269	.051	.068
T3 vs T4	.131	.019	.017
T0 vs T2	.001	.014	.032
T1 vs T3	.001	.002	.002
T2 vs T4	.008	.000	.000

Notes. This table reports p-values corresponding to hypothesis that the effect of treatment k is the same as the effect of treatment k', for all possible $k \neq k'$. We do this for the three specifications considered in Table 4.2.

Table C7: Comparisons (experiment 1)

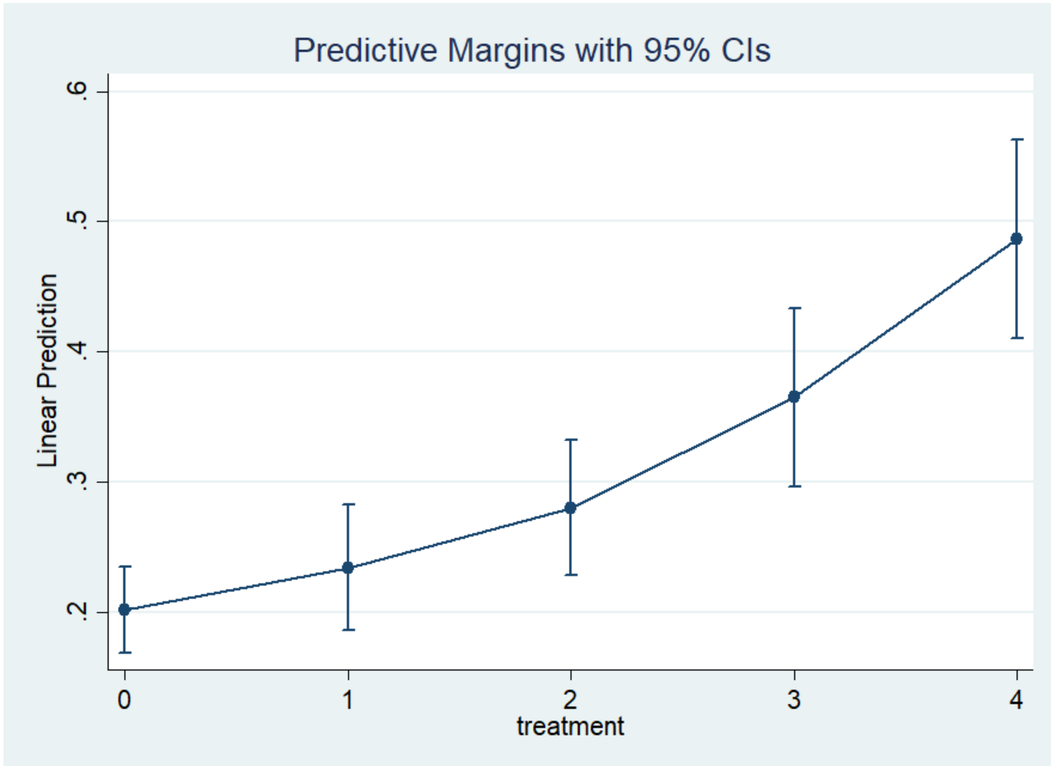

Notes. This figure shows the marginal effects of the treatments in the main specification, including the confidence intervals.

Figure C1: Mask wearing by treatment group (with confidence intervals)

	T0	T1	T2	T3	T4
Putting mask on	.028	.080	.106	.223	.368
Taking mask off	.056	.143	.059	.067	.037

Notes. The first row shows the share who put a mask on given that they entered the room without wearing a mask. The second row shows the share who took their mask off given that they entered the room wearing a mask.

Table C8: Changes (experiment 1)

Variable	Linear	Quadratic	Cubic
Masks	.070***	0.008	0.024
	[-.010]	[-.028]	[-.062]
Masksˆ2		.016**	.004
		[-.008]	[-.045]
Masksˆ3			.002
			[-.008]
Pre	.752***	.757***	.757***
	[-.029]	[-.029]	[-.029]
Age	.002	.002	.002
	[-.005]	[-.005]	[-.005]
Male	-.008	-.007	-.007
	[-.026]	[-.026]	[-.026]
Constant	-.022	.016	.014
	[-.102]	[-.107]	[-.107]
Joint test	.000	.000	.000
R^2	.491	.494	.494

Notes. In this table, we regress whether subjects chose to wear a mask on the number of experimenters wearing a mask, as well higher order terms to capture potential non-linearity (we also control for 'pre', age, and gender). The penultimate row reports p-values corresponding to the hypothesis that the coefficients on all mask variables are zero.

Table C9: Polynomial regressions (experiment 1)

Explanation	Frequency
Trying to avoid judgement	.148
Trying to cater to others' preferences	.511
Trying to follow rules	.148
Reciprocity	.023
COVID risks	.011
Not answering question	.159
n	88

Notes. This table shows the frequencies of the explanations given by subjects (see Appendix C.5 for a detailed description of the categories).

Table C10: Explanations from online survey

Treatment	Frequency	Percentage
0	232	20.8
1	204	18.3
2	223	20.0
3	241	21.7
4	213	19.1
Total	1113	100.0

Notes. This table shows how many subjects were allocated into each of the five treatments in the second experiment.

Table C11: Sample allocation (experiment 2)

Variable	Mean	Std. Dev.
Age	42.4	13.9
Male	.465	.499
n	1113	

Notes. This table shows the descriptive statistics for experiment 2.

Table C12: Descriptive statistics (experiment 2)

Variable	No controls	Main specification	All controls
Treatment 1	.077*	.127***	.133***
	[.043]	[.039]	[.039]
Treatment 2	.176***	.215***	.218***
	[.043]	[.039]	[.040]
Treatment 3	.281***	.314***	.323***
	[.043]	[.039]	[.045]
Treatment 4	.355***	.385***	.389***
	[.044]	[.041]	[.051]
Pre		.741***	.743***
		[.092]	[.092]
Age		.000	.000
		[.001]	[.001]
Male		.023	.023
		[.027]	[.027]
n	1,113	1,111	1,109

Notes. This table reports the exact same specifications reported on in Table 4.4, except using logistic instead of linear regressions. Robust standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table C13: Logit regressions (experiment 2)

Variable	No controls	Main specification	All controls
Treatment 1	.077*	.125***	.130***
	[.043]	[.039]	[.039]
Treatment 2	.176***	.216***	.218***
	[.043]	[.039]	[.040]
Treatment 3	.281***	.312***	.321***
	[.043]	[.039]	[.046]
Treatment 4	.355***	.385***	.389***
	[.044]	[.040]	[.052]
Pre		.701***	.699***
		[.075]	[.076]
Age		.000	.000
		[.001]	[.001]
Male		.024	.025
		[.027]	[.027]
n	1,113	1,111	1,109

Notes. This table reports the exact same specifications reported on in Table 4.4, except using probit instead of linear regressions. Robust standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table C14: Probit regressions (experiment 2)

Comparison	No controls	Main Specification	All Controls
T0 vs T1	.074	.003	.002
T1 vs T2	.035	.043	.051
T2 vs T3	.022	.028	.020
T3 vs T4	.116	.116	.152
T0 vs T2	.000	.000	.000
T1 vs T3	.000	.000	.000
T2 vs T4	.000	.000	.001

Notes. This table reports p-values corresponding to hypothesis that the effect of treatment k is the same as the effect of treatment k', for all possible $k \neq k'$. We do this for the three specifications considered in Table 4.4.

Table C15: Comparisons (experiment 2)

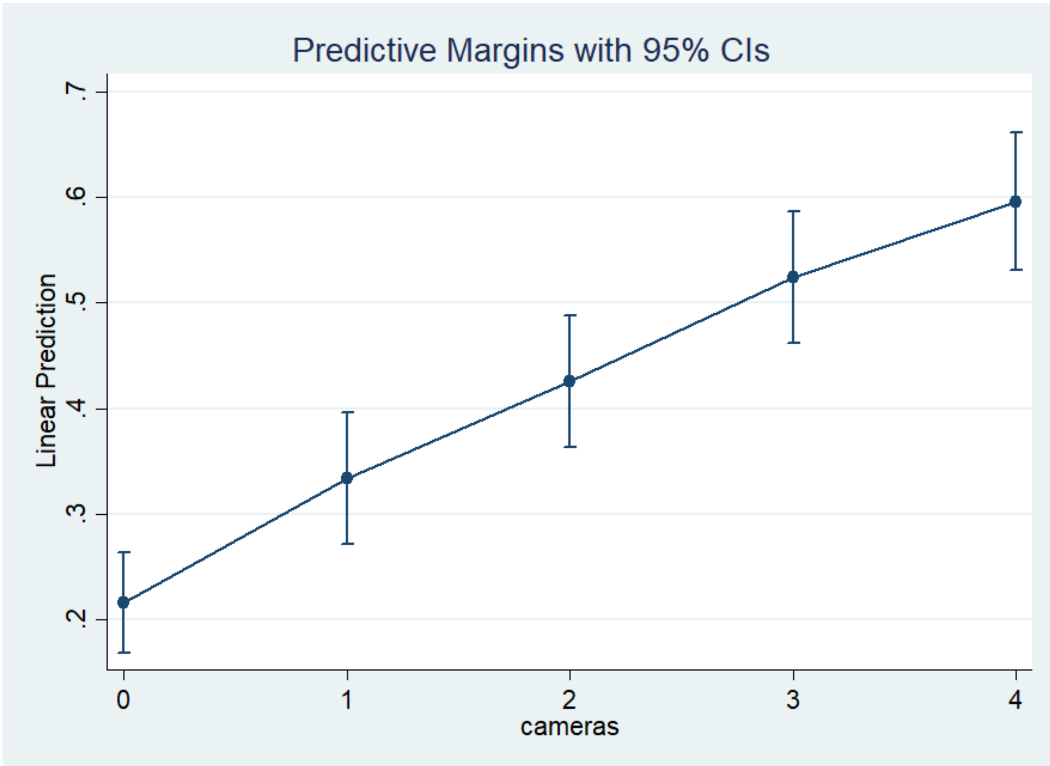

Notes. This figure shows the marginal effects of the treatments in the main specification, including the confidence intervals.

Figure C2: Camera use by treatment group (with confidence intervals)

	T0	T1	T2	T3	T4
Turning camera on	0.156	0.296	0.381	0.491	0.566
Turning camera off	0.111	0.125	0.000	0.059	0.000

Notes. The first row shows the share who turned their camera on given that they joined the call without video. The second row shows the share who turned their camera off given that they joined the call with video.

Table C16: Changes (experiment 2)

Variable	Linear	Quadratic	Cubic
Cameras	.095***	.119***	.119
	-.009	-.032	-.074
Cameras^2		-.006	-.006
		-.008	-.049
Cameras^3			.000
			-.008
Pre	.576***	.578***	.578***
	-.033	-.033	-.033
Age	.000	.000	.000
	-.001	-.001	-.001
Male	.023	.024	.024
	-.027	-.027	-.027
Constant	.169***	.156***	.156***
	-.046	-.047	-.047
Joint test	.000	.000	.000
R^2	.161	.161	.161

Notes. In this table, we regress whether subjects chose to use their camera on the number of experimenters using a camera, as well higher order terms to capture potential non-linearity (we also control for 'pre', age, and gender). The penultimate row reports p-values corresponding to the hypothesis that the coefficients on all camera variables are zero.

Table C17: Polynomial regressions (experiment 2)

C.4 Experimental protocols

In this section, we provide a more detailed outline of the experimental protocols followed in both experiments.

Experiment 1 (face masks). There were four experimenter roles, labelled 1 through 4. Experimenter 1's role was to greet the subject and (at the end) bid them goodbye. Experimenter 2's role was to record the data and ask some demographic questions. Experimenter 3's role was to ask the question about the lotteries. Experimenter 4's only role was to was to introduce themselves when asked to do so and wear a face mask when this was required by the randomisation.

Subjects were asked to arrive at a room within a particular time slot. Importantly, it was not possible to view inside the room without entering it; and once a subject had entered, the only people they could see were the experimenters inside the room. Before each subject entered the room, the number of the four experimenters in the room who were wearing a mask (and the allocation of masks to experimenters) had been randomised. Thus, there were five treatment groups, corresponding to: 0/4 masks, 1/4 masks, 2/4 masks, 3/4 masks, 4/4 masks. All four experimenters were seated in front of a table on which a box of face masks, hand sanitiser, and bag of checkers had been placed.

Once a subject arrived, the experiment proceeded in the following manner:

1. Experimenter 1 welcomed the participant in and asked all other experimenters to introduce themselves. The other three experimenters then did this by stating their name and subject of study.

2. Experimenter 2 asked the subject for their name, age, and academic division. They recorded these on a spreadsheet, along with their apparent gender and whether they had entered the room wearing a mask.

3. Experimenter 3 asked the subject the following question. '*As you may know, we have issued a fixed number of lottery tickets for an Amazon voucher. I am now going to give you two options to choose from. The first option is simply to get one lottery ticket for the voucher. The second option is a gamble between 2 and 0 lottery tickets. Specifically, if you take the second option, then you will take a checker from the bag in front of you. If you get a black checker — and there are six of these — then you will get two lottery tickets. However, if you get a white checker — and there are five of these — then you will not get any lottery tickets. So what do you choose — getting one lottery ticket for sure, or taking the gamble between 2 and 0 lottery tickets?*'

4. Person 1 thanked the participant for coming and told them that the experiment had concluded.

Occasionally, a subject asked if they should wear a face mask. In response to such questions, Experimenter 2 always replied: ‘it’s up to you’.

Experiment 2 (Zoom calls). As before, there were four experimenter roles, labelled 1 through 4. Experimenter 1’s role was to admit subjects into the Zoom room and guide them through the experiment. Experimenter 2 was the data recorder, and Experimenter 3 double checked all data. Experimenter 4 pasted a link to the survey in the Zoom chat just before the subject was asked to leave the room.

Subjects were asked to join the Zoom call at a particular time slot. Before the subject joined the call, the number of experimenters with their camera on, and which experimenters had their camera on, had been randomised. Once a subject arrived in the Zoom waiting room, the experiment proceeded in the following manner:

1. Experimenter 1 thanked the subject for joining and asked if whether they could hear the audio. They stated their name, and said that the other experimenters would now introduce themselves.

2. The other three experimenters on the call now introduced themselves by stating their name.

3. Experimenter 1 asked the subject for their age. Experimenter 2 recorded this on a spreadsheet along with their apparent gender, and whether they had joined the call with their video camera on.[3]

4. Experimenter 1 then asked the subject the following question: ‘*If we were to hypothetically give you a £10 bonus payment, would you choose to share half of it with the next person on the call?*’

5. If the subject’s camera had remained off throughout the call, Experimenter 1 asked them if there were any issues with their camera.

6. Experimenter 1 then thanked the subject for participating, asked them to click the survey link, and removed them from the room.

Some notes:

[3] In contrast to the first experiment, we did not require Experimenter 2 to ask the demographic questions due to the more rapid pace of data collection.

- All experimenters ensured that they did not have a Zoom profile photo; so when they turned their camera off, only the text of their name was visible.
- If a subject asked if they should turn their camera on, then Experimenter 1 told them that 'it's up to you'.
- If a subject turned their camera on when the host asked if they had issues with their camera, then we ignored this from a data recording point of view (see previous discussion).

C.5 Explanations

In this section, we elaborate on the way in which we categorise subject explanations for 'switching behaviour' (recorded in the online survey). The categories are as follows:

1. 'Trying to avoid judgement'
   Elaboration: if you see many others wearing a mask, you might infer that these others want you to wear a mask. This in turn might induce you to wear a mask if you do not want to be negatively judged by the others.
   Example from dataset: 'Don't see the point in wearing a mask now, but if everyone else was then social conformity and not wanting to be the odd one out would mean I probably would.'

2. 'Trying to cater to others' preferences'
   Elaboration: If you see many others wearing a mask, you might infer that these others want you to wear a mask. This in turn might induce you to wear a mask if you want to altruistically cater to their preferences (e.g. to make them feel more comfortable).
   Comment: Observe that, like the explanation before, this explanation is based on learning about the preferences of others through their actions.
   Example from dataset: 'if i see someone wearing a mask it makes me think that they might be uncomfortable about the virus so if i had one on me i would wear it.'

3. 'Trying to follow rules'
   Elaboration: If you see many others wearing a mask, you might conclude that a (formal or informal) rule requires wearing a mask — and you might generally try to follow rules.
   Comment: In practice, it can be hard to distinguish this from the first explanation: individuals may follow informal rules to avoid judgement. However, we used this category since some participants mentioned rules without mentioning a fear of being judged.
   Example from dataset: 'If majority of people wearing mask, I assume there is written/unwritten rules regarding this, in that room that I am not aware of'

4. 'Reciprocity'
   Elaboration: if you see many others wearing a mask, you might infer that they are trying to protect you. As a result, you might want to protect them (as in Rabin (1993)).
   Example from dataset: 'I want to protect others, but if they aren't willing to protect me then I'm not willing to protect them'

5. ‘COVID risks’

   Elaboration: if you see many others wearing a mask, you might conclude that the COVID risk around you is high: for example, these people might be wearing a mask because they have COVID. Assuming that you want to avoid COVID, you might therefore choose to wear a mask.

   (Only) example from dataset: ‘I’d think if anyone were wearing a mask they probably have a good reason to, like being a close contact of a positive tester. Or if someone is just being particularly careful I would also think they have a good reason to and try to respect that.’

6. ‘Not answering question’

   Elaboration: Some subjects explained the various factors which determine whether they choose to wear a mask, but did not explain why their decision to wear a mask would vary with the number of others wearing a mask (the question we were interested in).

   Example from dataset: ‘It depends on the setting, and if I were carrying a mask at the time. If asked in advance I would always wear a mask, and would never want to make someone feel uncomfortable. However if the situation was relatively safe, I would not feel a need to wear a mask.’

Appendix D

Chapter 4 Appendix

D.1 Proofs for pure strategies

In this section, we prove the results from Theorem 8 and 9.

With abuse of notation, in proofs, we denote by $\cdot$ the choices of all the other players when focusing on one single player. Moreover, when comparing equilibrium choices with coalition choices, we refer to $U_i(k^*, \cdot)$ as the utility from choosing social group k in equilibrium, and $U_i(k'_\epsilon, \cdot)$ as the utility of a coalition of mass ϵ choosing social group k'.

We start by giving some general results that apply both to the homophily case and to the benchmark case.

Lemma 16. *If DM chooses d_0, there cannot exist any pooling equilibrium on H.*

Proof.

This lemma is quite standard in signaling games, and the proof is straightforward. If DM chooses d_0 and blindly assigns H tasks to candidates, they have no incentive to invest in their human capital since they are assigned H tasks in any case. Therefore whenever DM chooses d_0, all candidates choose L, and DM's beliefs are not correct. Hence, the one described above is not an equilibrium.

$\square$

Corollary 10. *If DM chooses d_0, in equilibrium it must be that $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0, 0)$.*

The argument stands from the previous Lemma, therefore, the proof is omitted.

Lemma 17. *If* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (1,0)$*, or* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0,1)$*, then, it must be that* $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0,0)$*.*

Proof.

The argument of this proof stands from Proposition 1 of Bilancini and Boncinelli (2018c). When DM has beliefs $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (1,0)$ or $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0,1)$, her expected utility from choosing d_0 depends on $\mathbf{p}^*(\mathbf{s}|\mathbf{k})$. Since we assume both ζ and $\overline{\phi}^s$ to be homogeneous, it must be that $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0,0)$, or $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (1,1)$.

If $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0,0)$ and if $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (1,0)$ or $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0,1)$, then $V(\cdot, d_0, \mathbf{m}^*_{\mathbf{k}}) = \tau(1+\delta_L)$ and $V(\cdot, d_1, \mathbf{m}^*_{\mathbf{s}}) = \tau(1+\delta_L)-c$. If $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (1,1)$ and if $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (1,0)$ or $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0,1)$, then $V(\cdot, d_0, \mathbf{m}^*_{\mathbf{k}}) = \tau(1+\delta_H)$ and $V(\cdot, d_1, \mathbf{m}^*_{\mathbf{s}}) = \tau(1+\delta_H) - c$. Consequently, if $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (1,0)$ or $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0,1)$, DM always plays d_0, and hence, due to Corollary 10, it must be that $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0,0)$.

□

The next lemma will be useful in proving equilibrium strategies for candidates.

Lemma 18. *Given a candidate* i*, with* $s_i = s'$*:*

$$MU_i(t_i, x, s', d_1, \mathbf{m}^*_{\mathbf{s}}) = MU_i(t_i, y, s', d_1, \mathbf{m}^*_{\mathbf{s}}).$$

The proof is straightforward and therefore is omitted. The intuition is simple: when DM plays d_1, she always observes each candidate's type and skill. Hence, DM assigns a candidate to the same task regardless of the social group he chooses.

We now prove the principal statements in the main text.

Proof of Theorem 8.

Firstly, we prove that there cannot exist separating equilibria on social groups.

When DM has beliefs $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (1,0)$ ($\mathbf{p}(\mathbf{t}|\mathbf{k}) = (0,1)$), she assigns α (β) tasks to social group x, and β (α) tasks to social group y. For Lemma 17, DM plays d_0 this case, and $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0,0)$. Therefore, If $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (1,0)$ ($\mathbf{p}(\mathbf{t}|\mathbf{k}) = (0,1)$), the decision-maker plays d_0, and $\mathbf{m}_{\mathbf{k}} = (\alpha_L, \beta_L)$ ($\mathbf{m}_{\mathbf{k}} = (\beta_L, \alpha_L)$). Next, let us consider candidates' strategies. Thanks to Lemma 16, we know that they all choose L. Consider the situation in which each A type chooses x, and each B type chooses y. In this case, $U_i(x,\cdot) = \phi_{\alpha_L} + \eta, \forall i$, and $U_i(y,\cdot) = \phi_{\beta_L} + \eta, \forall i$.

Given that $U_i(x,\cdot) > U_i(y,\cdot)$, each i chooses x, and therefore, there cannot exists any separating equilibrium on social groups.

Next, let us consider the case of pooling equilibria on social groups. There can be two cases, $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (p_A, p(A|y))$ ($\mathbf{p}(\mathbf{t}|\mathbf{k}) = (p(A|x), p_A)$), or $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (p_A, p_A)$.

Let us consider the second case. First, let us consider the case of $\mathbf{p}(\mathbf{s}|\mathbf{k}) = (1,1)$. In this case, DM plays $\mathbf{m_k} = (\alpha_H, \alpha_H)$. Therefore, $E\left[V(\cdot, d_0, \mathbf{m_k})\right] = p_A\tau + \tau\delta_H$, if there is substitutability or $E\left[V(\cdot, d_0, \mathbf{m_k})\right] = p_A\tau(1+\delta_H)$, if there is complementarity. And in both cases, $V(\cdot, d_1, \mathbf{m}^*_\mathbf{S}) = \tau(1+\delta_H) - c$.

Therefore, she chooses d_1, if and only if $c < p_B\tau$ (substitutability) or $c < p_B\tau(1+\delta_H)$ (complementarity). If she plays d_1, due to Lemma 18, we know that each candidate earns the same material utility by choosing social group x, or y. Moreover, due to Assumption 4, all candidates choose H. To sustain such an equilibrium it must be that both A and B types choose x in a fraction q, and y in a fraction $1-q$. In such a case, $U_i(x,\cdot) = U_i(y,\cdot) = \phi_{\alpha_H} - \zeta + \eta, \forall i \in A$ and $U_j(x,\cdot) = U_j(y,\cdot) = \phi_{\beta_H} - \zeta + \eta, \forall j \in B$. Therefore, there exist infinite pooling equilibria on social groups ($\forall q, 1-q \in [0,1]$), pooling on H, where DM buys the information. Thanks to Lemma 16, we know that there could not exist any pooling equilibrium on social groups, pooling on H when DM plays d_0, and hence, there is no equilibrium when $c > p_B\tau$ ($c > p_B\tau(1+\delta_H)$).

With a similar reasoning it is possible to show that there exist infinite pooling equilibria on social groups, pooling on L, when $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0,0)$, $c > p_B\tau(1+\delta_L)$ (both under substitutability and complementarity) and DM plays $(d_0, \mathbf{m}^*_\mathbf{k})$, with $\mathbf{m}^*_\mathbf{k} = (\alpha_L, \alpha_L)$.

All the above equilibria are also CPE. Indeed, in the case of pooling equilibrium on H, by Lemma 18, we know that the material utility of each candidate is the same, no matter which social group they choose. Moreover, if there is no homophily, $SU_i(k'_\epsilon, \cdot) = \eta, \forall\epsilon$, which is the same utility they earn in equilibrium. Therefore, there are no profitable coalitions in case of pooling equilibria on social groups and on H. The same reasoning applies to a pooling equilibrium on L, given that DM assigns α_L to each social group.

Now, we consider a degenerate pooling equilibrium on x (the argument also stands for the case of a degenerate pooling equilibrium on y for symmetry in payoffs). In this case, it must be that $p(A|x) = p_A$, and it can be that $p(A|y) > p(B|y)$, or that $p(A|y) < p(B|y)$. If $p(A|y) > p(B|y)$, similar to the above cases, there exists a pooling equilibrium on x, such an equilibrium can be pooling on H, and or L depending on DM's beliefs, and information decision. Such equilibrium is also CPE.

Let us consider the case such that $p(A|y) < p(B|y)$. If $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (1, 1)$, DM plays $\mathbf{m_k} = (\alpha_H, \beta_H)$, if she plays d_0. Similar to what we prove before, there is a pooling equilibria on x and on H, where DM buys the information, and similarly to what we prove before, such an equilibrium is also CPE.

There also exists a pooling equilibrium on L. In fact, if $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0, 0)$, DM plays $\mathbf{m_k} = (\alpha_L, \beta_L)$ if she chooses d_0. Similarly to what we prove before, there exists no equilibrium if $c < p_B\tau(1+\delta_L)$. However, consider the case when $c > p_B\tau(1+\delta_L)$. $U_i(x, \cdot) = \phi_{\alpha_L} + \eta, \forall i$, and $U_i(y, \cdot) = \phi_{\beta_L}, \forall i$. Therefore, each candidate chooses social group x, and the one described above is an equilibrium.

Similar to what we prove before, these kind of equilibria are also CPE.

The uniqueness of equilibria as meant in the statement is a consequence of the above arguments.

□

Corollary 11.

Consider a game where there is not homophily, and types and skills are either complement or substitute.

- *Under substitutability, if $p_B\tau < c < p_B\tau(1 + \delta_L)$, there exists no equilibrium. If $c > p_B\tau(1+\delta_L)$, there only exist pooling equilibria on social groups, where all candidates play L, DM does not buy the information, and she assigns α_L to each candidate. These equilibria are CPE.*
- *Under complementarity, if $p_B\tau < c < p_B\tau(1 + \delta_L)$, there only exist pooling equilibria on social groups, where all candidates play H, and DM buys the information. These equilibria are CPE. If $p_B\tau(1 + \delta_L) < c < p_B\tau(1 + \delta_H)$, there exist pooling equilibria on social groups, where all candidates play H, and DM buys the information, and there exist pooling equilibria on social groups, where all candidates play L, DM does not buy the information, and she assigns α_L to each candidate. These equilibria are CPE. If $c > p_B\tau(1 + \delta_H)$, there only exist pooling equilibria on social groups, where all candidates play L, DM does not buy the information, and she assigns α_L to each candidate. These equilibria are CPE.*

The nonexistence of equilibria for $p_B\tau < c < p_B\tau(1 + \delta_L)$ under substitutability must not be a concern since it is a consequence of Assumption 4, that can be seen as a refinement of the possible equilibria. The multiplicity under complementarity instead is inevitable, but it does not change that qualitative result of the main paper.

Proof of Theorem 9.

We prove the theorem bullet point by bullet point.

Pooling equilibria on social groups, and on H if $c < p_B\tau$ ($c < p_B\tau(1+\delta_H)$).

To prove DM's strategies we refer to Theorem 8's proof. In case, $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (p_A, p_A)$ and $\mathbf{p}(\mathbf{s}|\mathbf{k}) = (1,1)$, she chooses d_1 if and only if $c < p_B\tau$ (substitutability) or if and only if $c < p_B\tau(1+\delta_H)$ (complementarity).

If DM plays d_1, candidates play H due to Assumption 4, and due to Lemma 18, their social group choice depends only on SU. Specifically, if every candidate chooses x,

$$U_i(x,\cdot) = \phi_{\alpha_H} - \zeta + p_A\eta, \forall i \in A,$$

$$U_i(y,\cdot) = \phi_{\alpha_H} - \zeta, \forall i \in A,$$

$$U_j(x,\cdot) = \phi_{\beta_H} - \zeta + p_B\eta, \forall j \in B,$$

$$U_j(y,\cdot) = \phi_{\beta_H} - \zeta, \forall j \in B.$$

Given that both $U_i(x,\cdot) > U_i(y,\cdot)$ and $U_j(x,\cdot) > U_j(y,\cdot)$, there exists a pooling equilibrium on x and pooling on H if $c < p_B\tau$ ($c < p_B\tau(1+\delta_H)$). However, this equilibrium is not coalition-proof. Indeed, consider a coalition of B types of mass ϵ that deviates to y.

$$U_j(y_\epsilon,\cdot) = \phi_{\beta_H} - \zeta + \eta, \forall j \in \epsilon.$$

Now consider a sub-coalition of mass $\epsilon' < \epsilon$, that chooses x.

$$U_j(x_{\epsilon'},\cdot) = \phi_{\beta_H} - \zeta + \left(\frac{p_B(1+\epsilon'-\epsilon)}{p_B(1+\epsilon'-\epsilon)+p_A}\right)\eta, \forall j \in \epsilon'.$$

Given that $U_j(y_\epsilon,\cdot) > U_j(x_{\epsilon'},\cdot)$, the coalition is self-enforcing. Moreover, $U_j(y_\epsilon,\cdot) > U_j(x^*,\cdot)$. Therefore, the coalition is profitable with respect to the equilibrium and no pooling equilibrium on social groups and H is CPE.

A similar reasoning applies to a pooling equilibrium on y and to all the other pooling equilibria on social groups such that A and B types choose x in a fraction q, and y in a fraction $1-q$.

Pooling equilibria on social groups, and on L if $c > p_B\tau(1+\delta_L)$.

Firstly, consider the case for which $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (p_A, p_A)$, and $\mathbf{p}(\mathbf{s}|\mathbf{k}) = (0,0)$. Coherently with Theorem 8, DM chooses d_0 when $c > p_B\tau(1+\delta_L)$. Thanks to Lemma 16, we know that there exists no pooling equilibrium when DM plays d_1; therefore, we only consider the case of $c > p_B\tau(1+\delta_L)$. When she chooses d_0, she chooses $\mathbf{m_k} = (\alpha_L, \alpha_L)$.

In this case, all candidates choose L, and they earn $MU = \phi_{\alpha_L}$ whichever social group they choose. Consider a case such that both A and B types choose x with frequency q. $SU_i(x,\cdot) = SU_i(y,\cdot) = \eta p_A, \forall i \in A$, while $SU_j(x,\cdot) = SU_j(y,\cdot) = \eta p_A, \forall j \in B$. Therefore, candidates choose social group x with frequency q, and the above described is an equilibrium. Similarly to the case when DM chooses d_1, such an equilibrium is not a CPE.

The same reasoning applies to a degenerate pooling equilibrium on x (y) s.t. $p(t|k) = (p_A, p(A|y))$ with $p(A|y) \geq p_A$ ($p(t|k) = (p(A|x), p_A)$ with $p(A|x) \geq p_A$).

Secondly, we consider the case when a degenerate pooling equilibrium on x (y) s.t. $p(t|k) = (p_A, p(A|y))$ with $p(A|y) < p_A$ ($p(t|k) = (p(A|x), p_A)$ with $p(A|x) < p_A$). As for Theorem 8's proof, DM chooses d_0 and $\mathbf{m_k} = (\alpha_L, \beta_L)$. In this case, from Lemma 16 all candidates choose L. Moreover,

$$U_i(x,\cdot) = \phi_{\alpha_L} + p_A\eta, \forall i \in A,$$

$$U_i(y,\cdot) = \phi_{\beta_L}, \forall i \in A,$$

$$U_j(x,\cdot) = \phi_{\alpha_L} + p_B\eta, \forall j \in B,$$

$$U_j(y,\cdot) = \phi_{\beta_L}, \forall j \in B,$$

Therefore, each candidate chooses social group x, and the above described is an equilibrium. To see if this equilibrium is coalition-proof, consider a coalition of B types of mass ϵ deviating to y.

$$U_j(y_\epsilon,\cdot) = \phi_{\beta_L} + \eta, \forall j \in \epsilon.$$

Consider a sub-coalition of mass $\epsilon' < \epsilon$ deviating to x,

$$U_j(x_{\epsilon'}, \cdot) = \phi_{\beta_L} + \left(\frac{p_B(1+\epsilon'-\epsilon)}{p_B(1+\epsilon'-\epsilon)+p_A}\right)\eta, \forall j \in \epsilon'.$$

Comparing the above equations we say that the coalition is self-enforcing if and only if $\eta < \frac{p_B(1+\epsilon'-\epsilon)+p_A}{p_B(1+\epsilon'-\epsilon)}\underline{\phi}^m$. Moreover, the coalition is profitable with respect to the equilibrium if and only if $\phi_{\beta_L} + \eta > \phi_{\alpha_L} + p_B\eta \leftrightarrow \eta < \frac{1}{p_A}\underline{\phi}^m$.

Next, consider a mixed coalition of mass $\epsilon = \epsilon_A + \epsilon_B$, such that ϵ_A of A types chooses y, and ϵ_B of B types chooses y.

$$U_i(y_{\epsilon_A}, \epsilon_B, \cdot) = \phi_{\beta_L} + \left(\frac{p_A\epsilon_A}{p_A\epsilon_A + p_B\epsilon_B}\right)\eta, \forall i \in \epsilon_A,$$

$$U_j(y_{\epsilon_B}, \epsilon_A, \cdot) = \phi_{\beta_L} + \left(\frac{p_B\epsilon_B}{p_A\epsilon_A + p_B\epsilon_B}\right)\eta, \forall j \in \epsilon_B.$$

If $\epsilon_A = \epsilon_B$, $U_i(y_{\epsilon_A}, \cdot) < U_i(x^*, \cdot)$. If $\epsilon_A > \epsilon_B$, $U_i(y_{\epsilon_A}, \cdot) > U_i(x^*, \cdot)$ but $U_j(y_{\epsilon_B}, \cdot) < U_j(x^*, \cdot)$. Finally, if $\epsilon_A < \epsilon_B$, $U_j(y_{\epsilon_B}, \cdot) > U_j(x^*, \cdot)$, but $U_i(y_{\epsilon_A}, \cdot) < U_i(x^*, \cdot)$. Therefore, this coalition is never profitable $\forall \epsilon_A \gtreqless \epsilon_B$.

Therefore, no mixed coalition is profitable compared to the equilibrium, but a coalition of B types is profitable and self-enforcing if and only if $\eta > \frac{1}{p_A}\underline{\phi}$. We conclude that the pooling equilibrium on x and on L is also CPE if and only if $\eta < \frac{1}{p_A}\underline{\phi}$.

A similar reasoning applies to a pooling equilibrium on y s.t. $\mathbf{m_k} = (\beta_L, \alpha_L)$.

Separating equilibria on social groups.

Consider the case when DM has beliefs $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (1, 0)$. Thanks to Lemma 17, we know that it must be that in equilibrium $\mathbf{p}(\mathbf{s}|\mathbf{k}) = (0, 0)$. Hence, it must be that DM chooses d_0 and $\mathbf{m_k} = (\alpha_L, \beta_L)$. Given DM's choices, candidates all choose L (Lemma 16). Consider the case s.t. all A choose x, and all B choose y.

$$U_i(x, \cdot) = \phi_{\alpha_L} + \eta, \forall i \in A,$$

$$U_i(y, \cdot) = \phi_{\beta_L}, \forall i \in A,$$

$$U_j(x, \cdot) = \phi_{\alpha_L}, \forall j \in B,$$

$$U_j(y, \cdot) = \phi_{\beta_L} + \eta, \forall j \in B.$$

Hence, each i always chooses x, while each j chooses y if and only if $\eta > \underline{\phi}^m$, and the one described above is an equilibrium. Is it CPE?

To prove this, we have to show that there cannot be any coalition that deviates from the equilibrium that has a strictly higher utility than in equilibrium. Let us first consider a mixed coalition of ϵ_A A types and ϵ_B B types. Consider the case of $i \in \epsilon_A$:

$$U_i(y_{\epsilon_A}, \epsilon_B, \cdot) = \phi_{\beta_L} + \eta \left(\frac{\epsilon_A p_A}{\epsilon_A p_A + (\epsilon_B - 1) p_B} \right), \forall i \in \epsilon_A.$$

Since $U_i(y_{\epsilon_A}, \epsilon_B, \cdot) < U_i(x^*, \cdot)$, there would never be any mixed coalition profitable with respect to the separating equilibrium. Similarly, there could never be any coalition of A types alone. Now let us think about a coalition of B types with mass ϵ to social group x. In this case,

$$U_j(x_\epsilon, \cdot) = \phi_{\alpha_L} + \eta \left(\frac{\epsilon p_B}{\epsilon p_B + p_A} \right), \forall j \in \epsilon.$$

Consider a sub-coalition of mass $\epsilon' < \epsilon$ deviating to y.

$$U_j(y_{\epsilon'}, \cdot) = \phi_{\alpha_L} + \eta \left(\frac{(1 + \epsilon' - \epsilon) p_B}{(1 + \epsilon' - \epsilon) p_B + p_A} \right), \forall j \in \epsilon'.$$

The coalition is self-enforcing if and only if $\eta < \frac{(\epsilon p_B + p_A)((1+\epsilon'-\epsilon)p_B + p_A)}{p_B p_A (1 - 2\epsilon + \epsilon')} \underline{\phi}^m$. Moreover, such a coalition is profitable for each member of the coalition if $\eta < \frac{\epsilon p_B + p_A}{p_A} \underline{\phi}^m$. Therefore, the separating equilibrium on social groups and pooling on L such that all A types choose x, and all B types choose y is also CPE if and only if $\eta > \frac{\epsilon p_B + p_A}{p_A} \underline{\phi}^m$. A similar reasoning applies to the other separating equilibrium on social groups given the symmetry in utility functions.

Lastly, we show that there cannot exist any semi-separating equilibrium on social groups.

Consider a situation where DM has beliefs $\mathbf{p(t|k)} = (p(A|x), p(A|y))$, where $p(A|x) > p(B|x)$, and $p(A|y) < p(B|y)$, and $\mathbf{p(t|k)} = (0, 0)$. In this case, there is no equilibrium when DM buys the information; hence, we only consider the case in which she plays d_0, and $\mathbf{m_k} = (\alpha_L, \beta_L)$. In this case all candidates choose L, and to be an equilibrium it must be that a fraction q_A $(1 - q_A)$ of A types chooses x (y), and a fraction q_B $(1 - q_B)$ of B types choose x (y) such that $q_A > q_B$ $(1 - q_A < 1 - q_B)$. Each $i \in A$ earns $U_i(x, \cdot) = \phi_{\alpha_L} + \eta \left(\frac{q_x^A p_A}{q_x^A p_A + q_x^B p_B} \right)$, or $U_i(y, \cdot) = \phi_{\beta_L} + \eta \left(\frac{(1-q_x^A) p_A}{(1-q_x^A) p_A + (1-q_x^B) p_B} \right)$. Similarly, each $j \in B$ earns $U_j(x, \cdot) = \phi_{\alpha_L} + \eta \left(\frac{q_x^B p_B}{q_x^A p_A + q_x^B p_B} \right)$, or $U_j(y, \cdot) = \phi_{\beta_L} + \eta \left(\frac{(1-q_x^B) p_B}{(1-q_x^A) p_A + (1-q_x^B) p_B} \right)$. Therefore, there exist a semi-separating equilibrium on social groups, pooling on L if and only if

$$\begin{cases} \phi_{\alpha_L} + \eta\left(\frac{q_x^A p_A}{q_x^A p_A + q_x^B p_B}\right) = \phi_{\beta_L} + \eta\left(\frac{(1-q_x^A)p_A}{(1-q_x^A)p_A + (1-q_x^B)p_B}\right) \\ \phi_{\alpha_L} + \eta\left(\frac{q_x^B p_B}{q_x^A p_A + q_x^B p_B}\right) = \phi_{\beta_L} + \eta\left(\frac{(1-q_x^B)p_B}{(1-q_x^A)p_A + (1-q_x^B)p_B}\right) \end{cases}$$

Given that the solution to the above system involves $\frac{q_B}{q_A} > \frac{1-q_B}{1-q_A}$, which would violate weak-consistency, the above one cannot be an equilibrium.

A similar reasoning applies to the case when $\mathbf{p}(\mathbf{s}|\mathbf{k}) = (1,1)$, and DM buys the information.

□

Corollary 12.

Consider a game where there is homophily, and types and skills are either complement or substitute.

- *Under substitutability, if $p_B\tau < c < p_B\tau(1+\delta_L)$, there exists no pooling equilibrium on social groups. If $c > p_B\tau(1+\delta_L)$, there exist pooling equilibria on social groups, where all candidates play L, DM does not buy the information, and $\mathbf{m}_\mathbf{k}^* = (\alpha_L, \alpha_L)$. These equilibria are not CPE. If $c > p_B\tau(1+\delta_L)$, there exists one pooling equilibrium on x and one on y, where all candidates play L, DM does not buy the information, and $\mathbf{m}_\mathbf{k}^* = (\alpha_L, \beta_L)$ ($\mathbf{m}_\mathbf{k}^* = (\beta_L, \alpha_L)$). These equilibria are CPE if and only if $\eta < \frac{1}{p_A}\underline{\phi}^m$.*
- *Under complementarity,*
  - *if $p_B\tau < c < p_B\tau(1+\delta_L)$, there exist pooling equilibria on social groups, where all candidates play H, and DM buys the information. These equilibria are not CPE.*
  - *If $p_B\tau(1+\delta_L) < c < p_B\tau(1+\delta_H)$, there exist pooling equilibria on social groups, where all candidates play H, and DM buys the information, and there exist pooling equilibria on social groups where all candidates play L, DM does not buy the information, and $\mathbf{m}_\mathbf{k}^* = (\alpha_L, \alpha_L)$. These equilibria are not CPE.*
  - *If $p_B\tau(1+\delta_L) < c < p_B\tau(1+\delta_H)$, there exists one pooling equilibrium on x and one on y, where all candidates play L, DM does not buy the information, and $\mathbf{m}_\mathbf{k}^* = (\alpha_L, \beta_L)$ ($\mathbf{m}_\mathbf{k}^* = (\beta_L, \alpha_L)$). These equilibria are CPE if and only if $\eta < \frac{1}{p_A}\underline{\phi}^m$.*
  - *If $c > p_B\tau(1+\delta_H)$, there exist pooling equilibria on social groups, where all candidates play L, DM does not buy the information, and $\mathbf{m}_\mathbf{k}^* = (\alpha_L, \alpha_L)$. These*

equilibria are not CPE.

– *If $c > p_B\tau(1+\delta_H)$, there exists one pooling equilibrium on x (and one on y), where all candidates play L, DM does not buy the information, and $\mathbf{m}^*_{\mathbf{k}} = (\alpha_L, \beta_L)$ ($\mathbf{m}^*_{\mathbf{k}} = (\beta_L, \alpha_L)$). These equilibria are CPE if and only if $\eta < \frac{1}{p_A}\underline{\phi}^m$.*

D.2 Mixed Strategies Equilibria

In this section we state and prove the existence of all the Mixed strategies Nash Equilibria (MNE from now on) of our model. We begin by saying that we exclude from the results those MNE that exist only under knife edge conditions.

To formalise the results we need to define what a MNE is in our model. We call $\boldsymbol{\Upsilon}^* = (q^*_{xA}, q^*_{xB}, q^*_H)$ the vector of mixed strategies in equilibrium of candidates, where q^*_H is the probability with which candidates choose H in equilibrium, and q^*_{xt} is the probability with which type t candidates choose x in equilibrium.[1] Similarly, we call $\boldsymbol{\Sigma}^* = (q^*_d, \boldsymbol{\omega}^*_k, \boldsymbol{\omega}^*_s)$ the choices of DM in equilibrium, where q^*_d is the probability with which DM chooses d_0 in equilibrium, and $\boldsymbol{\omega}^*_k$ and $\boldsymbol{\omega}^*_s$ are the tasks choices in equilibrium.[2] I will write the value of c for substitutability and in parenthesis the value of c for complementarity.

D.2.1 Benchmark

Theorem 10.

Consider a game where there is not homophily, and types and skills are either complement or substitute.

- *Take $q^*_H < \frac{1}{2}$. If $c = \tau\left[p_B\left(1 + (1 - q^*_H)\delta_L\right) + q^*_H\delta_H\right]$, there exist infinite pooling equilibria on social groups s.t. $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (p_A, p_A)$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (q^*_H, q^*_H,)$, $\boldsymbol{\Sigma}^* = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}, (\alpha_L, \alpha_L), \boldsymbol{\omega}^*_s\right)$, and $\boldsymbol{\Upsilon}^* = (q^*, q^*, q^*_H)$. These equilibria are CPE.*
- *Take $q^*_H > \frac{1}{2}$. If $c = \tau\left[p_B + (1 - q^*_H)\delta_L\right]$ ($c = \tau\left[p_B\left(1 + q^*_H\delta_H\right) + (1 - q^*_H)\delta_L\right]$), there exist infinite pooling equilibria on social groups s.t. $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (p_A, p_A)$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (q^*_H, q^*_H,)$, $\boldsymbol{\Sigma}^* = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}, (\alpha_H, \alpha_H), \boldsymbol{\omega}^*_s\right)$, and $\boldsymbol{\Upsilon}^* = (q^*, q^*, q^*_H)$. These equilibria are CPE.*

[1] Given Assumption 3, A types B types randomize the skill with the same probability.

[2] Because we exclude knife edge conditions, we are automatically excluding the strategies where DM mixes within the tasks assignments, hence, we do not need to have mixed strategies inside the tasks assignments.

Proof.

We will prove the second bullet point and the argument stands for the first by symmetry of payoffs.

Firstly, let us start by showing the conditions for DM. In order for the first kind of MNE to happen DM must have these kind of beliefs: $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (p_A, p_A)$ and $\mathbf{p}(\mathbf{s}|\mathbf{k}) > (\frac{1}{2}, \frac{1}{2})$. Note that there are two cases in which DM can have this beliefs in equilibrium: one where $q_{xA} = q_{xB} = q$, and the one where $q_{xA} = q_{xB} = 1$ ($q_{xA} = q_{xB} = 0$) and $p(A|y) = p_A$ ($p(A|x) = p_A$). If this is the case, $\mathbf{m_k} = (\alpha_H, \alpha_H)$. In this case $E[V(\cdot, \mathbf{m_k}, d_0)|x] = E[V(\cdot, \mathbf{m_k}, d_0)|y] = p_A\tau + q_H\tau\delta_H$, if there is substitutability, and $E[V(\cdot, \mathbf{m_k}, d_0)|x] = E[V(\cdot, \mathbf{m_k}, d_0)|y] = p_A\tau(1+q_H\delta_H)$. In both cases, $V(\cdot, \mathbf{m}^*_\mathbf{s}, d_1) = \tau(1 + q_H\delta_H + (1 - q_H)\delta_L) - c$.

If $c > p_B\tau + (1 - q_H)\tau\delta_L$ ($c > p_B\tau(p_B + p_B q_H\delta_H + (1 - q_H)\delta_L)$), DM chooses d_0 and candidates all choose L and the one described above cannot be an equilibrium. On the opposite, If $c < p_B\tau + (1 - q_H)\tau\delta_L$ ($c < p_B\tau(p_B + p_B q_H\delta_H + (1 - q_H)\delta_L)$), DM chooses d_1, candidates all choose H and the one described above cannot be an equilibrium. Therefore, there exists MNE with the above characteristics if and only if $c = p_B\tau + q_H\delta$. Note that in this case, DM chooses to play d_0 with probability q_d, and when she plays d_0 she plays $\mathbf{m_k} = (\alpha_H, \alpha_H)$.

Secondly, consider candidates. If $q_{xA} = q_{xB} = q$, each candidate earns the same utility, no matter which social group they choose. Moreover,

$$U_i(q, H, \cdot) = \phi_{\alpha_H} - \zeta + \eta, \forall i \in A$$

$$U_i(q, L, \cdot) = q_d\phi_{\alpha_H} + (1 - q_d)\phi_{\alpha_L} + \eta, \forall i \in A.$$

The first equation equals the second if and only if $q_d = \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}$ (equally for B types). If that is the case, candidates choose H with probability $q_H \in (\frac{1}{2}, 1)$. Given that all candidates are indifferent between x and y, hence, $q_{xA} = q_{xB} = q$ in equilibrium. Similar reasoning can be done for the case in which all candidates choose x (y).

As for previous proofs under the benchmark, note that these equilibria are also CPE since there is no coalition that gives a strictly higher utility to the candidates in the coalition than in the equilibrium.

□

D.2.2 Homophily

We now consider $f\left(n_{k'}^{t''}, n_{k'}^{t'}\right)$ as indicated in Definition 7. We give a unique statement divided into four parts. The first concerns pooling equilibria on social groups, the second concerns separating equilibria on social groups with randomized skill, the third concerns semi-separating equilibria on social groups, and the fourth concerns semi-separating equilibria on social groups with randomized skill. With randomized skill, we mean equilibria where candidates choose $q_H^* \notin \{0,1\}$. Note that throughout the statements we will refer to $\overline{\phi}^m = \phi_{\alpha_H} - \phi_{\beta_H}$.

Theorem 11.

Consider a game where there is homophily, and types and skills are either complement or substitute.

Pooling equilibria on social groups.

- *Take* $q_H^* < \frac{1}{2}$. *If* $c = \tau\left[p_B\left(1 + (1-q_H^*)\delta_L\right) + q_H^*\delta_H\right]$, *there exist infinite pooling equilibria on social groups s.t.* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (p_A, p_A)$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (q_H^*, q_H^*,)$, $\boldsymbol{\Sigma}^* = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}, (\alpha_L, \alpha_L), \boldsymbol{\omega}_s^*\right)$, *and* $\boldsymbol{\Upsilon}^* = (q^*, q^*, q_H^*)$. *These equilibria are not CPE.*
- *Take* $q_H^* > \frac{1}{2}$. *If* $c = \tau\left[p_B + (1-q_H^*)\delta_L\right]$ $\left(c = \tau\left[p_B\left(1 + q_H^*\delta_H\right) + (1-q_H^*)\delta_L\right]\right)$, *there exist infinite pooling equilibria on social groups s.t.* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (p_A, p_A)$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (q_H^*, q_H^*,)$, $\boldsymbol{\Sigma}^* = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}, (\alpha_H, \alpha_H), \boldsymbol{\omega}_s^*\right)$, *and* $\boldsymbol{\Upsilon}^* = (q^*, q^*, q_H^*)$. *These equilibria are not CPE.*

Separating equilibria on social groups, with randomized skill.

- *Take* $q_H^* < \frac{1}{2}$. *If* $c = \tau q_H^*\delta_H$ *and* $\eta > \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\underline{\phi}^m$, *there exist two separating on social groups s.t.* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (1,0)$ *or* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0,1)$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (q_H^*, q_H^*)$, $\boldsymbol{\Sigma}^* = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}, (\alpha_L, \beta_L), \boldsymbol{\omega}_s^*\right)$ *or* $\boldsymbol{\Sigma}^* = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}, (\beta_L, \alpha_L), \boldsymbol{\omega}_s^*\right)$, *and* $\boldsymbol{\Upsilon}^* = (1, 0, q_H^*)$ *or* $\boldsymbol{\Upsilon}^* = (0, 1, q_H^*)$. *These equilibria are CPE if and only if* $\eta > \frac{p_A + \epsilon p_B}{p_A}\left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\underline{\phi}^m\right)$.
- *Take* $q_H^* > \frac{1}{2}$. *If* $c = \tau(1-q_H^*)\delta_L$ *and* $\eta > \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\overline{\phi}^m$, *there exist two separating on social groups s.t.* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (1,0)$ *or* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0,1)$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (q_H^*, q_H^*)$, $\boldsymbol{\Sigma}^* = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}, (\alpha_H, \beta_H), \boldsymbol{\omega}_s^*\right)$ *or* $\boldsymbol{\Sigma}^* = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}, (\beta_H, \alpha_H), \boldsymbol{\omega}_s^*\right)$, *and* $\boldsymbol{\Upsilon}^* = (1, 0, q_H^*)$ *or* $\boldsymbol{\Upsilon}^* = (0, 1, q_H^*)$. *These equilibria are CPE if and only if* $\eta > \frac{p_A + \epsilon p_B}{p_A}\left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\overline{\phi}^m\right)$.

Semi-separating equilibria on social groups.

- *Take* $p^*(A|x) \in (p_A, 1)$ *or* $p^*(A|y) \in (p_A, 1)$, *and* $\eta < \frac{1}{p_A}\underline{\phi}^m$.

If $c > \tau(1+\delta_L)o_x^*$ *or* $c > \tau(1+\delta_L)o_y^*$, *there exist infinite semi-separating equilibria s.t.* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (p^*(A|x), 0)$ *or* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0, p^*(A|y))$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0,0)$, $\mathbf{\Sigma}^* = (1, (\alpha_L, \beta_L), \boldsymbol{\omega}_s^*)$ *or* $\mathbf{\Sigma}^* = (1, (\beta_L, \alpha_L), \boldsymbol{\omega}_s^*)$, $\mathbf{\Upsilon}^* = (1, q_{xB}^*, 0)$ *or* $\mathbf{\Upsilon}^* = (0, q_{xB}^*, 0)$, *and* $\eta = \frac{p_A + q_{xB}^* p_B}{p_A}\underline{\phi}^m$. *These equilibria are not CPE.*

- *Take* $p^*(A|x) \in (p_A, 1)$ *or* $p^*(A|y) \in (p_A, 1)$, $q_d^* > \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}$, *and* $\eta < \frac{1}{p_A} q_d^* \underline{\phi}^m$.
  If $c = \tau(1+\delta_L)o_x^*$ *or* $c = \tau(1+\delta_L)o_y^*$, *there exist infinite semi-separating equilibria s.t.* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (p^*(A|x), 0)$ *or* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0, p^*(A|y))$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (0,0)$, $\mathbf{\Sigma}^* = (q_d^*, (\alpha_L, \beta_L), \boldsymbol{\omega}_s^*)$ *or* $\mathbf{\Sigma}^* = (q_d^*, (\beta_L, \alpha_L), \boldsymbol{\omega}_s^*)$, $\mathbf{\Upsilon}^* = (1, q_{xB}^*, 0)$ *or* $\mathbf{\Upsilon}^* = (0, q_{xB}^*, 0)$, *and* $\eta = \frac{p_A + q_{xB}^* p_B}{p_A} q_d^* \underline{\phi}^m$. *These equilibria are not CPE.*

- *Take* $p^*(A|x) \in (p_A, 1)$ *or* $p^*(A|y) \in (p_A, 1)$, $q_d^* < \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}$, *and* $\eta < \frac{1}{p_A} q_d^* \overline{\phi}^m$.
  If $c = \tau o_x^*$ *or* $c = \tau o_y^*$ ($c = \tau(1+\delta_H)o_x^*$ *or* $c = \tau(1+\delta_H)o_y^*$), *there exist infinite semi-separating equilibria s.t.* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (p^*(A|x), 0)$ *or* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0, p^*(A|y))$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (1,1)$, $\mathbf{\Sigma}^* = (q_d^*, (\alpha_L, \beta_L), \boldsymbol{\omega}_s^*)$ *or* $\mathbf{\Sigma}^* = (q_d^*, (\beta_L, \alpha_L), \boldsymbol{\omega}_s^*)$, $\mathbf{\Upsilon}^* = (1, q_{xB}^*, 1)$ *or* $\mathbf{\Upsilon}^* = (0, q_{xB}^*, 1)$, *and* $\eta = \frac{p_A + q_{xB}^* p_B}{p_A} q_d^* \overline{\phi}^m$. *These equilibria are not CPE.*

Semi-separating equilibria on social groups with randomized skill.

- *Take* $q_H^* < \frac{1}{2}$, $p^*(A|x) \in (p_A, 1)$ *or* $p^*(A|y) \in (p_A, 1)$, $q_d^* = \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}$, *and* $\eta < \frac{1}{p_A} q_d^* \underline{\phi}^m$.
  If $c = \tau\left[(1 + (1-q_H^*)\delta_L)o_x^* + q_H^*\delta_H\right]$ *or* $c = \tau\left[(1 + (1-q_H^*)\delta_L)o_y^* + q_H^*\delta_H\right]$, *there exist infinite semi-separating equilibria s.t.* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (p^*(A|x), 0)$ *or* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0, p^*(A|y))$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (q_H^*, q_H^*)$, $\mathbf{\Sigma}^* = (q_d^*, (\alpha_L, \beta_L), \boldsymbol{\omega}_s^*)$ *or* $\mathbf{\Sigma}^* = (q_d^*, (\beta_L, \alpha_L), \boldsymbol{\omega}_s^*)$, $\mathbf{\Upsilon}^* = (1, q_{xB}^*, q_H^*)$ *or* $\mathbf{\Upsilon}^* = (0, q_{xB}^*, q_H^*)$, *and* $\eta = \frac{p_A + q_{xB}^* p_B}{p_A} q_d^* \underline{\phi}^m$. *These equilibria are not CPE.*

- *Take* $q_H^* > \frac{1}{2}$, $p^*(A|x) \in (p_A, 1)$ *or* $p^*(A|y) \in (p_A, 1)$, $q_d^* = \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}$, *and* $\eta < \frac{1}{p_A} q_d^* \overline{\phi}^m$.
  If $c = \tau\left(o_x^* + (1-q_H^*)\delta_L\right)$ *or* $\tau\left(o_y^* + (1-q_H^*)\delta_L\right)$ ($c = \tau\left[(1 + q_H^*\delta_H)\, o_x^* + (1-q_H^*)\,\delta_L\right]$ *or* $c = \tau\left[(1 + q_H^*\delta_H)\, o_y^* + (1-q_H^*)\,\delta_L\right]$), *there exist infinite semi-separating equilibria s.t.* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (p^*(A|x), 0)$ *or* $\mathbf{p}^*(\mathbf{t}|\mathbf{k}) = (0, p^*(A|y))$, $\mathbf{p}^*(\mathbf{s}|\mathbf{k}) = (q_H^*, q_H^*)$, $\mathbf{\Sigma}^* = (q_d^*, (\alpha_H, \beta_H), \boldsymbol{\omega}_s^*)$ *or* $\mathbf{\Sigma}^* = (q_d^*, (\beta_H, \alpha_H), \boldsymbol{\omega}_s^*)$, $\mathbf{\Upsilon}^* = (1, q_{xB}^*, q_H^*)$ *or* $\mathbf{\Upsilon}^* = (0, q_{xB}^*, q_H^*)$, *and* $\eta = \frac{p_A + q_{xB}^* p_B}{p_A} q_d^* \overline{\phi}^m$. *These equilibria are not CPE.*

The coefficients $o_x^* \in (0,1)$ *and* $o_y^*(0,1)$ *are:*

$$o_x^* = 1 - p_A p^*(A|x) - q_{xB}^* p_B p^*(A|x) - (1 - q_{xB}^*) p_B$$

$$o_y^* = 1 - p_A p^*(A|y) - (1 - q_{xB}^*) p_B p^*(A|y) - q_{xB}^* p_B.$$

Proof.

Pooling equilibria on social groups.

We prove the second bullet point, and the argument for the first stands by symmetry of payoffs. The part of this proof concerning DM strategies is the same as the one in Theorem 10 and hence, we omit it. Concerning candidates, we already know from Theorem 10 that they will choose H with probability $q_H \in (\frac{1}{2}, 1)$ if and only if $q_d = \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}$. Since candidates are indifferent between H ans L, for the remaining part of this proof, will consider that they choose L for simplicity. Consider the case in which candidates choose x and y in the same proportion.

$$U_i(x, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\alpha_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\alpha_L} + \eta p_A, \forall i \in A,$$

$$U_i(y, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\alpha_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\alpha_L} + \eta p_A, \forall i \in A.$$

$$U_j(x, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\alpha_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\beta_L} + \eta p_B, \forall j \in B,$$

$$U_j(y, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\alpha_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\beta_L} + \eta p_B, \forall j \in B.$$

Given that $U_j(x, L, \cdot) = U_j(y, L, \cdot)$ and $U_i(x, L, \cdot) = U_i(y, L, \cdot)$, $q_{xA} = q_{xB} = q$, and DM beliefs will be correct in equilibrium. The same can be said for the case in which all candidates choose x (y).

However, none of these equilibria is CPE. Indeed, consider the case in which a fraction ϵ of B types who were choosing identity x chooses identity y.

$$U_j(y_\epsilon, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\alpha_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\beta_L} + \eta \frac{(1 - q + \epsilon) p_B}{(1 - q + \epsilon) p_B + (1 - q) p_A}, \forall j \in \epsilon.$$

Consider a sub-coalition of mass $\epsilon' < \epsilon$ deviating from the coalition and choosing x.

$$U_j(x_{\epsilon'}, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\alpha_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\beta_L} + \eta \frac{(q + \epsilon' - \epsilon) p_B}{(q + \epsilon' - \epsilon) p_B) p_B + (1 - q) p_A}, \forall j \in \epsilon'.$$

The coalition is self-enforcing given that $\frac{(1-q+\epsilon)p_B}{(1-q+\epsilon)p_B+(1-q)p_A} > \frac{(q+\epsilon'-\epsilon)p_B}{(q+\epsilon'-\epsilon)p_B)p_B+(1-q)p_A}$, for $\epsilon > \epsilon'$. Moreover, $\frac{(1-q+\epsilon)p_B}{(1-q+\epsilon)p_B+(1-q)p_A} > p_B$ and the coalition is profitable.

Therefore, all the pooling equilibria on social groups s.t. $q^*_{xA} = q^*_{xB} = q$ are not CPE. Similarly for the pooling equilibrium on x (y).

Separating equilibria on social groups with randomized skill.

We prove the fourth bullet point and the argument for the third stands by symmetry of payoffs. Firstly, consider DM. For the equilibrium to exist, it must be that $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (1,0)$ and $\mathbf{p}(\mathbf{s}|\mathbf{k}) = (q_H, q_H) > (\frac{1}{2}, \frac{1}{2})$. Given these beliefs, she chooses $\mathbf{m_K} = (\alpha_H, \beta_H)$, if she plays d_0. Specifically, $E[V(\cdot, \mathbf{m_k}, d_0)|x] = E[V(\mathbf{m_K}, d_0)|x] = \tau(1 + q_H\delta)$, both under substitutability and complementarity.

Given that $V(\cdot, \mathbf{m}^*_\mathbf{s}, d_1) = \tau(1 + q_H\delta_H + (1 - q_H)\delta_L)$, she chooses $q_d \in (0, 1)$ if and only if $c = \tau(1 - q_H)\delta_L$. If $c > \tau(1 - q_H)\delta_L$, DM chooses d_0 and all candidates choose L, therefore, there cannot be any equilibrium with $q_H > \frac{1}{2}$. If $c < \tau(1 - q_H)\delta_L$, DM chooses d_1 and all candidates choose H, therefore, there cannot be any equilibrium with $q_H < 1$.

Secondly, consider candidates' choices. We know from Theorem 10's proof that candidates choose $q_H \in (\frac{1}{2}, 1)$, if and only if $q_d = \frac{\overline{\phi}^s-\zeta}{\overline{\phi}^s}$. Since candidates are indifferent between H and L, we show the case in which they all choose L. Consider the case in which $q_{xA} = 1$ and $q_{xB} = 0$.

$$U_i(x, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\right)\phi_{\alpha_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\right)\phi_{\alpha_L} + \eta, \forall i \in A,$$

$$U_i(y, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\right)\phi_{\beta_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\right)\phi_{\alpha_L}, \forall i \in A.$$

$$U_j(x, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\right)\phi_{\alpha_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\right)\phi_{\beta_L}, \forall j \in B,$$

$$U_j(y, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\right)\phi_{\beta_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s}\right)\phi_{\beta_L} + \eta, \forall j \in B.$$

Each $i \in A$ chooses x $\forall \eta > 0$. However, B types choose y if and only if $\eta > \frac{\overline{\phi}^s-\zeta}{\overline{\phi}^s}\overline{\phi}^m$. Therefore, there exists a separating equilibrium on social groups, provided that the above condition holds.

To show the coalition proofness, consider a coalition of mass ϵ of B types choosing x.

$$U_j(x_\epsilon, L, \cdot) = \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\alpha_H} + \left(1 - \frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \right) \phi_{\beta_L} + \eta \frac{\epsilon p_B}{\epsilon p_B + p_A}, \forall j \in \epsilon.$$

A sub-coalition of mass $\epsilon' < \epsilon$ deviating to y would earn

$$U_j(y_{\epsilon'}, L, \cdot) = \phi_{\beta_L} + \eta \frac{(1 + \epsilon' - \epsilon)p_B}{(1 + \epsilon' - \epsilon)p_B + p_A}, \forall j \in \epsilon'.$$

The coalition is self-enforcing if and only if $\eta < \frac{(\epsilon p_B + p_A)((1+\epsilon'-\epsilon)p_B + p_A)}{p_B p_A (1-2\epsilon+\epsilon')} \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \underline{\phi}^m \right)$. Moreover, the coalition is profitable if and only if $\eta < \frac{\epsilon p_B + p_A}{\epsilon p_B} \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \overline{\phi}^m \right)$. Therefore, the separating equilibrium on social groups described above is CPE if and only if $\eta > \frac{\epsilon p_B + p_A}{\epsilon p_B} \left(\frac{\overline{\phi}^s - \zeta}{\overline{\phi}^s} \overline{\phi}^m \right)$.

Semi-separating equilibria on social groups.

We now prove the fifth bullet point, and the argument stands for the sixth and the seventh, thanks to symmetry in payoffs.

Consider the case in which DM has the following beliefs: $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (p(A|x), 0)$, with $p(A|x) \in (p_A, 1)$, and $\mathbf{p}(\mathbf{s}|\mathbf{k}) = (0, 0)$. In such a case, $\mathbf{m_k} = (\alpha_L, \beta_L)$. Therefore,

$$E[V(\cdot, \mathbf{m_k}, d_0)] = (p_A + q_{xB} p_B) \left(p(A|x) \tau (1 + \delta_L) \right) + (1 - q_{xB} p_B)(\tau(1 + \delta)).$$

With due simplifications, it can be shown that

$$E[V(\cdot, \mathbf{m_k}, d_0)] = \tau(1 + \delta_L) \left(p_A p(A|x) + q_{xB} p_B p(A|x) + (1 - q_{xB}) p_B \right).$$

Note that $1 - o^*_x = p_A p(A|x) + q_{xB} p_B p(A|x) + (1 - q_{xB}) p_B$. Therefore, given that $V(\mathbf{m}^*_\mathbf{s}, d_1) = \tau(1 + \delta_L) - c$, DM will chose d_1 if and only if $c < o^*_x \tau(1 + \delta_L)$. However, in this case, there exists no equilibrium since candidates will all choose H. If $c > o^*_x \tau(1 + \delta_L)$, DM chooses d_0, and if $c = o^*_x \tau$, DM mixes. We only show the proof for the first case, but the argument for the second stands for symmetry in payoffs.

Consider the case when DM chooses d_0. Let $q_{xA} = 1$ and $q_{xB} \in (0, 1)$.[3] If $q_d = 1$, we know

[3] Note that any $q_{xB} \in (0, 1)$ satisfies $\mathbf{p}(\mathbf{t}|\mathbf{k}) = (p(A|x), 0)$, hence, any $q_{xB} \in (0, 1)$ sustain a semi-separating equilibrium.

from Lemma 16 that $q_H = 0$. Therefore,

$$U_i(x, L, \cdot) = \phi_{\alpha_L} + \eta \left(\frac{p_A}{p_A + q_x^B p_B} \right), \forall i \in A,$$

$$U_i(y, L, \cdot) = \phi_{\beta_L}, \forall i \in A.$$

$$U_j(x, L, \cdot) = \phi_{\alpha_L} + \eta \left(\frac{q_x^B p_B}{p_A + q_x^B p_B} \right), \forall j \in B,$$

$$U_j(y, L, \cdot) = \phi_{\beta_L} + \eta, \forall j \in B.$$

All $i \in A$ will choose x. However, $U_j(x, L, \cdot) = U_j(y, L, \cdot)$ if and only if $\eta = \frac{q_x^B p_B + p_A}{p_A} \underline{\phi}^m$, therefore, semi-separating equilibria exist if and only if this condition holds.

Let us check for coalition proofness. Consider a coalition of mass ϵ of B types previously choosing y that now chooses x. They will all earn

$$U_j(x_\epsilon, L, \cdot) = \phi_{\alpha_L} + \eta \left(\frac{(q_{Bx} + \epsilon) p_B}{p_A + (q_{Bx} + \epsilon) p_B} \right), \forall j \in \epsilon.$$

If a sub-coalition of mass $\epsilon' < \epsilon$ deviates to y, they earn

$$U_j(y_{\epsilon'}, L, \cdot) = \phi_{\beta_L} + \eta, \forall j \in \epsilon'.$$

Given that the semi-separating equilibria exist if and only if $\eta = \frac{q_{Bx} p_B + p_A}{p_A} \underline{\phi}^m$ and that $\frac{(q_{Bx}+\epsilon)p_B}{p_A+(q_{Bx}+\epsilon)p_B} > \frac{q_{Bx} p_B}{p_A + q_{Bx} p_B}$, $U_j(x_\epsilon, L, \cdot) > U_j(y^*, L, \cdot)$, and $U_j(x_\epsilon, L, \cdot) > U_j(y_{\epsilon'}, L, \cdot)$. Therefore, the coalition is both self-enforcing and profitable with respect to the equilibrium and all these semi-separating equilibria are never CPE.

Semi-separating equilibria on social groups with randomized skill.

This part of the proof refers to the eighth and ninth bullet points. The logic of the proof is the same as the fifth to seventh bullet points for the semi-separation on social groups, while the logic for the randomization of skill is the same as the third and fourth bullet points. For this reason, we omit the proof.

□

Estratto per riassunto della tesi di dottorato

Studente: Roberto Rozzi matricola: 956470

Dottorato: Economics

Ciclo: 34

Titolo della tesi: “Four essays on social conventions”

Abstract:

In this thesis, I analyze how and why social conventions emerge across distinct contexts and their impact on different economic environments. In the first and second chapters, I employ evolutionary game theory techniques together with stochastic stability to study the formation of social conventions in two different strategic situations: a coordination game and a conflict game. In the third chapter, I conducted two field experiments to assess how people react to different shares of the population engaging in a particular behavior, which is crucial in understanding the emergence of social conventions. In the fourth chapter, I use a signaling game to study the relevance of homophily in guiding the formation of social groups among employees and the consequent impact on labor market outcomes.

In questa tesi, analizzo come e perché le convenzioni sociali emergono in contesti distinti e il loro impatto su diversi ambienti economici. Nei primi due capitoli, utilizzo tecniche di teoria dei giochi evolutiva insieme alla stabilità stocastica per studiare la formazione di convenzioni sociali in due diverse situazioni strategiche: un gioco di coordinamento e un gioco di conflitto. Nel terzo capitolo, ho condotto due esperimenti sul campo per valutare come le persone reagiscano a diverse percentuali della popolazione che si impegnano in un comportamento particolare, che è cruciale per comprendere l'emergere delle convenzioni sociali. Nel quarto capitolo, utilizzo un gioco di segnalazione per studiare l'importanza dell'omofilia nella guida alla formazione di gruppi sociali tra dipendenti e l'eventuale impatto sui risultati del mercato del lavoro.

Firma dello studente

Roberto Rozzi