UNIVERSITAT DE BARCELONA

Essays on the Economics of Social Dilemmas

Giorgos Papadomichelakis

PhD in Economics

Essays on the Economics of Social Dilemmas

Giorgos Papadomichelakis

UNIVERSITAT DE BARCELONA

2020

PhD in Economics | Giorgos Papadomichelakis

UNIVERSITAT DE BARCELONA

PhD in Economics

Thesis title:

Essays on the Economics of Social Dilemmas

PhD student:
Giorgos Papadomichelakis

Advisors:
Concepció Patxot
Francesc Dilmé

Date:
June 2020

UNIVERSITAT DE BARCELONA

Dedicated to my father, mother and brother

“Thoughts without content are empty, intuitions without concepts are blind.”

Immanuel Kant,
The Critique of Pure Reason

Acknowledgements

First and foremost, my deepest gratitude goes to my Ph.D. advisors, Concepció Patxot and Francesc Dilmé for the many years of support, guidance and patience as I went through this challenging but and enjoyable period of my life. Special thanks to Ció for going far and beyond the duties of an academic advisor and doing her very best for me to have everything I needed in order to develop my ideas and my skills as a researcher. Also a special mention to Francesc's guidance, insight, and immense patience, especially through the long emailing and skyping sessions, since the distance did not allow us many opportunities for face to face interaction. I could not have asked for better advisors.

I would also like to thank a number of researchers whose contribution could not be overstated. I am very thankful to Santiago Sánchez-Pages who was kind enough to provide me with letters of recommendation in my Job Market process. Santi was also instrumental for my involvement in experiments, thanks to his encouragement, advice and his work that inspired me. I also thank him for taking the time to comment on my work despite the distance and busy schedule. I owe gratitude also to Esther her letters of recommendation, for reviewing my work and giving much needed and appreciated feedback, and for always being a pleasant colleague during the long hours at the department. I am also very grateful to Jesús Marin, who has been my external evaluator, and gave his feedback and insight numerous times.

During my doctoral studies, I spend a short, but very important for me personally and academically period, on a research visit at the University of California – San Diego in the Spring of 2019. I had the pleasure of being hosted by Professor Simone Galperti, who despite his overwhelming schedule, always found the time to discuss my ideas with me, offer me his deep insight, and some kind words that made my stay on the other side of the world enjoyable and constructive. I am also thankful for valuable time discussing with professors Joel Sobel, Augusto Nieto Barthaburu and the PhD students of UCSD's Economic department.

My thanks to Meritxell Solé, Guadalupe Souto, Elisenda Rentería and Alfonso Herranz-Loncán, with whom, and together with my advisor Concepció Patxot worked for years on a project that ran parallel to my main PhD research, to create the National Transfers Account historical time series for Spain. I learned a lot while working together and I appreciate the time spent with them.

A very special group of people that I am grateful towards, are my fellow Ph.D. students from a number of universities for helping me develop my research ideas, but most importantly making my life the last few years better with their friendship. I thank Natassa Papadopoulou, Athina Raftopoulou and Gianko Michailidis that have been great friends and colleagues from the start of my Ph.D. life. Another big thank you goes to Jessica Ellis, Ting Ting Wu and Adrià Bronchal who shared their knowledge about how to set up and run an experiment and made their best to help me conduct my own. My thanks to Cosima Steck and Shohei Yamamoto with whom we founded the Behavioural Design Unit, a research group focused on nudging and policy design in the summer of 2019. Finally, I thank Carmen García Galindo, Alessandro Ferrari and Matic Petricek for being very good friends during the intense time I spend at the European University Institute in Florence, Italy.

I was lucky to make some very good friends throughout my life as a Ph.D student, and I am particularly thankful to Kinga Tchorzewska, Niclas Poitiers, Diego Ocampo, Tania Fernandez, Till Hollstein, Cesar Blanco, Arianna Garofalo, Chris Koser, Pierre Magontier, Bernard Moscoso, Mariona Mas, Carlos Suarez, Nicola Rubino, Manish Singh, Ferran Armada, José Luis Castillo and Cynthia Armas and all the other colleagues at the UB School of Economics whom I didn't get the pleasure of spending as much time with as I would have liked other than the occasional chat during our PhD and Beer seminars, but still appreciate a lot.

A very special mention has to be made of the director of our program Elisabet Viladecans-Marsal and the head manager Jordi Roca Solaneles, who I thank for all their assistance in every step of the way and commend them for doing a fantastic work in making this Ph.D. program what it is today, and never stopping in improving it. I also thank Anna Alsina Keith, for taking care of a large volume of administrative work on my behalf for all those years, while always being efficient and kind.

I am leaving the most important acknowledgment of all for last. I want to thank with all my heart, my father Theocharis, my mother Maria, and my brother Ilias. Your love, support and patience made everything good I have in my life possible. Σας ευχαριστώ για όλα, και σας αγαπώ πολύ.

Funding Acknowledgement

This research received institutional support from the Spanish Science and Technology System (project number ECO2016-78991-R) as well as from the European Commission [Joint Programming Initiative (JPI) “More Years, Better Lives (2016)”, project WELTRANSIM].
I would also like to thank the Sociedad Económica Barcelonesa de Amigos del País (SEBAP) for the mobility grant that facilitated my research stay at the University of California – San Diego. Finally, I am thankful to Fundació Bosch I Gimpera and the University of Barcelona for the management of my funding.

Giorgos Papadomichelakis

Barcelona, June 2020

Contents

List of Figures

List of Tables

1 Introduction

The individual decision-maker is the fundamental building block of an economic system, in the same way that particles are the fundamental element of a natural system. Just as the interaction of particles, produces a vast array on wondrous phenomena, so do humans when interacting, can produce interesting, complex, and often paradoxical results. Despite this commonality, physics and economics, belong to distinct scientific branches, (the former being a natural science, and the latter a social one). This distinction is made to a large degree in recognition of the fact that understanding, and predicting human behavior entails a number of limitations, both practical, and moral, which that natural sciences do not have to be concerned with. Despite that, the economic literature has made vast leaps forward, towards understanding human behavior, and modelling it a way that has explanatory power and predictive value. That being said, many open questions remain, regarding how humans make decisions, what motivates them, and how these decisions are affected by the presence of others.

The purpose of this dissertation is to study specific cases of all those aforementioned questions in a number of different settings, all of which share one common characteristic of particular economic relevance. The settings that will be studied all involve a two-ways interaction of an individual with the group of his peers. More specifically, the essays in this thesis are concerned with a problem classified in the literature under the umbrella term of, "Social Dilemma" (Midgley and Olson (1969), Kollock (1998)). A social dilemma is any strategic situation (in the game-theoretic sense of the term) in which there is a conflict between individual and collective interest. To be more precise, the defining characteristics of a social dilemma situation is that everyone collectively is better off if everyone cooperates towards a specific action, but each one individually is better off defecting, regardless of the actions of others. The famous "Prisoner's Dilemma" is a classic example showcasing this conflict of interest.

The reason why social dilemmas are interesting from an economist's perspective, is because this conflict of interest is at the heart of many issues of economic importance, such as labor productivity, public finance, crime and environmental policy just to name a few. Understanding the drivers of individual behavior, as well as the interaction of individuals within a group, and identifying problematic points to these interactions is key for the design of the appropriate policies, and this thesis aims to contribute towards that goal. This thesis focuses on two distinct cases of social dilemma, represented mathematically as two games. These two cases are a standard public good game augmented by the presence of a mechanism for the enforcement of cooperation, and a social learning game, where information can only be communicated among the individuals of the group through their actions. As explained before, what connects these two games, is the interdependence of actions and payoffs of their players, as well as the misalignment between what is best for the individual, and what is best for the collective, creating a classical problem of externalities.

These two cases are formally defined, their respective equilibria and other features are identified, and potential welfare-improving mechanism are discussed. An additional common element of these two studies, that constitutes another point of contribution of this body of work, is that I study the role of mechanisms that do not solely rely on the provision of extrinsic (monetary) incentives, as it is the standard approach of the economic literature [1]. In the public good game case, I model the role of endogenous social norms that reinforce cooperation among agents and how it interacts with the presence of a punishment mechanism for defectors. In the case of the social learning game, the mechanism to address the inefficiencies that I shall demonstrate, is one that relies on the strategic disclosure of available information, instead of a provision of payoff-relevant incentives to implement the socially optimal behavior. Finally, this thesis which is primarily involved with theoretical analysis of social dilemma problems, is supplemented with a laboratory experiment on social learning, aimed both at testing empirically the propositions of the theoretical work, and other hypotheses that remain ambiguous in the literature. In the following paragraphs, I shall briefly outline each of the 3 chapters individually, motivating them, and discussing the main results and the contribution of each one.

In the first essay of my thesis and chapter 2 of this book, I am looking into the role that social norms play in the decision of agents to cooperate, in a public good game (Samuelson, 1954). It is a common theme in the empirical literature of public good games, that there exists significant heterogeneity in the observations, when it comes to cooperation of the subjects, compared to what the theory would predict. It has been fairly established, that in any strategic interaction that does not have the characteristics of a fully competitive market (Smith 1962), standard 'homo oeconomicus' models fail to predict real world behavior. A large number of alternative theories has been proposed to

[1] see Bénabou and Tirole (2006) for a seminal work on the interaction of extrinsic and intrinsic motives on economic behavior

reconcile this discrepancy, from altruism (Andreoni, 1990), to spite and envy (Levine, 1998), inequity aversion (Fehr and Schmidt, 1999) to name a few. However, in a very influential study across different types of societies across the world, Henrich et al. (2001) found that the degree of heterogeneity in cooperative behavior can be best explained by the different prevailing norms in each society, which act as an informal enforcement mechanism.

Driven by these findings, this chapter attempts to model the effect of a norm that discourages free-riding, and see what equilibrium predictions it generates, compared to what the standard model predicts. Two additional characteristics add to the novelty of this paper. First, the effectiveness of the norm is endogenous to the level of the share of agents that actually cooperate. This captures the notion that in order for a norm to be effective, a sufficient share of the population has to adhere to it, since it does not rely on any form of tangible punishment. Second, this 'informal' form of enforcement coexists with an explicit mechanism that imposes a penalty for those caught free riding.

The main takeaway from the analysis of this model is that multiple equilibria with levels of cooperative behavior arise, for given sets of policy parameters. There seems to exist a complementarity between norms and extrinsic enforcement where the optimal amount of cooperative behavior, and hence social welfare, is achieved for moderately strict levels of punishment that seem to reinforce the effectiveness of the norms. A very strict or a very lax enforcement mechanism will make defection more desirable for the individual, hence decreasing the strength of the norm as well. In addition, I look into what level of cooperation is sustainable if it was to be determined endogenously through a majority-rule and compare I to the optimal level that a benevolent social planner would impose. I find that a political outcome will result in under-provision of public good, but lower levels of defection compared to the case of the social planner. The second and third essays of this thesis, that constitute the third and fourth chapter respectively, are very closely related to each other, since are both concerned with the study of social learning games (Banerjee (1992), Chamley (2003)) where individuals make choices about an unknown state of the world based on their private information, and on what they observe other's around them do. These two chapters can be viewed as the theoretical and empirical counterparts of a larger research project on social learning. As it has been established in the literature, interactions such as these are prone to certain inefficiencies when it comes to aggregating individual information and making the optimal decision. This is due to the fact that individuals fail to internalize the benefit towards others, when acting on their private information instead of simply imitating others. Since the issue appears to be in the way information is being fed-back to the group. The relevance of social learning in the majority of the choices an individual makes every day cannot be understated, especially with the degree of integration every person has with information sharing technologies, such as the internet, mobile devices, social media etc. In fact, the majority of information platforms, such as Google, TripAdvisor, Reddit, IMDB, all rely harness social learning, since in order to fulfill their primary purpose, which is to provide their users with accurate information, they rely on the same user base in order to obtain that same piece of information in the first place. Hence, making sure these mechanisms aggregate and disseminate information optimally is of the outmost importance.

Chapter 3 is concerned with the theoretical study of social learning through a game. I model a situation where agents are faced with a one-off choice whose payoff depends on an unobserved state variable. Agents are endowed with their own private information, but also turn to their environment (particularly, what others have done before them) to reach a better-informed decision. This simple game captures the essence of a large class of social interactions, where we turn to others in order to obtain information that will supplement our decision regarding a pay-off relevant decision. Social learning occurs in the majority of decisions where uncertainty is involved, since people will naturally seek information from others to reach a decision. This is even more prevalent nowadays that the internet and advances in communication exposes individuals to much more information about what others are doing. Hence, the question that arises is whether more information is always better, or does moderation works best, at least in specific settings?

As it turns out, social learning is prone to some inherent inefficiencies with the most prominent being, that information obtained from what others are doing d tends to overwhelm any possible private belief an agent might hold, making it rational for her to imitate the actions of others rather than do what their personal belief would dictate. Although there is no irrationality in how an individual act, agents fail to consider the positive externality that them acting on their signal and hence revealing to other may incur to the group (Bikhchandani et al., 1992), (Smith and Sørensen, 2000). Since it is impossible to tell whether an agent simply imitates others, or just happen to have information that agrees with them (Herding vs. Informational Cascades) the entire group can be led to ill-informed decision and hence outcomes. I study a setting of social learning such at the one described, identify the and characterize the inefficiency in terms of a conflict between private and public signals and propose an informational design approach (Bergemann and Morris, 2019) that mitigates those inefficiencies. I assume a benevolent social planner than discloses information strategically to the agents. I find that despite the fact that the incentives of the agents collectively and those of the Social Planner align, it is still optimal for the Planner to refrain from fully disclosing the available information to some of the agents in order to induce them to act upon their own information hence reveal it fully to the designer who can then make better recommendations to subsequent agents.

Chapter 4 is the final chapter of my thesis and as mentioned above, is the empirical counterpart to chapter 3. Specifically, I design and carry out a controlled, lab experiment to test the main proposition derived from chapter 3, that there exist a potential benefit in obfuscating information in a social learning environment, since this will create the incentive to individuals to rely more in their private information when making their decisions, hence their choices carrying more useful information once observed by others. In addition to testing the validity of this proposition, I test whether this result is in fact achieved through the process hypothesized in the previous chapter, by observing the effect of the treatment on the weight each individual assigns to each piece of information they possess (the private and the public information).

Social learning, has been extensively studied in the lab, starting from Anderson and Holt (1997), who were the first to study learning and herding behavior when observing the actions of others, and continues to be studied to the present day (Chandrasekhar et al., 2020), thanks to the fact that social learning is more pervasive than ever now thanks to technological advances, but also because the role of information and the formation of beliefs has seen a resurgence thanks to some important advances in the literature, such as Bayesian Persuasion (Kamenica and Gentzkow, 2011) and Information Design (Bergemann and Morris, 2019) to name a few. However, the question of how information (or lack thereof) affects the efficiency of social learning has not been addressed, to the best of my knowledge. This study also contributes to an open debate about whether individuals under- or over-weigh their own information, with the classical theory point towards the former (Bikhchandani et al., 1992), but a large portion of the empirical literature pointing towards the latter (Nöth and Weber, 2003).

Another contribution of the paper lies in the fact that its design, constitutes a synthesis of Çelen and Kariv (2004) and Filippis et al. (2017) with additional features not present in any of the two, such as the introduction of noise in the treatment group, and the use of survey question to test a number of secondary hypotheses. This allows us to distinguish two types of behavior that are observationally equivalent, but are crucial to the efficiency of social learning, namely cascade and herd behavior. This is achieved by splitting the decision of each subject into two stages, to allow to measure the effect of each component of information, and measure the effect as the difference between prior, interim and posterior beliefs.

2 Laws, Norms and Cooperation in a Public Goods Game

2.1 Introduction

There exists a disparity across social sciences as to how individual behavior is being modeled which is mainly a reflection of the problems that each discipline traditionally is concerned with. On one extreme I have the *homo economicus* whose behavior is guided by instrumental rationality (i.e. behavior that leads to the optimal outcome based on a set of preferences) and can mainly be motivated by economic incentives. On the other hand I have the *homo sociologicus* (Benoit-Smullyan et al., 1938) whose decision making is led chiefly by social norms, which are the prescribed rules of acceptable behavior within the society along with an associated sanctioning system. Following the widespread introduction of controlled experiments to the study of behavior it appears that which model works best has mainly to do with the institutional setting under which the agents are operating. For example, the neo-classical homo economicus approach seems to work extremely well at predicting behavior in situations that resemble a fully competitive market [1] with minimal interaction among agents, whereas in situations of market failure (externalities, moral hazard etc.), and generally situations of strategic interdependence among agents, such models fail to fit empirical data [2]. What is being identified as the main contributors for this failure by the literature is mainly two: one the presence of social preferences, i.e. that the way preferences are modeled is incomplete, or the way incentives are, citing social norms as an informal incentive that needs to be studied further (Fudenberg and Levine, 2016). [3]

When studying cooperative behavior, an inherently socially-oriented activity with significant economic implications, it is far more likely that both extrinsic and intrinsic motives come into play in determining individual behavior. The simplest Prisoners' Dilemma game is enough to highlight the

[1] Seminal experimental study by Smith (1962)

[2] For an extensive survey of related literature the reader is referred to Fehr and Schmidt (2006)

[3] For a comprehensive discussion on the role of norms see Elster (1989)

potential disparity between individual and social optimality, a phenomenon usually referred to as "social dilemma" where every agent is better off when everybody cooperates, but each agent individually by defection. Many theories have been put forward to try and explain why do people cooperate, I will be highlighting the role of social norms, and how they work with explicit sanctions to enforce cooperation. Formally, a norm can emerge in a situation of strategic interdependence in which the decisions of individually rational actors lead to an inferior outcome for all or some parties, than the decisions of "collectively rational" actors. Social dilemma cases are such situations where the need for enforcement mechanisms arise, and social norms is one of them.

An intuitive explanation why norms can work in implementing compliance is that those who enforce the norms are closer to the individual compared to the law-enforcer. Proximity implies better information about circumstances and the behavior of other individuals therefore more likely to identify a free-rider. Individuals who deviate from social norms are discomforted by either "shame" emanating from the disapproval of others and/or "guilt" when they internalize the norm. Hence, a norm is an enforcement mechanism that does not require a centralized authority to implement it.

I focus on two types of choice related to public good provision, an economic and a political one since it hardly exists a market that determines a competitive price for such a good given that it is non-exclusive and beyond the means of a single individual. Contributing to a public good that benefits not just the individual but also the group at large appear to be one of those situations.

Following Becker's seminal work (Becker, 1974) I model the interplay of those two forces on enforcing cooperation on a public goods game where individuals are faced with the so-called "social dilemma". On top of the legal sanctions taken against law-breaker, there seems to exist in most societies an implicit social norm that atomistic behavior (in our case expressed as free-riding) is looked down upon. This normative pressure can be more or less strong and can take the form of "shame" or "guilt". For the sake of simplicity I will not distinguish between the two in this paper.

I assume that the intensity of the social norm against free ride depends in part on how many others in one's peer group choose to comply to it. Bentham (1789) referred to this effect as the syndrome of robberies without shame: "where robberies are frequent, and unpunished, robberies are committed without shame", p. 156. Free riding is less shameful the more others engage in the same activity. The intensity therefore of the norm is endogenous to our model[4] , i.e. the fraction of the population adhering to it is positively related to

[4] Fehr et al. (1997), find evidence supporting the idea of endogenously determined strength of the norm.

its intensity (the intensity of the norm becomes an equilibrium phenomenon, similar to the approach by Lindbeck et al. (1999).

Within this modeling framework, I first analyze individual decision-making and equilibrium outcomes under given policy parameters concerning level of contribution and law enforcement measures. Each individual then decides whether or not to engage in free-riding, and these decisions are based on knowledge of relevant policy parameters and on expectations about the going rate of free-riding. A level of free-riding is an equilibrium outcome if no individual wants to change her individual choice given this rate and establish conditions for existence and uniqueness of equilibrium. I also illustrate the equilibrium outcome and its comparative statics properties through graphical representations.

Having studied the equilibrium individual behavior, I turn to the question of how policy in determined. Agents choice is aggregated through a political process of simple majority voting, the level of contribution to a public good and are then are called to decide whether they are going to comply, or not (free-ride). The setup chosen is that of a public good game, because it best captures the situation of "social dilemma", a situation where the collective and the individual incentives are in conflict. Everyone is better-off when a certain public good is provided, but each agent is individually better off by not paying her own share of the cost and choose to free-ride instead.

The objective of this paper is to study how formal (laws) and informal (norms) institutions work to overcome the free-riders problem that is inherent in public good games and how they interact with each other. Incorporating the political process in the model it would also allow us to give theoretical support to experimental findings that participation rights foster cooperation observed in a large number of experiments. [5]This will be made possible by comparing the result of the voting process (political equilibrium) with that which would be expected should the decision was made by a social planner who is maximizing social welfare (Normative Benchmark). This law will then be taken as a given by each agent and they will rationally decide whether to cooperate or not. The social norm in this setting will be that free-riding is looked down up, causing disutility to the free-rider, reversely proportional to the share of the population that is also free-riding. The model in its simplest form can be summarized as follows: in the 1st Stage, agents express their policy preferences as voters, anticipating the consequences of chosen policy for their own economic choice and for the aggregate behavior-including the adherence to the norm. and in the 2nd Stage, agents maximize their utility subject to contribution, level of public good provision, norm probability of getting caught etc.

[5]For a representative experimental work, see Haigner et al. (2006)

The remainder of this paper is organized as follows: Some related literature is discussed in the next section. The basic model is presented in section 3, and the individual decision-making is further analyzed on Section 4. In Section 5 I move on to the policy analysis of the game and I discusses the outcomes between a normative social optimum, and a result from a voting process of the agents. I conclude in Section 6 with a summary of the main contributions and results along with a discussion of future extensions.

2.2 Overview of the Literature

The formal investigation of cooperation within groups of individuals can arguably trace its origins in Midgley and Olson (1969) where an analysis both analytically and with sociological and historical approaches of how groups (of different sizes and composition) tackle the issue of cooperation, the limitations that emerge and possible solution was first tackled. The issue has been more recently revisited by Tabellini (2008) where the author summarizes all the approaches up to that point, and proposes his own model of cooperation where private values individuals hold are highlighted as the main motive for them to engage in costly pro-social behavior.

Many different hypotheses have been suggested to explain the empirical fact that I observe far greater levels of cooperation than what the neo-classical models of economic behavior would predict. In (Becker, 1974) the idea of altruism as the main drive for cooperation was first discussed, which also gave way to the large strand of literature that is still expanding as of this day on social preferences. On that path, other theories with significant support are those of "Warm Glow" by Andreoni (1990) where public good provision was also used as an setting for his suggested theory. Other suggested motives for cooperation are reciprocal altruism by Levine (1998), inequity aversion by Fehr and Schmidt (1999), self-image concerns by Bénabou and Tirole (2006) and more recently peer-pressure by Levine and Modica (2016).

The role of norms as an incentive have been studied extensively as well, being cited as the main source for the observed heterogeneity in levels of pro-social behavior by influential experimental studies such as Henrich et al. (2001). Many attempts have been made to model norms explicitly especially on a game theoretic models. Okuno-Fujiwara and Postlewaite (1995) propose a Nash equilibrium that relies on substantially less information than common knowledge of the game which they name Norm Equilibrium. In their setting a social norm is defined as Social Standard of Behavior and a Transition Mapping. Takahashi (2010) show how cooperation can be sustained in a large community in a prisoner's dilemma repeated game by having cheaters punished not by the victims but by third parties.

A first attempt to incorporate specific social norms in economic analysis was done by Akerlof (1980) who studied the role of social customs in a model of fair wages and unemployment.Young (1993) formulated social norms as *conventions*: A pattern of behavior that is customary, expected and self-enforcing. Convention is an equilibrium everybody expects and (Bernheim, 1994) formulates the social norm of conformity where utility over consumption and social status where status is derived by public predisposition about the individual instead of his action. Predisposition is unobservable, actions signal predisposition therefore affect status. And in a recent paper, Ali and Miller (2016) look at a very specific norm with long history, which is ostracism.

Our formulation derives many elements from Lindbeck et al. (1999) who study the role of norms in deciding the size of the welfare state in an economy with a setup derived from works such as Meltzer and Richard (1983) in their seminal paper suggesting a rational about the size of the government in terms of redistribution.

The interaction of formal and informal incentives that I study in this paper have also been extensively featured in recent literature. In Benabou and Tirole (2011) the authors show how co decisions and public policies are shaped by personal and societal preferences ("values"), material or other explicit incentives ("laws") and social sanctions or rewards ("norms"). In a more recent study, Acemoglu and Jackson (2017) study the interplay of laws and norms in a game where agents choose their behavior and then are matched randomly with other agents in their group. On top of the probability of being caught by the law an additional enforcement mechanism is the ability of agents to give in their match if they observe her breaking the law while they do not. They find cases where laws and norms actually contradict each other and lead to more law-breaking than intended.

The assumption that I rely in our model that people have preferences or social approval has been supported by experimental evidence numerous times. For example, Gächter and Fehr (1999) conducted an experiment where in the treatment group, the identity and the amount of contribution of each subject was revealed to others after the end of a round giving them an incentive for social approval. When compared to the control group where the subjects maintained their anonymity they find a significant effect on the behavior of the treatment group, particularly among people who knew each other.

Finally, the motivation to test theoretically the idea that deciding endogenously on the level of public good provision leads to higher levels of participation came from experimental studies such as Haigner et al. (2006) where a public good game is played with the treatment effect being that agents in some groups are allowed to decide the enforcement mechanism endogenously while others are being given a scheme by the experimenters (exogenously).

The main finding is that participation rights foster cooperation i.e. a inclusive norm leads to higher participation. Similar findings were made in Dal Bó et al. (2010) also.

2.3 The Model

Consider a society populated by a continuum of agents playing a type of public goods game. Each individual is identically endowed and is faced with a binary choice: either to *cooperate* by contributing a fixed amount c of their endowment E towards the public good, or *free-ride* in which case they either get away with it and spend their entire endowment for private consumption with probability $(1-\pi)$ or get caught with probability $\pi \in [0,1]$ and are forced to pay a penalty p.

I assume that $E \geqslant p > c$ i.e., that the penalty that a free-rider is faced with if caught is strictly larger that the contribution that is asked of him as a cooperator, since otherwise the law would act as an incentive for free-riding instead of a deterrent (The same way that the fine on public transport if caught without a ticket is orders of magnitude largers than the price of the actual ticket.)

Taking into account the standard definition of a public good as being non-excludable and non-rivalrous, I assume that all agents cooperators or free-riders (even those who are caught) derive utility from the public good and the same amount is available to all.

All the above constitute the economic motives of behavior for an agent. However, on top of that, I assume that a social norm against free-riding is prevalent in the society, meaning agents who do not cooperate also experience a disutility associated with their actions that is independent of whether they are actually caught free-riding or not. This could be either be interpreted as an internalized feeling of guilt by doing something that goes against the societal interests, or the shame or negative reputational effect of being observed by peers performing something that poses an externality to others. The intensity of the disutility is endogenous to the level of free-riding which will be a source for multiple equilibria of free-riding, however as discussed above, it is an effect commonly observed both the the lab and the real world, therefore it seems to be a motive that affects behavior in settings such as the one I study. Agents are heterogeneous in the degree on importance they attach to this non-economic motive, again in line with empirical evidence that people exhibiting varying degrees of pro-social behavior, from full anti-social in one extreme to fully pro-social in the other and a distribution of agents across that range. I define that as a parameter $\theta \in (-\infty, +\infty)$ and I will refer to it as the degree of pro-sociality.

Additionally, agents also differ on the degree $\eta \in [0,2]$ which they care about the level of public good provided in relation the private consumption relative to private consumption which is normalized to 1 (therefore, $\eta = 1$ means that agents care equally about the two types of consumption goods), this is meant to capture the real world situation that on top of the aggregate social welfare, public good can have a heterogeneous impact on the welfare of particular individuals. For example, the building of a bridge can be beneficial for an economy, but its impact is felt more strongly by those who use it to commute to work everyday, that those who use it occasionally for recreational purposes.

The generic form of the value functions associated with each of the two choices can be summarized as follows:

$$U_c = u(E-c, \eta G) \tag{2.1}$$

if the agent chooses to cooperate, and

$$U_f = (1-\pi)u(E, \eta G) + \pi u(E-p, \eta G) - \theta n(x) \tag{2.2}$$

if the agent chooses to free-ride.

Where $u : \mathbb{R}_+^2 \to \mathbb{R}_+$ is the the utility function with two arguments related to consumption of goods (private and public), increasing in both arguments, and its the same regardless of the individual's state (cooperator, free-rider caught, free-rider at large) and $n(x) : [0,1] \to \mathbb{R}_+$ is the disutility experienced only by the free-rider which is non-increasing in its argument x. For the sake of tractability, I assume that all the agents have the same sub-utility functions u and n. Agents' vary only in their 'type' which is a vector of the two parameters I just discussed (θ, η). Additionally, I assume agents have 'rational expectations', meaning that they can correctly predict the level of free-riding x^*.

2.4 Equilibrium Free-Riding

Agents are rational, therefore each one picks the alternative with the highest expected utility. So an agent of type (θ, η) will choose to free ride if and only if,

$$u(E-c, \eta G) < (1-\pi)u(E, \eta G) + \pi u(E-p, \eta G) - \theta n(x) \tag{2.3}$$

We can see that for a given policy parameters, i.e. level of individual contribution c, level of public good G, probability of getting caught π, penalty p and level of free-riding x there exists a unique critical value of degree of pro-sociality θ^* which is determined at the point where an agent is indifferent between the two alternatives, namely where (3.2) becomes an equality.

$$\theta^* = w(\theta, x) = \frac{\Delta u}{n(x)} \tag{2.4}$$

where,

$$\Delta u = (1-\pi)u(E, \eta G) + \pi u(E-p, \eta G) - u(E-c, \eta G) \tag{2.5}$$

Δu is the expected consumption utility gain an agent can receive from free-riding given the expected equilibrium level of free-riding x^*. By rearranging the terms 2.5 we get a more clear image of how the economic motives:

$$\Delta u = u(E, \eta G) - u(E-c, \eta G) - \pi[u(E, \eta G) - u(E-p, \eta G)]$$

where the first two terms are the gain from being a free-rider that doesn't get caught vs. being a cooperator, and the second two the expected loss by being a free-rider and getting caught.

As showing in equation (2.4) who free-rides or not is dictated from both the explicit incentives (law) and the implicit (social consideration). We can see moreover, that the decision to cooperate or free-ride is both dependent on the agent's own type, as well as the expectation regarding what the others will do expressed here as the expected level of free-riding x^*. Agents whose type is such that $\theta < \theta^*$ will ultimately free-ride, and all others will cooperate (given that we have a continuum of agents we can safely ignore the possibility of indifference). Unsurprisingly,how much the individual values the public good (the parameter η) is irrelevant at the individual level of decision-making since the good is non-excludable and the marginal contribution of each individual is insignificant towards the total good, thus highlighting the core of the 'social dilemma' problem in this particular setup.

We make the assumption that individuals all make their choice simultaneously for given policy parameters. A strategy profile here is one that contains all the individual choices for a given game. A strategy profile is a *Nash equilibrium profile* if and only if every individual's choice is optimal, given the choices of all the others'. For a free-riding level x to be an *equilibrium free-riding level* it has to be consistent with a Nash equilibrium profile. Let Φ be the cumulative distribution of types (θ, η) in the population and x be the free-riding level that is associated with it. To establish the existence of an equilibrium x we must first the define

$$F(x) = \int_{\eta=0}^{\eta=2} \int_{-\infty}^{\theta^*} \Phi(\theta, \eta) \tag{2.6}$$

Proposition 1: A free-riding level $x \in [0,1]$ is a Nash free-riding level if and only if it satisfies the fixed-point equation $x = F(x)$ where $F : [0,1] \rightarrow [0,1]$ is defined in (2.6) and there exists at least one Nash equilibrium.

Proof. For any expected level of free-riding $x^* \in [0,1]$ the subset of individual types (θ, η) that choose to free-ride is Borel measurable, by equation (2.4), and the population share is equals to $x = F(x^*)$. Hence a strategy profile is a Nash equilibrium if and only if $x^* = F(x^*)$. F being continuous has at least one fixed point □

In short, $F(x^*)$ is the level of free-riding that occurs if all individuals expect the free-riding to be x^*. Equation $x = F(x)$ expresses that requirement, i.e. that individuals have 'rational expectations'.

Due to the relation of the intensity of the norm with x we can expect multiple equilibria as it will be demonstrated below. Only, in the special case where the sub-utility n is a constant we get a unique equilibrium level of free-riding as we can see in Figure 2.2.

Figure 2.1: Equilibrium Free-riding level with endogenous norm intensity

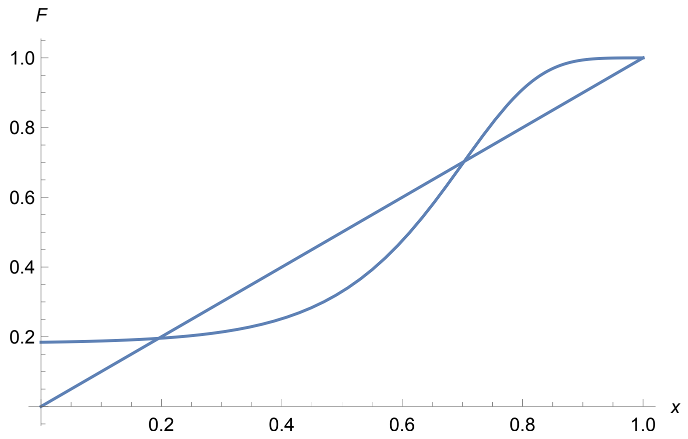

In Figure 2.1 we can see graphically the result of having an endogenous intensity of norm. In this particular example we can see 3 equilibrium levels of free-riding. The multiplicity of equilibria has an intuitive explanation that can offer some explanation as to real-world occurrences. When free-riding

Figure 2.2: Equilibrium Free-riding level with exogenous norm intensity

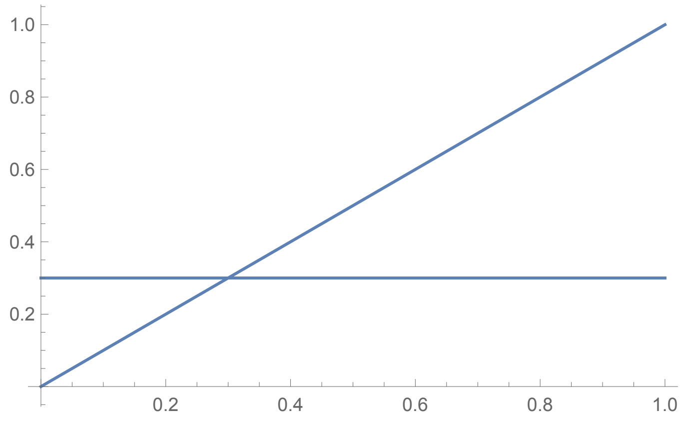

is low, then also the disutility experienced by free riding is high, therefore only those with very low consideration about social norms would choose to comply in our particular setup. Conversely, when free-riding is prevalent, even more pro-social individuals (those of high θ) would choose to be free-riders. The main takeaway from this graph is that even in the presence of similar policy parameters (e.g. a similar legal framework) the policy outcomes could vary significantly, an common occurrence both in the field and the lab. We highlight here how norms can play a big role in explaining such paradoxical outcomes.

We now move on to examine how the equilibrium free-riding is affected by changes in Δu. This will allow us to see how the economic incentives and therefore the law affects the level of cooperativeness under the presence of the norm.

Figure 2.2 highlights the effect of the presence of the norm on the levels of free-riding. Free-riding increases as the utility gains from free-riding increase but for an interval of Δu around 0.1 and 0.4 approximately, there exists a fold where the norm leads to multiple equilibria. In the classical case without a norm, that figure would simply be a step function jumping from 0 to 1 as soon as the gains from free-riding turn positive.

The presence of the fold for moderate levels of utility gain from consumption gives us some interesting insights as the complementarity of laws and norms, and how moderate policy interventions to deter free-riding can lead to significant decreases in the equilibrium level x, a result that is also corroborated by the findings in Acemoglu and Jackson (2017) where in a similar setup where laws and norms coexist found that a gradual change in laws that are not contradicting the prevailing norm strongly lead to higher levels of compliance than an abrupt tightening of the legal enforcement. It is evident that the reverse also holds, namely that a relatively small loosening of law enforcement could

Figure 2.3: Correspondence between NE level of free-riding to changes in consumption utility

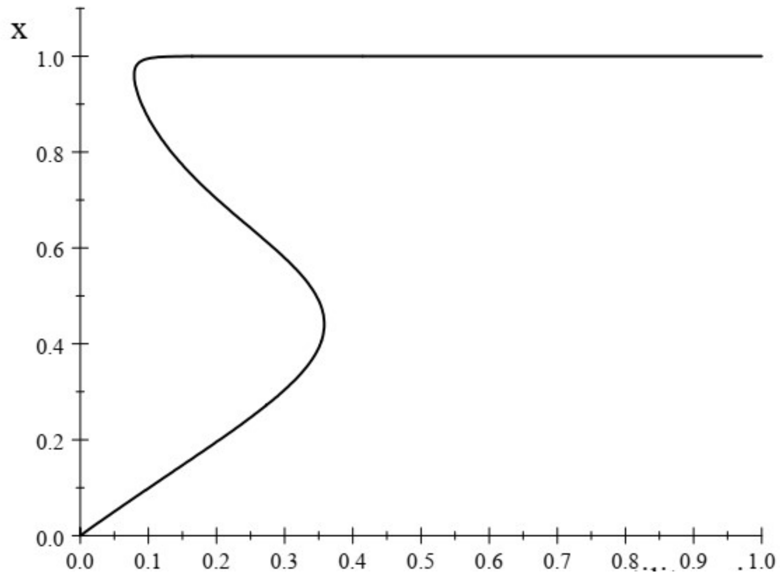

lead to a disproportionately high increase in free-riding.

We conclude the study of decision making at the individual level by looking into how each of the policy parameters affect the consumption gains from free-riding. In terms of comparative statics, its straightforward to see that Δu is decreasing in π and p and increasing in c, meaning that a stricter law in terms of probability of getting caught and/or harsher penalties deters free-riding whereas an increase in the required contribution makes free-riding more profitable.

2.5 Policy Analysis and Political Equilibrium

After analyzing individual decision making, we can extend this model for policy analysis, and further capture the inherent conflict of individual and social optimum in a social dilemma situation by predicting policy as an outcome of a democratic voting process vs the policies that would maximize a given welfare function (positive vs. normative analysis).

But first, we have to establish a connection between the individual contribution c and the amount of public good that agents will enjoy G. For the sake of simplicity we will focus only on the set of policy outcomes that maintain a balanced budget i.e. the provided public good is simply the sum of the contributions of the cooperators (Of course in reality budgets are not always balanced, but such a study would be more relevant on an inter-temporal setting of this particular model instead of the one-shot game we study here)

As an extra twist to the model we introduce costly law enforcement, i.e. the cost is increasing in π, since monitoring and implementing punishment requires resources especially in larger groups and societies. In this setup, there are two potential policy instruments to prevent free-riding, probability of getting caught π and the penalty to be paid in this case p.

So we assume a formal institution (e.g. a government) that handles the provision of the public good and the enforcement of the law. The budget of that institution is shaped as follows:

$$R(c,x,\pi) = \int_{\theta^*}^{+\infty} c d\Phi = G + C(\pi,x) \tag{2.7}$$

In words, the revenue is the contribution of all the cooperators, which is used to fund the enforcement mechanism, with the residual being the amount of public good to be provided. We define a quintuple $s = (c, G, \pi, p, x)$ as a *state* that satisfies the budget equation (2.8) and the fixed-point equation $x = F(x)$ mentioned above, is a *balanced-budget equilibrium state* denoted by $s^* \in S^*$ where S^* is the set all the possible balanced-budget equilibrium states. By defining a social welfare state, which will be done later, we can perform a normative analysis about the policy mix that maximizes said function.

Now to move towards a positive analysis, and specifically what would happen if the policy outcomes were endogenously decided on by the agents (as it is the case in most society, either via direct or representative democratic processes) we first have to define what a *political equilibrium* is. Specifically, we have to look into the set S^* defined above and see is there exists a policy $l = (c, \pi, p)$ that is a political equilibrium, meaning that there does not exist an other policy that is preferred a majority of the voters (i.e. more than $1/2$ of them). All voters know the current level of free-riding under incumbent policy $'$ and they can predict the free-riding rate x' under any other candidate policy l'. The expected utility of an individual of type (θ, η) in any state is:

$$U(s,\theta,\eta) = max\{U_c, U_f\} \tag{2.8}$$

where U_c and U_f were as defined in (3.3) and (3.4) respectively. Namely, each individual anticipates to make an optimal individual choice for any given state s.

Finally, we define a policy $l = (c, \pi, p)$ to be a *political equilibrium policy* or a *Condorcet Winner*, as it is often referred to, under majority rule, if the incumbent state $s = (c, G, \pi, p, x)$ is a balanced equilibrium state and is preferred by a (weak) majority over any other alternative state $s' = (c', G', \pi', p', x')$ that is if:

$$\int_{\eta=0}^{\eta=2}\int_{-\infty}^{\theta^*}\mathbb{I}[(s,\theta,\eta)-U(s',\theta,\eta)]d\Phi(\theta,\eta)\geqslant\frac{1}{2}\forall s'\in S^* \tag{2.9}$$

Where $\mathbb{I}$ is the indicator function where $\mathbb{I}=1$ for $x\geqslant 0$ and $\mathbb{I}=0$ otherwise.

The mathematical machinery outlined above can help us identify the conditions under which an equilibrium policy exists that is also a political equilibrium which is $S^p\subset S^*$ i.e. those equilibrium states that also satisfy equation (2.9). Note that the set S^p can be potentially empty. I shall now discuss an example case, to showcase the use of this setting

Political Support for Cooperation: An Example

I assume that all agents have a logarithmic utility function with regards to consumption, and that are also motivated by the existence of the norm, which is the parameter η discussed above. I assume it is distributed according to an arbitrary continuous cumulative probability distribution function H. To simplify, agents either full free-ride or fully cooperate. Moreover, I make the simplifying assumption that the fine paid if caught does not go back to the budget (even though it is a plausible and interesting scenario to analyze). The probability of getting caught is fixed (parameter π) and the cost of maintaining the enforcement mechanism is an affine increasing function of the level of free-riding: $K=K_0+kx$ where $K_0>0$ a fixed cost and $K>0$ the marginal cost of enforcement. So in this example, the capacity for enforcement can be adjusted to changes in the level of free riding.

Therefore, for any level of free-riding x, the public revenue is derived only from one source, those who contribute, and this revenue has to finance the expenses:

$$(1-x)c=K_0+kx+G \tag{2.10}$$

Where $K_0+G>0$ is fixed. The only remaining policy variable is c, the contribution to be paid, and p the penalty to pay if caught. G being the public good that has to be provided.

Combining the public budget balance equation with the equilibrium equation 2.5 we obtain the following equation relating c and p

$$\frac{c(1-x)-K_0+G}{c(1-x)+k}=H\left(\left[u\left(\frac{c(1-x)-K_0+G}{c(1-x)+k}\right)\right]^{-1}\left[ln(1-p)-\frac{1-\pi}{\pi}ln(1-c)\right]\right. \tag{2.11}$$

Note that for a maximal penalty of $p=1$, the consumption utility of someone caught is minus infinity. Hence free riding is zero and from 2.10

contribution is:

$$c_0 = \frac{K_0 + G}{(1-x)c} \tag{2.12}$$

Policy $(c,p) = (c_0,1)$ is ideal for cooperators, since contribution is minimal and free-riding is zero, hence minimizing the cost of enforcement and also maximising the utility from norm adherence.

The question that arises is, under what conditions is the policy $(c,p) = (c_0,1)$ a political equilibrium? The answer is there should exist no alternative balanced-budget equilibrium policy under which a majority would be free riders and obtain a higher expected utility than when contribution is c_0. To be more specific, let $\tilde{\eta}$ be the median value of the norm parameter and assume no other externalities from free-riding exist. Then

Proposition 3: $(c,p) = (c_0,1)$ is not a political equilibrium policy if and only if the following hold concurrently (2.13) and (2.14) hold for some policy (c,p) that satisfies (2.11)

$$(1-c)(1-p) > (1-c_0)^{1/\pi} e^{\tilde{\eta}u(0)} \tag{2.13}$$

$$c \geq 2c_0 + \frac{c}{G} \tag{2.14}$$

Proof. The policy $(c_0,1)$ cannot be beaten under majority vote by a policy that results in a majority of cooperators, since they should be better off with $(c,p) = (c_0,1)$. Hence, only a policy (c,p) satisfying (2.11) and such that (2.10) gives $x \geq 1/2$ can beat $(c_0,1)$. Note that $x \geq 1/2$ holds in a balanced budget equilibrium if and only if (2.14) holds. This inequality is obtained by solving for x in 2.10 and requiring a majority of free-riders. Hence a competing policy (c,p) results in $x \geq 1/2$ iff (2.14) holds, and such a policy beats $(c_0,1)$ if and only if the expected utility to a free-rider is higher than contributing under policy $(c,p) = (c_0,1)$. This boils down to

$$(1-\pi)lnc = \pi ln((1-p)(1-c)G > ln((1-c)G) + \eta\pi u(0)$$

or equivalently

$$\pi ln((1-p)(1-c)) > ln((1-c_0)G) + \eta\pi u(0)$$

This inequality needs to hold for all $\eta \leq \tilde{\eta}$ in order for x to be at least 1/2. This is condition (2.13) □

This proposition has some interesting comparative statics implications, summarized as follows. Let

$$P_0 = \{(c,p) \in [0,1]^2 : \text{eqs. 2.13,2.14hold}\}$$

We can say that $(c_0,1)$ is a more 'sustainable' political equilibrium, the smaller P_0 is. It follows from the above proposition that the policy $(c_0,1)$ is a more sustainable , the stronger the social norm is against free-riding, the higher the probability is for a free-rider to be caught and the higher the marginal cost (as a share of revenue) of law enforcement.

Corollary1: $(c_0,1)$ is a more sustainable political equilibrium, the higher $\tilde{\eta}u(0)$ is, and higher π is and the higher k/G is.

The question that naturally follows, is whether there exists any other political equilibrium. Since $(c_0,1)$ is the ideal policy for cooperators, political equilibria with a majority of free-riders is the only possible alternative. IN order to address this question, note that (2.11) defines the penalty rate p as a function of contribution c. Specifically (2.11) is satisfied if and only if $p = G(c)$, where

$$G(c) = \max\{0, 1-(1-c)^{\frac{1-\pi}{\pi}} exp\left[H^{-1}\left(\frac{c-c_0}{c+c/w}u\right)\left(\frac{c-c_0}{c+c/w}\right)\right] \tag{2.15}$$

Let $\hat{c}$ be the lowest level of contributions that is compatible with a (weak) majority of free riders. By (2.10), $x \geq 1/2$ requires $c \geq \hat{c}$ where

$$\hat{c} = 2c_0 + K/G \tag{2.16}$$

Hence a necessary condition for a political equilibrium with a majority of free-riders to exist is $\hat{c} < 1$. Let

$$c* = \min \arg \max_{c \in [\hat{c},1]} \{(1-c)^{\frac{1-\pi}{\pi}} exp\left[H^{-1}\left(\frac{c-c_0}{c+c/w}u\right)\left(\frac{c-c_0}{c+c/w}\right)\right] \tag{2.17}$$

c^* is the minimal level of contribution among those that maximise free riders expected utility in balance budget equilibrium where these constitute a weak majority. For continuous function H^{-1} and u, the maximand is continuous hence the set of maximizers in non-empty and compact (due to Weierstrass's' Maximum Theorem). Therefore, c^* is the well define, as long as $\hat{c} < 1$. Let $p* = G(c)$. This is the unique penalty rate that makes $(c*, p*)$ a balanced equilibrium policy. If there is a political equilibrium with a majority of free-riders, then the level of contribution in the equilibrium needs to be $c*$ and the penalty needs to $p*$ since otherwise $(c*, p*)$ would defeat that policy under majority rule. We can see that the opposite also holds:

Proposition 4: Assume that $\hat{c} < 1$ and B^{-1} is continuous. Then $(c*, p*)$ is a political equilibrium if and only if

$$(1-c^*)(1-p^*) \geq (1-c_0)^{1/\pi} e^{\overline{\eta} u(0)} \tag{2.18}$$

Proof. Suppose $\hat{c} < 1$, B^{-1} is continuous and 2.18 holds. Then $(c*, p*)$ results in a balanced equilibrium state, that commands a majority of free-riders, and from 2.18 their expected utility is at least as high as under the optimal policy for the cooperators.As a result, no policy (c, p) with $c < \hat{c}$ since $c*$ maximizes free-riders expected utility across all balanced budget equilibria that constitute a (weak) majority □

This more complicated in terms of comparative statics, than the previous proposition, since the candidate level of contribution c^* is more indirectly defined in terms of the primitives that c_0. There is however, now special case that makes the analysis more straightforward.

Assume that the majority attaches no weight to the norm, i.e. $\tilde{\eta} = 0$ while the minority has a positive weight $\eta_1 > 0$ on the norm. In this Case $x_1 = H(0) > 1/2$. The minority that still care about the norm would still contribute even if the majority won't, as long as

$$\eta_1 u(x_1) \geq (1 - \frac{1}{\pi} ln(1-c) \tag{2.19}$$

where c is the level of contribution at that state. Assume this level c_1 is such that the revenue from this 'honest' minority is enough to keep the balanced budges. In the absence of a penalty for free riders, the 'dishonest majority will free right and thus $(c_1, 0)$ is a political equilibrium. We can see that this equilibrium is easiest to obtain the larger the utility from adhering to the norm from the 'honest' minority, allowing us to conclude that the more they are vulnerable to be exploited by the free rider majority.

2.6 Conclusion

The main contribution of this paper is the development of a model that captures how formal and informal institutions (laws and social norms respectively) interact to enforce cooperation in a situation of social dilemma. Given how the traditional models that do not take non-economic incentives into account perform poorly with real world data particularly in situations where externalities are present, we provide a model that could potentially perform better in that respect.

The multiplicity of equilibrium free-riding that the model predicts can offer another view as to also why different groups of people or societies might exhibit different results for similar policy parameters. Finally, I show what kind of political equilibria just kind of societies can maintain.

Although it is a common criticism that any outcome is possible in a model where the preferences are cherry-picked towards that end, the assumptions that we relied on, namely, the presence of a social norm which whose intensity is endogenous is founded on extensive empirical evidence and well established theories from different disciplines of the social sciences, as it is discussed in sections I and II. As highlighted in Fudenberg and Levine (2016), the effect of social norms especially in the field of game theory is still not fully explored and provide a fertile area for exploration.

Possible extensions can go in several directions. Arguably the most apparent, is to introduce a temporal dimension to this game, and look into the outcomes when the game is played multiple times, and what happens when overlapping generations of agents interact and how cooperation evolves through time. This would also allow us to relax the balanced-budget requirement and study also how debt could come into play. What policy mix leads to the highest theoretical level of cooperation and at what rate.

Another extension of interest is to make a more realistic representation of the political process by including a setting of representative democracy which is more in line with how laws are determined in modern societies. Also dropping the simplifying assumption of a binary choice between cooperating and free-riding,would make the model even more realistic, since there exist strategic motives to partially comply especially when that leads to lower disutility from the social norm. Finally, the model could possibly be applied to study other situations of social dilemma of important policy significance such as the Prisoner's dilemma or the Tragedy of the Commons.

2.7 Appendix

Here are the numerical specifications for the figures used throughout this paper.

Figures 1 and 2

We hold parameter η fixed, and assume the degree of pro-sociality across the population is distributed with mean $\mathbb{E}(\theta) = 1$ and the norm function $n(x)$ is a logistic function. Figure 2.4 shows the graph of the following logistic function

$$n(x) = \frac{1}{1+exp(8x-4)}$$

Figure 2.4: Logistic Function

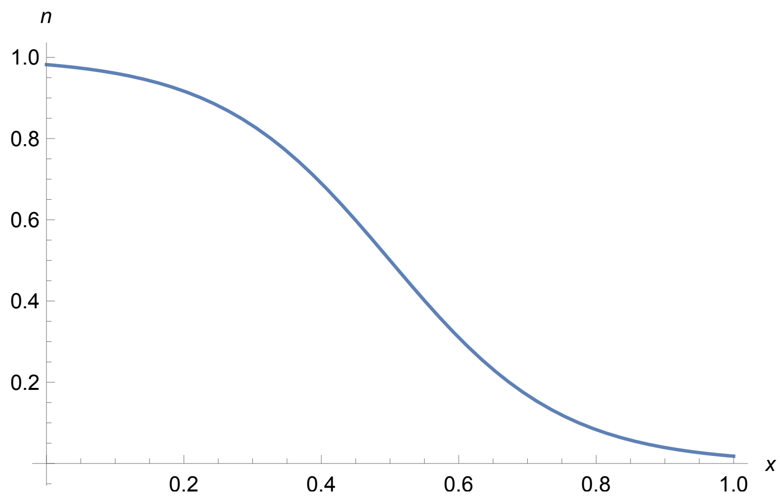

We use this utility and the exponentially distributed parameter and plot the fixed-point equation $x = F(x)$

$$x = 1 - exp(-\frac{\Delta u}{\mathbb{E}(\theta)}[1+exp(8x-4])$$

I defined Δu in equation (2.5) and Figure 2.1 is drawn for $\Delta u = 0.2$. Figure 2.2 uses the same equation defined above but allowing Δu to vary with the crime rate x.

3 Information Design and the Implementation of Efficient Social Learning Outcomes

3.1 Introduction

When faced with a decision that involves uncertainty, individuals turn to their environment to obtain additional information. The act of agents drawing inference from others either by communicating with them, observing the realization of their actions or those actions directly, is commonly referred to as *Social Learning*. Such behavior has always been prevalent in a large variety of economically relevant information, from buying a new product, adopting a new technology or even deciding which politician to vote for. In fact, owing to the advances in technology (particularly, the Internet and Social Media), people are exposed to an unprecedented amount of information about others, so the question arises as to whether this information does lead to better informed decisions.

The question of whether information across the members of an economy is efficiently aggregated has been central for as long as economics as a formal field of inquiry has existed, with Hayek (2009) raising the question more than 70 years ago[1]. The prevailing notion of the "wisdom of the crowd" a term first coined by Galton (1907)[2] posits that although individuals may make errors or be misinformed, pooling all the individually-held information together has to lead to a correct conclusion. Hence, the more people make a certain choice, the more confident we can be that this choice is correct, therefore it might be a good idea to imitate it. This is also the rationale that lead to the ubiquity of

[1]To directly quote: *"The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form., but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess* (Hayek, 1945, p. 519)

[2]Galton observed a contest calling on participants to guess the dressed weight of a live ox. Examining the nearly 800 entries, He wrote that "the middlemost estimate expresses the vox populi" and reported this value (the median) as 1207 pounds, within about 0.8% of the actual weight of 1198 pounds

online information aggregation platforms that collect the first-hand experience and evaluation of users in order to better inform others (e.g. TripAdvisor for travel destinations, Yelp for shops and restaurants etc.).

Although it would be very hard to argue against the positive impact that social learning has on improving economic performance and peoples' livelihood, the question as to whether social learning operates as efficiently as possible is one with a less clear answer. The main caveat of social learning is that, in order for information to be aggregated efficiently and provide useful information to the group, it has to be so that each agent's action reveals (at least, partially) some of their private information. The incentive structure however, is such that an agent endowed with private information that might contradict the prevailing public belief (which is the most useful type of information) would often prefer to imitate others (herd) rather than take the action their private information dictates. The rationale being, that is very unlikely that she is right when all others around her are wrong. This phenomenon in stead acts as a vicious circle that leads to the whole group's behavior being drive by very limited real data and by very few critical agents, with the risk of herding at a sub optimal decision. So a tension exists in terms of how social learning works, namely that when exposing people to information on others, the incentive to experiment goes away, cutting of the main source of information that is beneficial for all others in this group. This situation has all the characteristics of a typical market failure, therefore, the natural question that arises is whether there is room for some efficiency-improving intervention.

The aim of this paper is to study a simple game of observational social learning, i.e. a situation where only the actions of agents are observable to others[3], characterize the aforementioned tension between acting on what you know versus what others did before you[4] and measure the associated inefficiency compared to a theoretical benchmark where all private information is directly observable. Then, I propose a mechanism that mitigates this inefficiency, by assuming a social planner in the form of an *'Information Designer'* that a priori commits to mechanism that will strategically disclose the available information to each of the agent taking an action. I make the simplifying assumption that the designer can mediate between the available information and what agents will receive, in order to make exposition easier and focus completely on the question of how observability of other's might harm social learning in certain cases. However, although the designer holds a

[3]From all the variations of social learning scenarios, learning through other's actions is arguably the most prevalent way of social learning, since its a much more credible signal to others than what they may communication through words, or as the saying goes *'actions speak louder than words'*

[4]A more detailed discussion of the problem and the literature follows on Section 3.2

'monopoly' on public information, it will be made clear that he may not have the complete freedom as to how much of the information he can manipulate (the designer needs to maintain a reputation for being informative in order for his recommendations to be credible, especially in this case that is assumed to be benevolent).

The reasoning behind what makes the proposing mechanism optimal is that it addresses the unaccounted-for positive externality an agent revealing his signal can pose which he tends to disregard in a standard social learning game where the weight of public information overwhelms his own information. The designer strategically reduces the informativeness of the public signal, so agents are more likely to act upon their own information (which I assume its more likely to be right than wrong), so the designer in turn has access to better quality information that will eventually fully disclose, causing agents to herd to the right action based on the underlying state of nature. I focus on a beliefs-based approach rather than a more traditional mechanism-design-like incentive scheme to address the inefficiency, since the nature of the problem is one of an inefficient underlying information structure, but also because in terms of real-world application, providing explicit monetary incentives might make inference from observable actions even harder, or might outright be unfeasible and/or illegal.

The contribution of this paper is two-fold: First, it provides a simple model, that although fairly stylized[5], it manages to identify how individuals weigh their own information against that obtained from others, and captures the 'market failure', in terms of social welfare, since individual and social optima fail to align. Secondly, it provides a mechanism to address this inefficiency, that instead of trying to 'price' the externality, it takes a Bayesian persuasion approach of privately informed agents, which to my knowledge has not been attempted before.

The main insights drawn from this study are the following. Firstly, social welfare is not monotonic to the amount of available information, since past history might stifle any incentives the agent might have in following their own information when its too strongly opposed to public belief. Secondly, herding i.e. imitating what others are doing even when it does not agree with private information need not be always sub-optimal, even though herding can lead to inefficiencies as we will show. In fact the optimal mechanism is one that will eventually induce a herd, but at a later time than what happens in the standard

[5]The assumption of bounded private signals drives a lot of the result but is mainly used to simplify exposition. The existence of unbounded signals on the other hand would turn the problem from one of avoiding wrong cascades, to making convergence to the correct state arbitrarily slow which can be equivalently stated as welfare loss as well (Smith and Sørensen, 2000)

case. Finally, although counter-intuitive, garbling information can make the group better-off, as long as the amount of noise balances the benefits and costs of exploration.

The immediately apparent applications of those insights are the design of more efficient algorithmic procedures that online platforms implement (such as Google, TripAdvisor etc.) in order to optimally tackle the exploration/exploitation dilemma they face in their operation. The findings of this paper suggest a moderation in the initial available information about a product or a service to users maintaining the balance between being informative, and allowing users to explore and reveal their genuine information about it. Studying the mechanisms through which individual informational endowments are aggregated within a group has important economic implications and its relevance is becoming even greater as technology continuously increases the number of people an individual can interact with.

Beyond online platforms, the same principles apply to public policy as well. Think of a government that looks to implement a new environmental policy that promotes the adoption of a costly technology. Despite what the government might know from research and past experience, it relies heavily on the implicated firms, to go ahead and experiment with that new technology on the field, before the government can legislate mass adoption if the technology turns out to be beneficial, or withdraw the policy if the results are conclusively unfavorable. However, if initial results turn out to be overwhelmingly positive (negative), it might lead to others to rush to adopt the technology (or not willing to try it). How the government chooses to disclose initial information is hence key to learning about the unknown state effectively. Managing the beliefs of the interested parties is crucial and in fact it is common practice in the field of clinical trials in medicine, where the true type of the treatment a subject is given is withheld from him in order to identify placebo effects, hence assuring the clearest possible evaluation of the impact of a treatment.

The rest of the paper is organized as follows. Section 3.2 discussed the place of this paper in the literature. Section 3.3 presents the social learning model utilized. Section 3.4 presents the information design mechanism and its properties. Sections 3.5 and 3.6 characterize the optimal policy in 2 different contexts, being the main analysis, and Section 3.7 presents some concluding remarks.

3.2 Literature Review

This paper relates mainly to two strands of literature. First, that on social learning and in particular to observational learning with emphasis on the role of the underlying informational structure. The rigorous study of social learning, from a game-theoretic perspective started with Bikhchandani et al. (1992) and Banerjee (1992) who are the first to point out the potential inefficiency of herding, when a series of agents take actions myopically, ignoring their effects on the learning and welfare of agents in the future, and how information cascades can form in settings where private signals are bounded[6]. Smith and Sørensen (2000) extend the study into unbounded signals and establish the 'overturning principal' which states that an agent with a sufficiently strong private signal will eventually arrive and break an information cascade, therefore the group always learns the true state.

This however still leaves the question of efficiency open, since the speed of convergence can get arbitrarily slow with the welfare loss (Hann-Caruthers et al., 2018). Rosenberg and Vieille (2017) identify the two determinants of social learning efficiency as the distribution of private signal and the nature of the informational feedback available which the present paper contributes to.

Other important papers that have looked into the effect on efficiency of different information structures, are Smith and Sorensen (2011) and Çelen and Kariv (2004), the former looking into a mechanism disclosing a sample of past plays to the agent making the move and the latter a mechanism that only exposes the action of the previous player. Although the nature of the question is similar, the current paper a more systematic approach to finding the optimal information structure instead of testing specific schemes. What the current article adds to the literature, is the study of efficiency, and the proposal of efficiency-improving mechanism.

The second field that this paper contributes to is that of information design and strategic communication. The study of a sender-receiver relations in terms of influencing beliefs about an uncertain state has been the center of interest in microeconomics for as long as the field existed. Crawford and Sobel (1982) formalized the idea of cheap talk and Myerson (1983) formalizing the problem of "communication in games" where the designer cannot control outcomes but can elicit information from players and pass it to other players. In recent years the literature has received renewed attention, thanks to a large extend to a paper by Kamenica and Gentzkow (2011) that phrased the optimal design of information as a "Bayesian persuasion" problem between a sender and single receiver, closely related to previous work by Brocas and Carrillo (2007) and Rayo and Segal (2010). Kolotilin et al. (2017) looks into a situation that resembles the one of this paper, where a sender designs the optimal persuasion mechanism when facing a receiver that holds private information and has the ability to elicit that information via a revelation-principle-type

[6]An approach this paper takes as well

setup. I do not rely on the assumption of a two-way communication channel between parties, with the designer being the only one with the ability to send messages. Finally, for a comprehensive overview of the field of information design Bergemann and Morris (2019) offers an excellent survey of both the literature and applications of the theory.

The economic applications of information design have also been considered, from inducing higher voter turnout (Alonso and Câmara, 2016), to rating systems (Horner and Lambert, 2016), matching markets (Ostrovsky and Schwarz, 2009), regulatory transparency (Asquith et al., 2013) and in price discrimination at financial markets (Bergemann et al., 2015)

Finally, disclosure of information in the form of action recommendations to induce experimentation have been used by Kremer et al. (2013) who study the optimal mechanism that induces agents to explore two products of unknown qualities. As in this paper, the designer can incentivize agents to explore by manipulating their beliefs, and her ability to do so increases over time. While these themes are similar, there are differences especially the fact that the designer has a complete informational advantage and the agents do not hold any private information. So although the objectives of the designer are the same, and the method of influence is also a recommendation mechanism the policy method that information is being manipulated is substantively different. Another paper that studys social learning, and arguably the one closest to this present work, both in terms of methodology and of substantive question is that of Che and Hörner (2018). This article also studies the problem of inducing experimentation using a recommender system that can 'spam' agents by recommending to try a product of unknown quality. The key difference being once again the informational advantage the designer holds, where the leverage the designer has on influencing the agents is greater since the agents have no private information. Another difference, is the fact that information comes in the form of breakthrough news, with the actions of agents increasing the likelihood of news arriving, unlike the present paper that looks into efficiently aggregating privately-held information.

3.3 A Model of Observational Learning

Here I present the simplest possible setting of observational learning of two possible payoff-relevant states, two actions, and two possible signals the agent may receive, in order to fix ideas, and to demonstrate the concepts of *Herding* and *Information Cascade* that are a property of such games.

Agents are faced with the choice to buy a good or not the quality of which they do not know for certain. If they choose to buy it and the product turns out to be good they will enjoy a positive payoff. If the product turns out to be of bad quality they will suffer a negative payoff. Finally, they can choose not to buy it at all in which case they receive 0 payoff regardless of state. They have their private information about the quality but they also rely on the information they can obtain from others to guide their decision.

There are two possible payoff-relevant states of the world $\theta = \{0,1\}$ with common prior μ_1 that $\theta = 1$ and either finite N or a countably infinite population of agents, indexed by the period t each is faced with the decision $\alpha_t \in A = \{0,1\}$. The order of arrival is exogenous. Each agent is endowed with a private signal s_t of precision $q > 1/2$, with $q = P(s_t = \theta|\theta)$

The payoff of the agent is as follows:

$$u(\alpha,\theta) = \begin{cases} 0 & if \ \ \alpha = 0 \\ \theta - c & if \ \ \alpha = 1 \end{cases}$$

or equivalently, $u(\alpha,\theta) = (\theta - c)\alpha.$ where $0 < c < 1.$ The available public information at period t is $h_t = \{\alpha_1, \alpha_2, ...\alpha_{t-1}\}$, i.e. all the past plays of t's predecessors. Public belief at period t is denoted as $\mu_t = P(\theta = 1|h_t)$.

The information structure is that of binary signals about the possible state of the world. So $s_t \in (0,1)$. In that case 1 is considered a 'good' signal i.e. that $\theta = 1$ and 0 a 'bad' one. Hence, I call agent t whose $s_t = 1$ an "optimist" and those whose $s_t = 0$ a "pessimist". Agent combines s_t with μ_t to form his private belief denoted as μ^+ and μ^- for the optimist and the pessimist respectively.

At any period, the following relation holds for each type of belief $\mu^- < \mu < \mu^+$[7] A pessimist will only buy if $\mu^- > c$ since μ^- is a function of μ or equivalently if $\mu = \mu^{**} > c$ (particularly for $c = 1/2, \mu^{**} = q$). If a pessimist finds it optimal to buy then it must be that the optimist does so as well. Hence $\forall \mu_t > \mu^{**}$ all agents will buy. If $\mu_t \leq \mu^{**}$ only those with a good signal will buy.

Likewise, let μ^* be a value of the public belief such that $m^* = c$. If $\mu_t \leq \mu_*$ then t does not buy regardless of his signal. If $\mu_t > \mu^*$ he buys only if he is an optimist.

The decision rule of the agent can be summarized as follows: For all t and public belief μ_t.

- if $\mu^* \leq \mu_t \leq \mu^{**}$ agent t buys only if he is an optimist $(s_t = 1)$
- if $\mu_t > \mu^{**}$ agent t buys regardless of signal
- if $\mu_t < \mu^*$ agent t does not buy no matter his private information $(s_t = 0)$

[7] from Bayes' Rule $\mu^- = \frac{\mu(1-q)}{\mu(1-q)+\ (1-\mu)q} < \mu < \mu^+ = \frac{\mu q}{+(1-\mu)(1-q)}$

Some definition follows:

Definition 1: Agent t *herds*, on the public belief μ_t when his action is independent of his private signal.

Definition 2: If all agents herd, we have an *information cascade*

Definition 3: A herd takes place at $t = T$ if all actions after T are identical, i.e. $\forall t > T$ we have that $\alpha_t = \alpha_T$

A cascade occurs whenever $\mu_t > \mu^{**}$ or $\mu_t < \mu^*$ and social learning occurs only if $\mu^* \leq \mu_t \leq \mu^{**}$. When social learning occurs we have that $\alpha_t = s_t$ so the actions reveals perfectly the signal. The history of actions is in that case equivalent to the history of signals $(s_1, s_2, .., s_{t-1})$. It is easy to see in the case that two consecutive identical signals were to arrive, an information cascade would start, either of $\alpha = 1$ or $\alpha = 0$. In order not to have a cascade a necessary condition is that signals alternate consecutively between 0 and 1. The main takeaways from this stylized model is (i) that t the probability that a cascade has not started by t converges to 0 exponentially with time, (ii) there is a positive probability that we have a wrong cascade ($a_T = 0|\theta = 1$ or $a_T = 1|\theta = 0$) and (iii) beliefs do not change after a herd starts.

Social Learning with Bounded Beliefs

We can now generalize the model to admit any distribution of private beliefs $F^\theta(\mu)$ that depends on θ while still looking at the case of 2 possible states. We further assume the conditional density functions satisfy the proportional property, more formally:

$$\frac{f(\mu|\theta_1)}{f(\mu|\theta_0)} = \chi \frac{\mu}{1-\mu}$$

with χ a constant. This simply means that one of the two cdf's First-order Stochastically Dominates the other $(F(\theta_1) > F(\theta_0))$. We will represent with $F^\theta(\mu)$ the distribution of belief measured as the probability of θ_1 and $F^\theta(\lambda)$ the cdf of the distribution of the Log Likelihood Ratio between θ_1 and θ_0.

Learning within this setting takes place as follows. For every t, agents observe action α_t and update their belief via Bayes' Rule as follows:

$$\lambda_{t+1} = \lambda_t + v_t \text{ where } v_t = log\left(\frac{P(\alpha_t|\theta_1)}{P(\alpha_t|\theta_0)}\right)$$

Agent t buys if and only if the probability of $\theta = 1$ is greater than the cost, i.e. $\lambda > \gamma = log\frac{c}{1-c}$. The updating term is independent of λ_t, so the distribution of beliefs is transformed by a random term v_t from period t to $t+1$. The probability to buy depends on the state and is equal to $\pi_t(\theta) = 1 - F_t^{\theta}(\gamma)$.

The action α_t provides a binary random signal on θ with probability as follows:

	Actions	
State of Nature	$a_t = 1$	$a_t = 0$
$\theta = 1$	$1 - F_t^{\theta 1}(\gamma)$	$F_t^{\theta 1}(\gamma)$
$\theta = 0$	$1 - F_t^{\theta 0}(\gamma)$	$F_t^{\theta 0}(\gamma)$

Because $F^{\theta 1}$ stochastically dominates $F^{\theta 0}$ there are more optimistic agents on the good than in the bad state. Hence the probability of buying in higher in state $\theta = 1$ than $\theta = 0$, so (3.3) becomes:

$$\lambda_{t+1} = \lambda_t + v_t \; with \; v_t = \begin{cases} log\left(\frac{1-F_t^{\theta_1(\gamma)}}{1-F_t^{\theta_0(\gamma)}}\right) & if \quad \alpha_t = 1 \\ log\left(\frac{F_t^{\theta_1(\gamma)}}{F_t^{\theta_1(\gamma)}}\right) & if \quad \alpha_t = 0 \end{cases}$$

We can see that $v_t \geq 0$ if $a_t = 1$ and $v_t \leq 0$ if $a_t = 0$. So observing α_t conveys some information on θ as long as $F_t^{\theta_1}(\gamma) \neq F_t^{\theta_0}(\gamma)$. The distribution of LLR is invariate up to a transformation so it is sufficient to keep track of only one of the beliefs, therefore we can set public belief to be the midpoint of the support.

We can express the Social Learning process discussed above as a Markov process. We see that the position of the distribution of public beliefs can be characterized by one point λ_t. Let μ_t be the belief of an agents with $LLR = \lambda_t$. The expression above then takes the general form: $\mu_{t+1} = B(\mu, a_t)$ and a_t a random variable that takes values 0 or 1 according to the above table. The probability of a_t depends on μ_t and θ_t, hence the process of social learning can be summarized as:

$$\begin{aligned} &\mu_{t+1} = B(\mu, a_t) \\ &P(a_t = 1) = \pi(\mu_t, \theta) \end{aligned} \tag{3.2}$$

with B being the general form of the Bayesian formula presented above. The combination of the two equations defines a Markov Process for μ_t that can be used to further study social learning and particularly convergence.

Social Learning and Cascades with Bounded Beliefs

Assume the initial distribution of beliefs being bounded $\lambda \in (\underline{\lambda_1}, \overline{\lambda_1})$. Let λ_t be the public belief at t which is the midpoint of the support $\lambda_t = (\underline{\lambda_1} + \overline{\lambda_1})/2$ and let $\sigma = (\overline{\lambda_t} - \underline{\lambda_t})/2$ a constant.

If $\lambda_t > \lambda^{**} = \gamma + \sigma$ the support is above γ and we have an information cascade where t and everyone after him buy whatever their beliefs. Likewise if $\lambda \leq \lambda^* = \gamma - \sigma$ no agent buys.

Social learning only occurs as long as $(\lambda^*, \lambda^{**}) = (\gamma - \sigma, \gamma + \sigma)$, the complement of the cascade set. From the Martingale Convergence Theorem, we can obtain that the probability $\mu_t = \frac{et}{(1+et)}$ converges in probability. So λ^t must converge to some value.

Suppose the limit is not in the cascade set, then asymptotically the probability that $x_t = 1$ remains different in states θ_1 and θ_0. Therefore, with strictly positive probability the common belief is updated by some non-vanishing amount thus contradicting the convergence of the martingale. This argument can be used to prove that λ_t must converge to a value in the cascade set.

We arrive to the result that, for a support of private beliefs $I = [_1 - \sigma, _1 + \sigma]$, λ_t converges almost surely to a limit $\lambda_\infty \notin (\gamma - \sigma, \gamma + \lambda)$ with $\gamma = log\frac{c}{1-c}$.

Right and Wrong Cascades

The core issue this paper is concerned with is the fact that a cascade may arise where the wrong action is played by the agents. Agents learn rationally so the probability of a wrong cascade is small if agents have a wide diversity of beliefs as measured by the support of the distribution.

Suppose the initial distribution in LLR is symmetric around 0 with support of length 2σ. For an agent with initial belief $\frac{1}{2}$ the probability of a wrong cascade can be computed. A cascade where no one buys arises if $\lambda_t < \gamma - c$ i.e. $\mu_t \leq \beta = \frac{e^{\gamma-\sigma}}{1-e^{\gamma+\sigma}}$. We can then show that for $\mu_1 = 1/2$ we know that $P(\mu_t \leq \beta | \theta_1) \leq 2\beta$. Analogously for a cascade where everyone buys.

If the support of the initial distribution of LLR contains the interval $[-\sigma, +\sigma]$ then for an observer with initial belief of $1/2$ the probability of a wrong cascade is at most 4β where $\beta = \frac{\epsilon^{-\sigma}c}{1-c+\epsilon^{-\sigma}c}$

Cascades need not occur in cases of unbounded beliefs (non-atomistic distributions of prior beliefs). However, social learning can still be problematic due to slow convergence to the true state. In a way main question remains since we are interested in improving the efficiency of social learning, whether we have a quick convergence to the wrong state or a slow convergence to the right one still entail welfare loss.

3.4 The Information Design Mechanism

The Information Design approach can be seen as a parallel to Mechanism Design, following the same principles, but where rather than changing any of the parameters of the game (in our case the Social Learning game presented above), the interventions center on seeking to induce desirable posteriors for the agents. A detailed discussion of Information designer follows below.

The Designer does not possess any informational advantage over the agents, nor has direct access to their private information. Rather, I will be looking at the optimal feasible way that this private information agents hold can be elicited. A substantive difference of this paper to the standard approach in this case the designer cannot condition directly on the true state, but to the realization of a stochastic process. The intuition behind the solution is that by weakening the public information, via a mechanism that obfuscates information, agents are more likely to act upon their own information and hence feed back to the designer more accurate public information that can be used to inform subsequent agents.

Since the designer here acts merely as an information mediator, this mechanism can be seen as a case of information design without elicitation. Namely, the designer cannot condition what he reveals on the reported type, hence he will have to a contingent policy, that is a vector of disclosure policies, where each individual entry is an action recommendation for a specific type weighted by the likelihood of that type of agent realizing. Additionally, the constraints the designer is faced with regarding the mechanism he can implement, is that of obedience, i.e. that the mechanism needs to be such that the agent will respond to the information he will receive in the desirable way, and Bayes plausibility, i.e. that the expected value of the induced posterior has to equal to the prior.

I will now outline the model that captures the interaction between the designer and the agents, the information structure and the timing, A product is released at time $t = 0$ and for each time $t \in [0, \infty)$ an agent arrives and decides whether to consume the product or not. Agents are myopic[8]and homogeneous in preferences (varying only in terms of their private information). The setup of the game at each t in terms of payoff and actions are as discussed in Section 3.3, with the Designer repeating the same game with each arriving agent[9] (one at each t). The quality of a product is the underlying state of nature that is realized at $t = 0$, but is unknown to both the designer and the agents. The a priori probability of the product being good, i.e. the prior is μ_0 and is common knowledge. The prior can take any value but for simplicity of exposition I

[8]Meaning they have no considerations past their immediate payoff

[9]I will discuss the case for a continuum of agents on a latter stage

shall be focusing on the case of $m_0 = \frac{1}{2}$ meaning that at the starting point, all parties are completely have no information as to which state is more likely.

Unlike the standard observational learning game discussed above, agents do not observe previous agent's decisions or payoffs. Instead the designer mediates social learning by collecting information from past agents and disclosing all or part of that information to the agents.

The designer's information and beliefs: The designer receives information about the product by observing the actions of the agents and using them to infer the likelihood of the true state being either 1 or 0 (i.e. good or bad) representing the quality of the product. The designer has no informational advantage besides the observation of past actions which can be summarized by μ_t which is the likelihood of the true state given the actions observed. The designer's beliefs that coincide with the public belief in the standard game have the following the law of motion:

$$\mu_{t+1} = \mu_t + v_t \text{ where } v_t = log\left(\frac{P(\alpha_t|\theta_1)}{P(\alpha_t|\theta_0)}\right) \tag{3.3}$$

Designer's disclosure policy: Based on the information received, the designer provides feedback to the agents. Since agents' decisions are binary, without loss of generality, the designer simply decides whether to recommend the product or not. The designer commits to the following policy: At time t, the designer recommends the product or not, based on his available information. However, the designer has the discretion to not perfectly reveal the information he possesses. Mathematically, this is done by recommending the opposite action than the one his information deem to be the optimal one. The strategy of the designer is a pair (k_t, l_t) where l_t represents the probability that the opposite action is recommended, for example, if the optimal recommendation would be $a = 1$ the designer would recommend $a = 0$ with probability k_t. This technology allows the designer to garble information and hence weaken t he public signal so that to maximally induce the agent to act upon her own signal. The designer maximizes the inter-temporal net surplus of the agents, discounted at rate $\delta > 0$, over the measurable functions (k_g, l_b) what determine strategy at each t, where $k := \{k_{gt}\}_{t\geq 0}$, $\{l_{gt}\}_{t\geq 0}$

Agent's Beliefs and Incentives: In this model, agents do not directly observe the designer's information or beliefs. However they can form a rational belief about the designers belief. They know at which point in the game their turn is so it is becoming increasingly less likely that a recommendation contrary to their belief will come. In addition, for the disclosure policy to be implementable the agents must have an incentive to follow the recommendation.

Since the exact circumstances surrounding the recommendations (whether the recommendation because of good news or despite no news) are kept hidden from the agents, their incentives for following the recommendation depend on their posterior regarding the designer's information.

The probability $_t$ is pinned down by the martingale property, i.e. that the designer's posterior must on average, equal the prior:

$$g_t * 1 + (1 - g_t)\mu_t = \mu^0 \tag{3.4}$$

After forming their posterior regarding the public signal $\tilde{\mu}$ she then combines it with their private information s_t to form the belief about the state μ^s. The agent will consume the product if and only if, their posterior private belief is such that:

$$\mu^s > c \tag{3.5}$$

The designer objective: The designer chooses a (measurable) policy $(k,l) := \{k_{gt}, l_{gt}\}_{t\geq 0}$

$$W(k,l) := \int_{t\geq 0} e^{-\delta t} g_t \text{ł}_t (1-c) d_t + \int_{t\geq 0} e^{-\delta t}(1-g_t)k_t(\mu_t - c)dt$$

Where (μ_t, g_t) must follow the required laws of motions described above[10]. Welfare consists of the discounted value of consumption $1-c$ in he event of good news and $m_t - c$ in the event of no news - for those recommended to consume the product.

To facilitate the characterization of the optimal policy, it is useful to consider the following benchmarks:

- **No Social Learning:** In which case, each agent follows their signal, and do not benefit from any social learning. Their expected payoff is simply the precision of their signal q.
- **Full Transparency:** The designer fully discloses his information, which corresponds to a strategy of $(k = 0, l = 0)\forall t$.
- **First-best Policy:** The Designer optimizes her policy (k,l) to maximize W, and ignoring any incentive compatibility constraints.

[10] It is feasible to randomize over the policy parameters, but it can be shown that this is never optimal

- **Second-best Policy:** Here we study how the incentive compatibility constraints limit the freedom the designer has to gable, and compare the welfare properties vs. the first-best policy

The Stage Game

Now I describe how the disclosure mechanism affects belief formation and learning dynamics in each specific period t. The designer knows public belief, derived from t he observation of past actions taken. Instead of fully disclosing the μ_t, designer commits to a policy that discloses $\tilde{\mu}$ instead, which is generated by:

$$\pi(|\mu) = \begin{cases} \overline{\mu_t} & \text{with probability} \quad \pi \\ \underline{\mu_t} & \text{with probability} \quad 1-\pi \end{cases}$$

such that, $\overline{\mu_t}\pi_t + \underline{\mu_t}(1-\pi_t) = \mu_t$

So, agent observes (μ_t, s_t) and knowing also the strategy of the designer (as shown above) forms his posterior and takes optimal action based on that. The above setting generates an ex-ante distribution of outcomes that differ from t he Social Learning Game, without the Designer. Under the presence of garbling, then at each stage the game falls under 3 regimes depending on μ_t
Regime A: μ_A^*, μ_A^{**} is the interval where everyone reveals their signal, with its boundaries defined as follows

$$\mu_A^* = \gamma - log\left(\frac{\underline{\mu}}{1-\underline{\mu}}\right)$$

$$\mu_A^{**} = \gamma + log\left(\frac{\underline{\mu}}{1-\underline{\mu}}\right)$$

where $\gamma = \log\left(\frac{c}{1-c}\right)$

Whenever $\mu_t \in (\mu_A^*, \mu_A^{**})$ everyone reveals their signal. Observing action a_t is equivalent to observing that action in the case of full disclosure, since signal averages to μ_t. So updating for the designer goes as follows:

$$\mu_{t+1} = \begin{cases} \mu_t + k & \text{if} \quad a_t = 1 \\ \mu_t - k & \text{if} \quad a_t = 0 \end{cases}$$

with

$$k = log\left(\frac{(1-\pi)\underline{\mu} + \pi\overline{\mu}}{(1-\pi)(1-\underline{\mu}) + \pi(1-\underline{\mu})}\right)$$

Regime B: In the second regime, μ_t is bounded by μ_B^* and μ_B^{**} such that $\mu_B^* < \mu_A^* < \gamma < \mu_A^{**} < \mu_B^{**}$. In this regime only those that receive a recommendation that contradicts their private signal do not herd. This regime is bounded by:

$$\mu_B^* = \gamma - log\left(\frac{\overline{\mu}}{1-\overline{\mu}}\right)$$

$$\mu_B^{**} = \gamma + log\left(\frac{\overline{\mu}}{1-\overline{\mu}}\right)$$

Regime B is further divided into 2 sub-cases:

The first the one where $\mu_t \in (^{**}_{A}, ^{**}_{B})$. In t his case, agents whose signal is congruent with recommendation herd. Agents with private signal that is not congruent to the recommendation reveal their signal. Public belief updats as follows

$$\mu_{t+1} = \begin{cases} \mu_t + l & \text{if} \qquad a_t = 1 \\ \mu_t - log\left(\frac{\overline{\mu}}{1-\overline{\mu}}\right) & \text{if} \qquad a_t = 0 \end{cases}$$

with

$$l = log\left(\frac{(1-\pi)\overline{\mu} + \pi\underline{\mu}}{(1-\pi)(1-\overline{\mu}) + \pi(1-\underline{\mu})}\right)$$

The log-likelihood changes more when $a = 0$ is observes

The second case is that where $\mu_t \in (\mu_B^*, \mu_A^{**})$ and

$$\mu_{t+1} = \begin{cases} \mu_t + log\left(\frac{\overline{\mu}}{1-\overline{\mu}}\right) & \text{if} \qquad a_t = 1 \\ \mu_t - l & \text{if} \qquad a_t = 0 \end{cases}$$

In this case, public belief changes more when a action against the herd is taken.

Regime C: In the third regime where μ_t lies on μ_B^*, μ_B^{**} and everyone herds on the public information

As is shown, the mechanism garbles information in such a way that contrarian information have a much greater impact on beliefs and conversely downplays the impact of repeated actions. This decreases the likelihood of an early information cascade, but as the Martingale Convergence Theorem states, herding with eventually occur, but with greater precision than the standard case of social learning.

3.5 Optimal Mechanism: Private Disclosure

To understand how this policy affects the decision of the agents, we start by describing explicitly how garbling affects the receivers decision to take action or no. If the designer decides to garble, he changes the belief associated with each signal. By choosing a garbling strategy (k,l) the agent forms belief $\tilde{\mu} \neq \mu$ when receiving signal μ. This relationship between $\tilde{\mu}$ and μ I call *belief transformation*, is stated below and holds for any pair of values (k,l)

Lemma 1 (Belief Transformation): Under disclosure policy (k,l) signal μ induces belief $\tilde{\mu}$ where

$$\mu = \mu_0 \frac{(1-\mu_0)\tilde{\mu} - \mu_0(1-\tilde{\mu})l - (1-\mu_0)\tilde{\mu}k}{\mu_0(1-\mu_0) - \mu_0(1-\tilde{\mu}l - (1-\mu_0)\tilde{\mu}k} \tag{3.6}$$

The function has a fixed point μ_0. It is increasing in $\tilde{\mu}$ if $k+l<1$ decreasing if $k+l>1$ and constant to μ_0 otherwise. The range of beliefs $\tilde{\mu}$ is the interval where $\underline{\mu} = \frac{\mu_0 l}{\mu_0 l+(1-m_0)(1-k)}$ and $\overline{\mu} = \frac{\mu_0(1-l)}{\mu_0(1-l)+(1-m_0)l}$

Depending on whether $k+l \leq 1$, then the agent associates higher signals μ with higher beliefs $\tilde{\mu}$, but this is reversed when $k+l>1$. The belief transformation is illustrated in Figure 1. Whenever $k>0$ and $l>0$ the agent can never be sure that each recommendation is truthful. The receive buys when her belief exceeds $\hat{\mu}$. That is when her signal μ exceeds threshold $\hat{\mu}(k,l)$ obtained from the belief transformation. The following proposition characterizes the optimal cutoff above which agent takes action under garbling.

Proposition 1: There exists a threshold

$$\hat{\mu}(k,l) = \frac{(1-\mu_0)\hat{\mu} - \mu_0(1-\hat{\mu})l - (1-\mu_0)\hat{\mu}k}{\mu_0(1-\mu_0) - \mu_0(1-\hat{\mu}l - (1-\mu_0)\hat{\mu}k} \tag{3.7}$$

such that

1. If $k < \frac{0(1-\hat{\mu})}{\hat{\mu}(1-\mu_0}(1-l)$ then (k,l) is increasing in both arguments, and the agent buys for any signal $\mu \geq \hat{\mu}(k,l)$
2. If $k > 1 - \frac{0(1-\hat{\mu})}{\hat{\mu}(1-\mu_0}l$ then (k,l) is decreasing in both arguments, and the agent buys for any signal $\mu \leq \hat{\mu}(k,l)$
3. Otherwise, agent never buys

We have established how the designer can affect agents' posterior belief. Moreover, the advanced of this approach is that it does not require an elicitation mechanism that requires the agent reporting their information. It is enough for the designer to induce an environment where experimentation is more likely, since herding will eventually occur by the Martingale Convergence Theorem.

All the above, is feasible under the first-best policy. Additionally in the second-best case, the strategy of the designer and hence, the freedom to implement different posteriors is limited by the belief the agents hold at each stage. This limit is determined by the martingale property.

$$\overline{\mu_t}\pi_t + \underline{\mu_t}(1-\pi_t) = \mu_t$$

Now that the optimal strategy for any given stage is determined, the final step is to characterize the optimal inter-temporal policy for the designer. Both the first-best and second-best policies have a cutoff structure. They induce maximal feasible exploration, which equals 1 under the first-best policy and the maximal feasible level under the second-best policy. as long as the designer's posterior remains above the threshold levels μ_B^* andμ_B^{**}. Otherwise, no exploration is any longer optimal, garbling stops and the full information is disclosed inducing a herd.

Proposition 2: (i) The first-best policy describes full garbling $(k = l = 1)$and therefore full exploration as long as μ_B^* and μ_B^{**}. As soon as μ_t enters the complement of the above set, then the optimal policy is to fully reveal the true state, hence $k = l = 0$

(ii) The second-best prescribes exploration at

$$k^{SB}(\mu_t) = \begin{cases} \hat{k}(\mu_t) & \text{if} \quad \mu_t \in (\mu_B^*, \mu_B^{**}) \\ 0 & \text{if} \quad \mu_t \in (\mu_B^*, \mu_B^{**})^c \end{cases}$$

and

$$l^{SB}(\mu_t) = \begin{cases} \hat{l}(\mu_t) & \text{if} \quad \mu_t \in (\mu_B^*, \mu_B^{**}) \\ 0 & \text{if} \quad \mu_t \in (\mu_B^*, \mu_B^{**})^c \end{cases}$$

where

$$k_t \leq \hat{k}(\mu_t) := \min\{1, \frac{(1-c)(\mu^0 - \mu_t)}{(1-\mu_t)(c-\mu_t)}\} \tag{3.8}$$

and

$$l_t \leq \hat{k}(\mu_t) := \min\{1, \frac{(1-c)(\mu_t - \mu^0)}{(1-\mu_t)(\mu_t - c)}\} \tag{3.9}$$

The optimality of the cutoff policy and the associated cutoff can be explained by the main trade-off that the designer faces for any given belief μ, which the potential gain in precision from an additional agent contradicting the current recommendation and revealing his information, vs. the potential cost of taking the wrong action.

The optimal polices have a cutoff structure. They induce mainly feasible exploration, as long as the designer's posterior remains about the threshold level μ^*. otherwise no exploration is chosen. The optimal policies produce very interest learning trajectories depicted in Figure 2

3.6 Optimal Mechanism: Public Recommendation

Thus far, I assumed that the Designer communicates to an individual agent so agents can be kept in the dark on the recommendations that others may have received. This is a reasonable assumption to make when thinking of an internet platform for example (Netflix, or Spotify sends personalized recommendations for example). However this is not a reasonable assumption to make in a situation of the government agency I discussed in the introduction, since any information needs to be make towards the entire public, and is not possible to tailor it to each one individually.It is clear to see that such a mechanism is not as effect as private communication when incentives agents to experiment.

I now turn to the same model but with a continuum of agents at each stage with unit mass. The designer cannot communicate to each agent individually[11]. Public recommendations are clearly not as effective as private recommendations in terms of incentives user exploration. Indeed, the optimal private recommendation identified earlier is not incentive compatible when made public. If only some fraction of agents are recommended to explore the action reveal the designer's private information and agents will immediately recognise that the recommendation is false and thus ignore it. Hence if the designer wishes to trigger user exploration she must adopt a different approach. I show that although garbling becomes less effective as an incentive when the recommendation is public, it is still part of the optimal policy. However for the policy to effectively induce experimentation, there must exist a component of randomness in it.

To see why the policy must be random, suppose that the designer commits to garbling -i.e. to recommend the product to agents despite the available information suggesting otherwise - at some determined time t for the first time. Since the recommendation is public, all agents observe it. Since the probability that public belief is still in the experimentation interval is negligible the agents will put the entire probability weight on the recommendation being false and

[11] or alternatively, she has no way to ensure that each individual recommendation could remain truly private

ignore it. Hence deterministic garbling will not work. Consider the random policy described by $F(t)$, the probability that the designer starts to garble by time t. Lets, focus, on the case where $\mu^* < \mu_t < c$.

As before, once the designer is confident enough about the state, he will always recommend the corresponding action to all subsequent agents. To see why the disclosure policy must be random, suppose the designer commits to garbling, i.e., recommendation the opposite of the optimal action the designer's current belief suggest, at some deterministic time t for the first time. Since recommendation is public, it is observed by all agents. Since the probability that the designer is not confident about the state becomes negligible, the agents will put the entire probability weight to the recommendation being misleading and will ignore it, hence a deterministic garbling policy will not work. $F(t)$ describes the probability that the designer starts to garble on time t.

First, if at any point in time the designer's belief falls below μ^* the designer stops experimenting. This follows from the optimal trade-off between exploitation and exploration identified earlier under the optimal (private) recommendation policy. Let t^* bet the time at which the designer's posterior reaches the threshold belief μ^*. Clearly, if the designer does not trigger garbling by time t^*, she will never trigger garbling after that time. This implied that the distribution F is supported on $[0, t^*]$. Additionally, once the optimal policy sends completely uninformative signal to all agents at some random time $t < t^*$, continuing garbling from then on does not change the agents' beliefs, thought the agents have no ground to update their beliefs. Hence, once they have incentives to explore, all subsequent agents will have the same incentive. Consequently, the optimal policy will continue to recommend the action until the designer's belief falls to μ^*.

Given these characteristics, the distribution F must be chosen to incentivize users to follow their own signal. To see how, the agents' belief are:

$$g_t = \frac{\mu_0 e^{-\lambda\rho t}(\rho + h(t))}{\mu_0 e^{-\lambda\rho t}(\lambda\rho + h(t) + (1-\mu^0)h(t)}$$

where $h(t) := f(t)/(1-F(t))$ is the hazard rate of starting to garble. This formula is a consequence of Bayes' Rule. The denominator accounts for the probability that the recommendation is made for the first time at t, which arises from either the designer receiving news at time t. The numerator accounts for the probability that the recommendation is made for the first time at t and state is good. For the agents to have incentives to explore, the posterior g_t must be no less than c, a condition which yields an upper bound on the hazard rate:

$$h(t) \leq \frac{\lambda\rho\mu^0(1-c)}{(1-\mu^*)(c-(1-c)e^{\lambda\rho t}) - \mu^0(1-c)}$$

This implies that the distribution F must be atomless. It follows that the incentive constraint is binding at the optimal policy -i.e. $g_t = c$- which gives rise t o a differential equation for F, alongside the boundary condition $F(0) = 0$. As mentioned, g_t remains at c from then on until experimentation stops. The unique solution is:

$$F(t) = \frac{\mu^0(1-c)(1-e^{-\lambda\rho t})}{(1-\mu^0)c - \mu^0(1-c)e^{-\lambda\rho t}} \tag{3.10}$$

for all $t < t^*$. Since the designer never garbles after t^*, $F(t) = F(t^*)$ for $t > t^*$

Examining F reveals various features of the optimal policy. First, the exploration as measured by F(t) under the second-best policy is single-peaked, just as with private recommendation, but in a probabilistic sense: The (expected) exploration, or the spam campaign, starts 'small' initially but accelerates over time as the designer builds credibility (i.e. F(t) is strictly increasing in t) and it stops altogether when μ^* is reached.

While garbling is part of the optimal public recommendation, its randomness makes it less effect for the designer in converting a given public belief into incentives for exploration, leading to a reduced level of exploration. The can be seen by the fact that:

$$F(t) = \frac{(1-c)\mu^0 - (1-c)\mu^0 e^{-\lambda\rho t}}{(1-\mu^0)c - (1-c)\mu^0 e^{-\lambda\rho t}} < \frac{(1-c)\mu^0 - (1-c)\mu_t}{(1-\mu^0)c - (1-c)\mu_t} < \frac{(1-c)(\mu^0 - \mu_t)}{(1-\mu^0)(c-\mu_t)} = \hat{\alpha}_t \tag{3.11}$$

Where both inequalities use $\mu^0 < c$, and the first follows from:

$$\mu_t = \mu^0 e^{-\int_0^t \lambda(\rho+\alpha_t)dt < \mu^0 e^{-\lambda\rho t}}$$

Hence the speed of learning is slower on average under public recommendation than under private recommendation. This is formally stated in the following proposition.

Proposition 3: Under the optimal public recommendation policy, the designer recommends action $a = 1$ at time time if she is confident that $\theta = 1$. If the designer is not received and a recommendation is not made by time $t \leq t^*$, she triggers garbling according to F(t) (3.11) and the garbling lasts until her belief reaches μ^*. The induced experimentation under optimal public recommendation is on average slower and the welfare attained is strictly lower than under optimal recommendation.

3.7 Conclusion

In this paper, I analyzed a simple social learning game where agents combine their own information with that inferred from the actions of their predecessors. A key inefficiency arises, namely that past history quickly overwhelms any privately held information leading to agents herding on an action early instead of acting on what they know, hence revealing their information to others. A mechanism of strategic information disclosure is proposed and characterized, that leads more efficient aggregation of a groups' privately held information. A strength of a beliefs based mechanism is the fact that it does not rely on traditional incentives schemes usually employed in mechanism design approaches, since in the majority of the applications discussed they would either be unfeasible, undesirable, or outright unethical and illegal.[12]

Despite the fact that the designer is benevolent, he still has to misinform a number of agents in order to induce them to act upon their own information and reveal it through their actions by adding noise to the public signal, so he can learn the true state as quickly as possible. In fact, the amount of noise is shown to follow a hump-shaped trajectory, since initially herding is less likely and the designer has limited access to information hence limited capacity to obfuscate, but as the likelihood for a sub-optimal herd increases in time the designer increases the noise and eventually fully reveals his information and induces the herd once confidence about the underlying state has been established.

While this paper provides some insights on how to improve the efficiency of social learning by implementing alternative information structures that strike the optimum point between exploration and exploitation, it leads to a number of very interesting research questions worth investigating. Arguably the most relevant question is how an optimal mechanism would be in the case where the underlying state of the world changes with time. [13] The dynamic properties of an optimal information mechanism will be interesting and relevant to investigate. Additionally, it is interesting to explore how a group learns in the presence of heterogeneously of preferences or private information quality. Mathematically, this would correspond to a model of

[12] Consider the extra layer of complication, trying to infer private information from incentivized actions, for example users getting payed to try a new restaurant.

[13] For example, it is common for technology companies to release updated versions of their products every year in order to fit the changing circumstances brought about technological improvements, customer preferences and competition. We often see in the real world that past reputation of the company and product matters a lot, meaning that the risk that a bad product might still end up being popular because it's based on a successful predecessor or a reputable company, or a new product innovative product might end up failing due to lack of customers willing to experiment, with both results being sub-optimal in terms of social welfare

misspecification, where the assumption of common knowledge in terms of preferences ceases to hold and agents (and the designer) need to take that into account when interpreting signals[14] Both these extension are left for a follow-up study.

[14] Think of the problem of trying to figure out if you are personally going to like a film, based on the choices of others, who, to some degree at least, might have different preferences than yourself.

4 Is More Information Always Better? An Experimental Analysis of Social Learning

4.1 Introduction

In situations of incomplete information, agents turn to their surroundings for information that could supplement their own privately held one to make a more informed decision. The underlying principle leading to this behavior is often referred to as the 'Wisdom of the Crowd' (Galton, 1907) meaning that an action that is repeated by many agents, each endowed with their own private information, is a very strong signal as to the underlying state of the world [1]. Such behavior pervades a very big portion of everyone's daily decision-making, and it has become even more present in the current era of increased interconnection and exchange of information that new technology facilitates. For example, one might rely on the best sellers' chart when deciding which book to buy, or to the ratings of TripAdvisor (that are derived from user-reported information) to choose her next holiday destination where popularity and number of reviews is considered a strong sign of reliability of the information.

This type of group behavior is usually referred to by the literature as Social Learning which is concerned with how individuals process public information and combine them with their own information to make decisions. Hence, Social Learning has an impact both at the individual level of decision-making as well at a group-level with important welfare properties depending on how efficiently (if at all) they aggregate dispersed and privately held information (Weizsäcker, 2010). The field of social learning is relatively recent, at least as far as economists have been concerned with it, with the first theoretical studies on by with Banerjee (1992) and Bikhchandani et al. (1992) and the first experimental one by Anderson and Holt (1997), it has however, received an extensive attention regarding potential failures and efficiency, especially with the increased relevance that social networks and the Internet has on our current society. The main caveat associated with Social Learning, is that each agent is

[1] The main source of uncertainty in such settings

both a consumer and a producer of public information, making the impact of their choices on themselves and to others, less clear ex-ante. The problem lies in the fact that when observing the choice of an individual, especially when it comes from a small set of possible actions (e.g. buy or not buy a product), it is hard to infer whether the agent has clear private information about the state, or is simply imitating others. This can lead to imperfect learning, where potentially erroneous decisions are propagated throughout the group, and are being reinforced, making harder for anyone with contrarian information to act upon it and make this information knows to her peers. This is a classical problem of externalities, since agents do not internalize the social benefit of them revealing their action has on others, even when it contradicts public knowledge.

In such settings, it is obvious that the structure of information, particularly the mechanism that the public information feeds back to each agent can potentially affect the efficiency of Social Learning. The role of information on influencing behavior over pay-off relevant incentives has received a large amount of attention by the literature recently, especially after the work on Bayesian Persuasion by Kamenica and Gentzkow (2011) [2]. One of the most robust and policy-relevant findings from the theory is that efficiency and social welfare need not be strictly increasing by increasing the amount of available information, in many different applications including social learning, identifying obfuscation from the part of the designer as an optimal strategy over full disclosure. Despite the volume of theoretical work, up to the time of the writing of this article, little empirical evidence exists to support this theory, especially in the setting of Social Learning.

The intended contribution of this paper is two-fold. Primarily, to test the hypothesis that the relationship between more information about the actions of others and how accurate the subjects choices are need not be monotonic, especially when the information they receive is an imperfect signal of the underlying true state, since information about others acts in two ways upon the action of an agent. Additionally, this experimental setting allows for a very precise measure of the decision process of Social Learning, both at the individual level, and at the group one. Specifically, I go about answering the following questions: i) How likely in each setup is to make the right choice? ii) Are subjects more or less likely to rely on their own information instead of that from their environment when introducing noise to the private signal? iii) Whether a cascade can be broken once formed (overturning principle as defined by Smith and Sørensen (2000) or is it stable (Bikhchandani et al., 1992) iv) Are subjects rational Bayesian in how they make their choices? Or

[2] See Bergemann and Morris (2019)) for a comprehensive survey of the field of information design

some systematic bias is involved.

The way I go about answering those questions is by implementing an experiment of sequential observational learning, where subjects take turns to make a choice regarding an imperfectly known state of their world, based on their privately endowed signal about that state, and an imperfect signal about the choices of the preceding subjects. I rely on a setting adjusted from Filippis et al. (2017), that presents the information to the agents in two steps and asks them to make their guess about the state based on this information. In the first step, only their private information is revealed to them, and in the second, also the public information about others is shown. This makes possible to disentangle the effect of each piece of information and estimate the weight they place on their private and public information in their decision making, as well as measure how they update this information. In my setting however, unlike the aforementioned experimental study, agents learn only the final choice of others and not their exact belief about the underlying state, which is a setting that much better mimics real world situation, and the inability to perfectly convey what individuals know trough their actions. In the treatment group, this information about others is noisy, hence less informative, and that allows for the testing of the hypothesis that less information about others might lead to a better-informed group overall.

I find that reducing informativeness of the publicly available information, by introducing noise to it, can in fact lead to more efficient leaning , especially when agents are aware of the inherent noise, reinforcing the idea of the effect of commitment when attempting to persuade an agent [3]. I find some sizeable departures from the way subjects take decisions compared to Bayesian models, however, in general subjects seem to react in a consistent manner to the arrival of signals, and how they update their information. I find also that the risk preferences of individuals seem to have a role to play with the Weight they attach to their own information versus the ones they receive about others. At the group level, I find that cascades become less stable, and although not conclusive, I find support for more efficient decisions under the treatment.

The remainder of this paper is organized as follows: Section 2 reviews the related literature and how this paper relates to it, Section 3 presents the formal hypothesis to be tested, Section 4 presents the experimental setting. Section 5 presents the results derived from the analysis of the experimental Data and the paper concludes with concluding remarks and future extensions in Section 6.

[3] As opposed to cases of cheap talk such as Crawford and Sobel (1982)

4.2 Literature Review

This article is an experimental study on the efficiency of social learning. The formal study of the phenomenon in the economics literature started with Banerjee (1992) and Bikhchandani et al. (1992) who were the first to model rational Bayesian agents, sequentially learning from the actions of others (referred to as Observational Learning) and identified pervasive phenomena that may lead to inefficient learning, namely herding and informational cascades. Herding refers to agents who, after taking their private information into account, rationally choosing to imitate their predecessors, and informational cascade (IC) doing so by completely ignoring their privately held information. The issue here is that although both phenomena are observationally equivalent, once an IC starts, no new information is being generated to make the public belief better informed and based on what the associated action to that IC is, it might even be the wrong one. Additionally, they point out that in such an environment the cascades are impossible to overturn. Smith and Sørensen (2000) and Guarino and Jehiel (2013) further study the problem, and show that the richness of the action space affects learning, and cascades may be cascades can be broken (the name this phenomenon the overturning principal, and a group may not converge to the true state, however this convergence might take an arbitrarily long time to occur, hence impactful in terms welfare nonetheless, in a dynamic setting. This study also looks into the stability of the herding behavior, offering additional empirical evidence as to the stability of herds and cascades. Finally, Rosenberg and Vieille (2017) look into the impact of different informational structures, both in terms of the degree of the observability of public information and different distributions of private information and attempt an ordering of different settings in terms of efficiency. The essay in Chapter 3 of this thesis proposes a mechanism based on information design that can lead to such an improvement by introducing noise to the public signal. One of the main contributions of the present study is to test this proposition in the lab.

In terms of laboratory studies, social learning has been studied extensively as well. The first experiment, which has since become canonical, Anderson and Holt (1997), where subjects take turns to make a guess about an unknown state, based on their private signals drawn from urns with known probabilities, and from the choices made by others. In their experiment, informational cascades happen consistently, though are not always stable as the theory up to that time suggested, a result which they attribute to a number of biases and errors, such as employing heuristics. The main limitation of this pioneering paper is the use of empirical methods for eliciting the beliefs the subjects form during the duration of the experiment instead of a more direct form of elicitation. Goeree et al. (2007) revisit Anderson and Holt (1997) using a Quantal Response Equilibrium model and longer sequences of decision makers. They find that QRE well explains the data when extended so to incorporate base rate neglect: from this they draw the conclusion that agents tend to overweight their private information, a result that seems to contradict the majority of mainstream social learning models. Nöth and Weber (2003) tried varying the quality of the signal and the ordering, and found similar results, namely that people rely too much on own signal (overconfidence), which tends to reduce welfare in general. Our paper allows to disentangle the effect of the two pieces (private and public) and can offer additional evidence towards this controversy of whether agents over- or under- weight their own information.

The main challenge of empirically studying social learning is the inability to directly infer the beliefs of the acting agent and hence separate a Herd from an Informational Cascade situation. A number of different approaches have been employed to address that Nyarko and Schotter (2002) elicit beliefs using proper scoring rules instead of empirically derived proxies for beliefs. Çelen and Kariv (2004) and Çelen and Kariv (2005) propose a unique method to elicit beliefs and distinguish Informational Cascades from herd behavior in a Continuous signal and discrete action signal setting by asking subjects to pre-commit to a choice for any given public signal they receive, before seeing that signal and after seeing their own information[4].

The experimental setting most closely related to mine is that of Filippis et al. (2017). They propose a novel experimental design to study social learning in the lab. Subjects have to predict the value of a good in a sequential order. They elicit each subject's belief twice: firstly (prior belief) after the subject observes the choices of the preceding subjects, and secondly (posterior belief), after observing their signal realization, hence can disentangle learning from private signal learning. Subjects update on their private signal in an asymmetric way, they weigh the private signal as a Bayesian agent would when the signal confirms their prior, and they overweight the signal when it contradicts the prior. Where this work deviates is that I only disclose the final choice of previous agents, and the entire series of the choices of the predecessors, and not fully reveal the belief of the preceding agent as they do in their setting. I find this to be a more faithful adaption of observational learning models, and one that best represents real world situations.

Finally, the role of information instead of incentives to motivate individuals in situations of incomplete information, has also received wide-spread attention. The literature on strategic communication is vast, from the cheap talk model of Crawford and Sobel (1982), to Bayesian persuasion Kamenica

[4]For an extensive review on belief elicitation methods, see Schotter and Trevino (2014)

and Gentzkow (2011) to the more recent field of information design (Bergemann and Morris, 2019) and information management (e.g. recommender systems, Horner and Lambert (2016)). One of the main takeaway from this filed is the trade-off faced by the sender of information between committing to a level of transparency in advance, that limits her ability to send any information she likes, vs. not choosing to commit, but facing the downside that her message will have a lesser impact on the receiver. Despite the extensive treatment from the theoretical perspective, little empirical work exists on the impact of different informational structures on Social Learning outcomes. This experiment makes a contribution by fill that gap in the literature, especially that of empirical evidence of games of strategic communication , since the amount of noise that I introduce to the game is fixed and ex-ante known to the subjects.

4.3 Theory and Main Hypothesis

The main goal of this paper is to measure the effect of different information structures on the decision making of subjects, particularly how they react to different information about others, how they update their beliefs, and whether they tend to make the optimal choice, especially with regards to a fully Bayes-rational benchmark. Moreover, I look into how the group performs, in terms of the likelihood of making correct choices, when herds and cascades form, and whether they are stable, or tend to be broken. This experiment is designed to answer these questions. I refrain from presenting the model here, since it is discussed extensively in sections 3 and 4 of Chapter 3 of this book.

Instead, I will outline the main hypotheses to be tested, and briefly motivate them below, to maintain context.

Hypothesis 1 (Individual Effect): Introducing noise to public information will make agents put greater weight on their own information. The rationale being, that given some noise on the public signal, will counteract the tendency of subjects to disregard their own information vs. what others do, especially as the size of history increases, hence the role their private signal plays in their final decision increases.

Hypothesis 2 (Group Efficiency): Obfuscation leads to increased social welfare, i.e. a reduction in the precision of the publicly available information, can lead to more efficient learning, measured as how often members of the group make correct guesses, and how often subjects herds towards the right choice

The rationale behind the expected hypothesis, is that agents that upon their private signal and what their observe from others, will rationally choose to disregard, or put a lower weight, to their own signal, as the amount of information coming from the actions of others increases. Such conformity prevails even without assuming any behavioral biases, and in fact it is a unique equilibrium outcome in a setting such as the one I focus on, referred to in the literature as the Martingale Convergence Theorem (Chamley, 2003). However, there does exist a caveat, that despite each individual is endowed with a signal that in expectation is more likely to be right than wrong, convergence to the true state at the equilibrium is not guaranteed. This occurs because agents might suboptimally end up imitating the wrong actions, creating a cascade of wrong decisions, because public information overwhelms and privately held decision. The proposed treatment aims at creating an incentive for the subjects to rely more upon their own information, hence providing a more informative signal to the subsequent players. This experiment aims to determine whether subjects will in fact respond to the treatment in the way the theory predicts, and if the result will be the one predicted.

4.4 Experimental Setting

The experimental sessions were run at the Experimental Laboratory (Espai Lab) at the Department of Economics at the Autonomous University of Barcelona (UAB) during the months of February and March of 2020. The subject pool mainly consisted of undergraduate students of various disciplines at UAB recruited through ORSEE (Greiner, 2015). In total, 80 subjects participated, with each participating in only one session (the number of subjects was lower than the original planning, but the additional sessions had to be cancelled do to the lock-down imposed as part of the COVID-19 pandemic. Additional sessions and alternative treatments will be run as soon as conditions allow it). The sessions started with written instructions given to all subjects and read out loud along with a consent form. Subjects could ask clarifying questions, which were answered privately. The experiment was programmed and conducted using the program Z-Tree. Each session consisted of 8 participants (with one exceptions of 7 due to no-shows). The number of rounds were 10 for the control group and 9 for 8 the treatment one[5]. The with the first one being a practice one for the subjects to fully internalize the procedure (they were clearly informed about this fact).

To make the experiment more engaging, the subjects were given the scenario of deciding on whether they should buy a product or not, whose quality, either Good or Bad, is unknown to them. If the quality if Good, the product is worth 100 points, and 0 if Bad. So, the optimal strategy is to Buy if

[5]This decision was made after observing that rounds took longer than expected, so to ensure no time pressure, or potential mistakes from inattention of the subject the rounds were reduced by one in the treatment group

quality is Good and Not Buy if quality is bad. However, the point of interest is finding out the subject's belief about the likelihood of the state being either Good or Bad expressed on a scale from 0 to 100 (with 100 being certainty that state = Good). Additionally, I want to control for any possible strategic reporting, and make it Incentive Compatible for the subjects to reveal their true beliefs.

To elicit the beliefs, I used a quadratic scoring function, a quite standard elicitation method. The specific formula is based on Nyarko and Schotter (2002) and is structured as such as to make a dominant strategy for the participants in terms of expected payoff, to report truthfully the exact number that is their belief about the underlying state. It was explicitly stated to the subjects, and was demonstrated with examples, why truthful reporting is the best way to maximize their payoff. The explicit formula that determines their payoff is the following:

$$\textbf{Round Payoff} = 100 - 0.01 \times (\theta - c)^2$$

Where θ is the true quality of the product. (Good Quality Good Quality =100 and Bad Quality =0) and c your choice in the given round, (from 0 to 100, integers only).

Here I describe the procedure in the control group (afterwards I will present the changes for the treatment group)[6]. The steps were the following:

1. Each session consisted of several rounds. At the beginning of each round, the computer program randomly chose the quality of the product. The value was either Good or Bad with the same probability, independently of previous realizations.

2. In each round I asked all subjects to make decisions in sequence, one after the other. For each round, the sequence was randomly chosen by the computer software. Each subject had an equal probability of being chosen in any position in the sequence.

3. Participants Ire not told the value of the good. They knew, however, that they would receive information about the value, in the form of a symmetric binary signal. If the value was equal to Good, a participant would receive a message saying 'Good' with probability 2/3 and a 'Bad' with probability 1/3; if the quality was good, the probabilities were inverted.

[6]The entire instruction set the subjects received is available in the appendix of this chapter

4. As I said, each round consisted of 9 and 8 periods for the control and the treatment group respectively. In the first period a subject was randomly chosen to decide. He received her private signal and chose a number between 0 and 100, to reflect her belief about the true state

5. In the next period, and every subsequent period after that and until the end of the round, the choice is divided into two steps. First step, the subject was first shown her private signal, and was called to choose a number between 0 and 100, to reflect her belief about the true state. Then, in a second step, the subject is shown a message reflecting the choice the previous subjects made based on their reported confidence. Meaning that if their reported belief is over 50, they chose to Buy the product, since they believe it's more likely to be Good than Bad, they chose Not Buy if the reported belief is under 50, and a random message of one of the two if they reported exactly 50. After, observing this information (along with their private beliefs), the subjects have the option to update their belief.

6. At the end of the round, after all 10 subjects had made their decisions, subjects observed a feedback screen, in which they observed the value of the good and their own payoff for that round.

The payoffs were computed based on the formula showed above, with the parameters calibrated as such to make truth-telling optimal. The payoff was shown in 'experimental tokens' where 100 tokens correspond to 1.2 euro to be converted at the end of the experiment. After participants had observed their payoffs and clicked on an OK button, the software moved to the next round. Some general information about the control and the treatment groups is presented in Table 4.1.

At the end of the experiment, subjects were asked to fill in a questionnaire, to obtain some supplementary background information to check for possible covariates that may be associated with different social learning outcomes. I collected information about the subjects' age and gender, since, especially gender has been often cited as a variable that affect outcomes in experiments. I also asked about the faculty that the participant is enrolled in, since it has been often observed that students from different disciplines tend to have different behavior in the same experimental setting. Additionally, I asked questions regarding the risk preferences of the individual, which constituted of a single question asking subjects to rate themselves on a scale from 1 to 5 with the extremes being fully risk-averse to fully risk-taker, with the middle being risk-neutral. Finally, I ask them to rate themselves with respect to how they believe in the rationality of others when making decision (to obtain a rough

Table 4.1: Summary of Control and Treatment groups

	Private Signal Precision	Sessions	Subjects per Group	Informativeness	Participants	Rounds
Control	0.66	6	8	1	48	9 + 1
Treatment	0.66	4	8	0.9	32	8 + 1

estimate of whether they anticipate errors in others' choices, and whether they tend to trust information they receive, to roughly gauge whether subjects follow the procedure, and trust the information they are exposed to.

After all the rounds were played, the tokens were converted to euros, and a 5 euro participation fee is added, and at that point the final payoff for each subject is revealed to them. On average, the subjects earned 14.5 euros for approximately 90 minutes of experiment.

Treatment Group

The treatment groups played the exact same game with the added caveat that the information conveyed might not be 100% accurate. Specifically, and in order to make setting as easy to understand as possible for the subjects, there was a 10% probability that the opposite of the actual choice of the immediately preceding subject will be reported. The number was chosen such that the individual probability is relatively low, so it will only marginally affect the choice of each subject, however, the noise compounds as more subjects take actions and the length of history increases, hence maintaining a balance between the reducing the informativeness of the public information without pushing the subjects to ignore it excessively.

4.5 Results

In this section, the results of the experiments are presented and discussed, broken down into three subsections. First, I will present the main findings at the individual level, how individuals update their beliefs at each stage, and how they weigh their own information versus that from others in both the treatment and control group and relate it to how accurate their guesses are with regards to the true state. Secondly, we will look at how the group performed under the two different regimes: How often the right guesses were made, how quickly herds formed, and whether they were stable or broken at any point and compare the outcomes between the two groups. Finally, I present the results from the survey the subjects took at the end of the experiment and relate their reported characteristics to the previously presented results.

Individual Level Results

How do subjects make inference from their own signal only?

At the first stage, the subjects make their decision based solely on their private signal regarding the underlying state. Since the parameters for this decision round do not change between the control and treatment group, we pool the data together (for a total of 187 observations).

Figure 4.1 presents the distribution of choices at stage one, presented separately for the cases in which the signal the subjects received as good or bad. The top graph refers to the case of a good signal. A relatively high percentage of choices (34.5%) are consistent with standard Bayesian updating, deviating no more than 5 units; 19.5% of actions are lower than the Bayesian one and 43.3% of actions are larger. Interestingly, in 9.4% of the cases, subjects did not update their belief at all after seeing the signal, choosing an action exactly equal to 50. On the other hand, in 13% of the cases, subjects went to the boundary of the support, choosing an action 100. Finally, there is a small proposition of 2.8% of actions in the wrong direction (i.e. subjects that updated towards 0 rather than 100). The bottom graph refers to the situation where subjects received a bad signal. The situation is the nearly the inverse of that described by the previous case of good signal, with a mode of about 30, masses of 12.8% in 50 and 12.4% in 0, with other actions distributed similarly to the good signal case.

I endeavor attempt to interpret the result of the subjects', choices of assigning different weight to a similar signal (i.e. attributing different precision to the received signal) with a simple model.

$$\alpha_{1i} = 100\left(s_{1i}\frac{q^{\alpha_{1i}}}{q^{\alpha_{1i}}+(1-q)^{\alpha_{1i}}}+(1-s_{1i})\frac{(1-q)^{\alpha_{1i}}}{q^{\alpha_{1i}}+(1-q)^{\alpha_{1i}}}\right)$$

Where α_{1i} is the weight put on the signal in observation i and the precision of the signal q, considered to be always, 2/3[7]. Note that for $\alpha_{1i} = 1$ the expression gives the Bayesian updating formula, therefore $\alpha_{1i} = 1$ is the weight that a Bayesian agent would put on the signal. A value higher (lower) than 1 indicates that the subject overweight (underweights) the signal. For example, for $\alpha_{1i} = 2$, the expression is equivalent to Bayesian updating after receiving two conditionally independent signal and can therefore be interpreted as the action of a Bayesian Subjects action upon observing it ($\alpha_{1i} = 50$), whereas a

[7]Recall that a subject made many choices in the same experiment, since he participated in several rounds; the index i refers to the observation i at time 1, and not to the subject acting at that time. Of course, the same subject could have chosen different weights in different decisions.

Figure 4.1: Distribution of actions at time 1. The top (bottom) panel refers to actions upon receiving a good (bad) signal

Table 4.2: Distribution of weights on private signal for actions at time 1. The table shows the quartiles of the distribution of weights on private signal for actions at time 1

	1st Quartile	**Median**	**3rd Quartile**
α_{1i}	0.71	1.0	2.02

subject who puts an infinite weight on it chooses an extreme action ($\alpha_{1i} = 0$ or $\alpha_{1i} = 100$) as if he were convinced that the signal fully reveals the quality of the good. Finally, a negative value for α_{1i} indicates that the subject misreads the signal, e.g., interpret a good signal as a bad one.

Table 4.2 reports the quartiles of the distribution of the computed α_{1i} . Note that the median is $\alpha_{1i} = 1$, indicating that the median subject is actually Bayesian.

Of course, other ways to view the results are possible. One may, for example, argue that the fact that a subject chooses 65, while compatible with Bayesian updating, is not necessarily indication that he is a proper Bayesian: he may be choosing 65 simply because that is the precision of his signal. The fact that the median subject is Bayesian for a bad signal too, however, lends some credibility to the fact that the subjects are doing more than just inputting their signal precision. Action 50 may also be the result of different heuristics. A subject may feel that one signal alone is not enough for him to make any update; or perhaps he is happy to choose the least risky action. The extreme actions, on the other hand, may be the expression of a “guessing type” who, despite the incentives given in the laboratory, simply tries to guess the most likely outcome. It should be noticed, though, that of all subjects who acted at time 1 more than once, only one chose an extreme action (0 or 100) every time; similarly, only 5.7% of them chose the action 50 every time. In previous social learning experiments, deviations from equilibrium have been interpreted sometimes as subjects being overconfident in their own signal. The current analysis shows that there is much heterogeneity in the way subjects update their beliefs after receiving a signal. Despite these subjective beliefs, there is no systematic bias to overweight or underweight the signal. As a matter of fact, the median belief is in line with Bayesian updating.

How do subjects make inference from their own signal only?How do subjects make inference from their predecessors’ actions?

I now analyze how subjects make inference about the underlying state when only observing the actions of others. I use the 1st stage of decision of the subjects from the 2nd round. I will focus for ease of exposition only of the 1st stage action of the 2nd subject, since his decision is the simplest to analyze given that it is predicated on a single piece of information, namely the information he received on what the 1st subject did. I denote the choice by α_2^1 (meaning the 1st stage decision of the 2nd subject in line in each period).

A subject at time 2 has to infer what information the previous subject holds based only on the message he received (’Good’ if the predecessor chose >50, and ’Bad’ if he chose <50). We saw from the previous section that it is rare for subjects to update in the ’wrong’ direction (i.e. chose an action over 50 upon receiving a bad signal and vice versa), so the subject at time 2 could consider a message Good to imply the signal of the previous subject is Good, and a ’Bad’ message to imply that the previous subject holds a Bad signal.

I group the observations into two groups. The first being the one where the message from the first subject was ’Good’ and one where the message was ’Bad’. We also further subdivide the groups into control and treatment, since in the treatment group, there is a 10% probability that the subject at time 2 may receive the opposite message than the true one. The results are summarized in the following figures and discussed below.

As we can see from Figure 4.2 and compared with Figure 4.1 from the previous section, we see more concentration of mass around 50 in both the cases that the message subject 2 received was ’Good’ and ’Bad’ and a lower

Figure 4.2: Distribution of actions at time 1. The top (bottom) panel refers to actions upon receiving a good (bad) signal

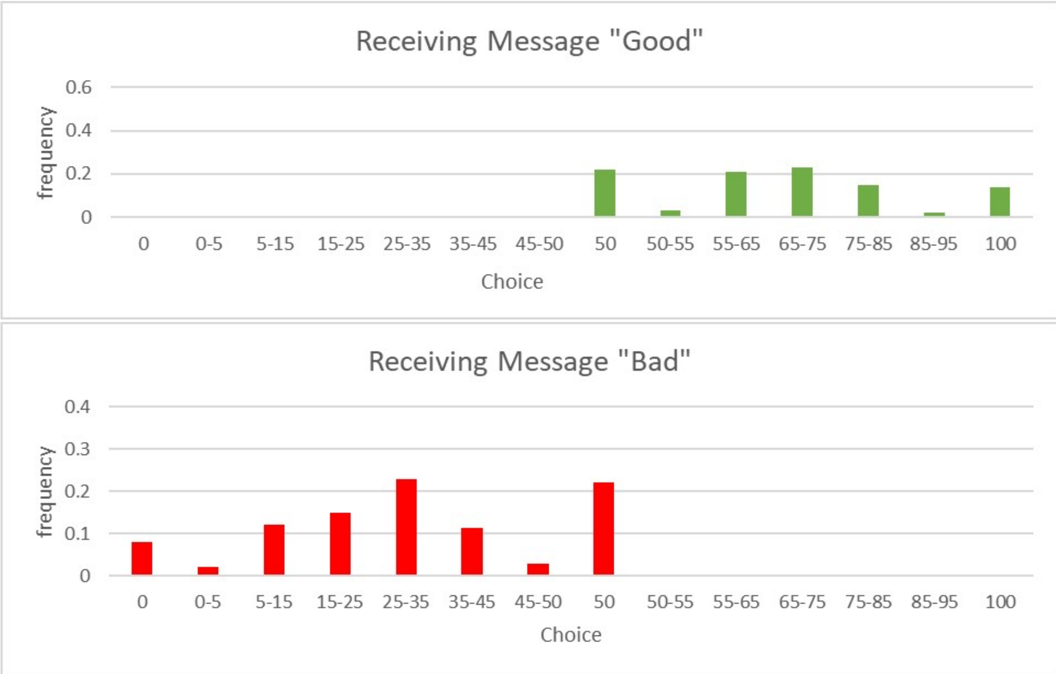

one around the 65 and 35 respectively. This is to be expected, since subject 2 at stage 1 is faced with the same decision as subject 1, and based on the same piece of information, with an additional degree of ambiguity from having to infer this piece of information indirectly. Figure 3 shows us the distribution of first actions in the treatment group, where we don't see substantial differences in choices compared to the control group, since the amount of additional noise is only marginal.

In Figure 4.4 I present the difference between the actions α_2^1 and action α_1 that the predecessor has chosen (remember that this is not directly observable by subject 2, only an imperfect signal of it). Due to the very small differences observed between the control and treatment groups in this respect, I pool all the data from both groups in the same figure for ease of exposition. A value around zero would signify that subject 2 imitated perfectly the decision of the 1st subject to play. We observe the main mass of subjects to follow closely the decision of their predecessors (about 1/3 of the cases), with also a considerable degree of deviation, which is also to be expected given the inherent noisy learning. I also, quite predictably, see that the choices of the 2nd subjects are skewed below the mean, indicating that on average, the 2nd subjects will choose a lower value to his predecessor, consistent with Bayes rationality.

I modify the model from the previous section, to generate some predictions regarding the choice subject 2 is faced with, namely that he observes an

Figure 4.3: Distribution of first actions at time 2 (the top panel refers to the case where message received is Good and bottom to message is Bad) in the treatment group

imperfect signal of subject 1's choice. If message is 'Good', he knows that $\alpha_1 > 50$ and chose α_{2i}^1 such that:

$$\alpha_{2i}^1 = 100\frac{q^{\alpha_{2i}^1}}{q^{\alpha_{2i}^1} + (1-q)^{\alpha_{2i}^1}}$$

and analogously for the case of message 'Bad', he knows that $\alpha_1 < 50$ chose α_{2i}^1 such that:

$$\alpha_{2i}^1 = 100\frac{(1-q)^{\alpha_{2i}^1}}{q^{\alpha_{2i}^1} + (1-q)^{\alpha_{2i}^1}}$$

Based on this model, we could expect the same distribution of choices for both the 1st and the 2nd subject, since the 2nd subject can perfectly infer the private information of the 1st subject by the message he receives. As we will see, the data does not support such prediction. Given the very well defined and controlled setting, and the lack of any incentives for subjects to be deceitful towards each other, the most plausible explanation for these deviations is that there a subject might hold subjective beliefs about the ability

Figure 4.4: How subjects reacted to the information received from the predecessor. Control and Treatment

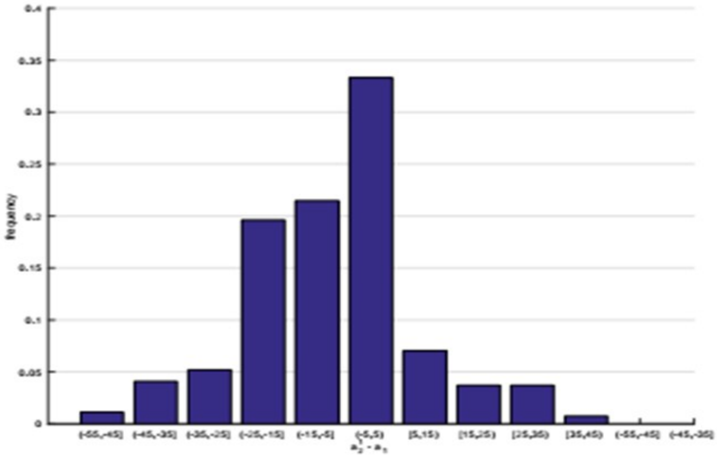

of his predecessor to correctly learn from the information he has available. I include a question in the questionnaire that subjects answer at the end of the experiment, to test if there is a potential link, and we present the results at a subsequent section of this article.

By applying this simple model, I obtain the results shown in Table 4.3 (for $q = 0.66$). The median weight is marginally lower than 1 and the 1st and 3rd quartiles are 0.13 and 1.4 (versus 0.81 and 2.05 at time 1), capturing the notion that subject seem to 'discount' in a sense the information transmitted in the action of their predecessors.

As mentioned earlier, in the presentation of the results regarding learning from the actions of others only, I decided to pool the data from the treatment and the control group due to the lack of significant differences in distribution of outcomes, and to take advantage of the larger sample size for our analysis. We have to point out that this lack of statistically significant difference in average choices when introducing additional nice might not be the case if the degree of noise increases, and remains to be investigated in subsequent studies using the same experimental setting and a treatment with a higher degree of signal noise about others' action.

How subjects weigh their private information against the actions of others?

We will now look into how subjects weigh their private information against that which they receive from others. As discussed in a prior section of this article, there does not seem to be a consensus in the literature regarding how

Table 4.3: Distribution of weights for first actions at time 2. The table shows the quartiles of the distribution of weights for first actions at time two. The message at time 1 is considered as a signal of precision 0.66 for the subject at time 2

	1st Quartile	**Median**	**3rd Quartile**
α_2^1	0.13	0.93	1.4
α_2^1 **with good message**	0.4	0.48	0.9
α_2^1 **with bad message**	0.9	1.31	2.8

subjects assign weight to each component, with a number of empirical studies indicating[8] that subjects tend to be 'overconfident' as compared to a full rational Bayesian benchmark. The previous results that looked into learning separately from private and public information, failed to detect any systematic biases within each component. I now take advantage of the setup that exposes the subjects to each of the two pieces of information in different stages, to identify the weight each component holds. I will focus again on the 2nd subject case only since it allows clearer identification. Subsequent players' weighting will inevitably be affected by the sequence of messages received.

I refer to the 1st stage decision of the subjects at time 2 as their 'prior belief' and the 2nd stage decision as their 'posterior belief'. I will start with the control group where no noise exists regarding the action of the predecessor and move on to that for the treatment group where the signal is noisy. Figure **??** shows the frequency of posterior belief conditional on whether the subject receives a signal that agrees with his prior belief (i.e. The message receives agrees with the private signal he receives), or in disagreement (i.e. message and private signal are contradicting each other)

This figure is obtained after transforming an action $\alpha_{2i}^1 < 50$ into $100 - \alpha_{2i}^2$ and the corresponding signal s_{1i} into $1 - s_{1i}$.

If subjects acted as in the Perfect Bayesian Equilibrium , in the case of confirming signal we would observe the entire distribution concentrated on the range above 66. However, the empirical evidence displays significant heterogeneity. Despite that, the median action is close to the PBE. On the other hand, When the signal is contradicting the public information, while the PBE would suggest a choice centered around 50, we observe significant asymmetry in our data, with more than 70% or the mass below the midpoint. I

[8]See for example, Nöth and Weber (2003)

Figure 4.5: Distribution of α_2^2 given $\alpha_2^1 > 50$ and confirming (top) or contradicting private signal in the control group

modify the previous model to try and put some structure to our data.

$$\alpha_{2i}^2 = 100 \frac{q^{\alpha_{2i}^2} \frac{\alpha_{2i}^1}{100}}{q^{\alpha_{2i}^2} \frac{\alpha_{2i}^1}{100} + (1-q)^{\alpha_{2i}^2} (1 - \frac{\alpha_{2i}^1}{100})}$$

When good signal is observed, and analogously, when observing a bad signal

$$\alpha_{2i}^2 = 100 \frac{(1-q)^{\alpha_{2i}^2} \frac{\alpha_{2i}^1}{100}}{(1-q)^{\alpha_{2i}^2} \frac{\alpha_{2i}^1}{100} + q^{\alpha_{2i}^2} (1 - \frac{\alpha_{2i}^1}{100})}$$

Table 4.4 presents the results for the control group. In the case of a confirming signal the median subjects put only a slightly lower weight on the signal than a Bayesian agent would do theoretically (median of 0.96). In the case of a contradictory signal, the observed weight is quite higher, namely 1.15. A different weight is observed also for the 1st and 3rd quartiles. Essentially, subjects update in a symmetric way, depending on whether their private signal agrees or disagrees with their prior held beliefs from observing the action of

Table 4.4: Distribution of weight assigned on the own signal in the control group

	1st Quartile	Median	3rd Quartile
α_2^2	0.68	0.96	1.89
α_2^2 **with confirming signal**	0.54	0.82	1.35
α_2^1 **with contradicting signal**	0.9	1.15	2.73

others. Contradicting signals are being over-weighed with respect to the fully Bayes-rational benchmark.

Figure 4.6 presents the distribution of posterior beliefs for the treatment group, in the case of confirming (upper panel) and contradicting signal. We see a bigger concentration of mass towards the two extremes of the distribution, above the 65-75 mode we observed in the control group for confirming signals, and below the 25-35 mode for the case of contradicting signals. Such an observation leads us to conclude that there was a strongest reaction to their respective private signals for the subjects in the treatment group, since they were aware, that the public information involved a degree of noise.

In Table 4.5, we can confirm that by comparing the results with the predictions from the model for the treatment group. We still see the same tendency to react stronger to a contradicting signal compared to one that is agreeing with public information as we see in the control group, with a median of 1.29 on average, and 1.70 in the case when receiving a contrarian private signal. However, we also see that the reaction to any type of private signal is significantly stronger than what the corresponding weights are in the control group.

Finally, I present the results of an independent-sample t-test, to determine whether the difference observed in the weighting of the private information vs. the public one between the treatment and the control is significant. I compare the two distributions using a two-sample Kolmogorov-Smirnov test of equality of distribution functions, which turns out to confirm the differences in distributions of posteriors between control and treatment are significant (p=0.05, corrected). This allows us to Result 1 presented below.

Result 1: Overall, the data suggest that subjects have the tendency to place higher weight to their private information in the presence of noisy public information in accordance with hypothesis 1. Moreover, subjects react more

Figure 4.6: Distribution of α_2^2 given $\alpha_2^1 > 50$ and confirming (top) or contradicting private signal in the treatment group

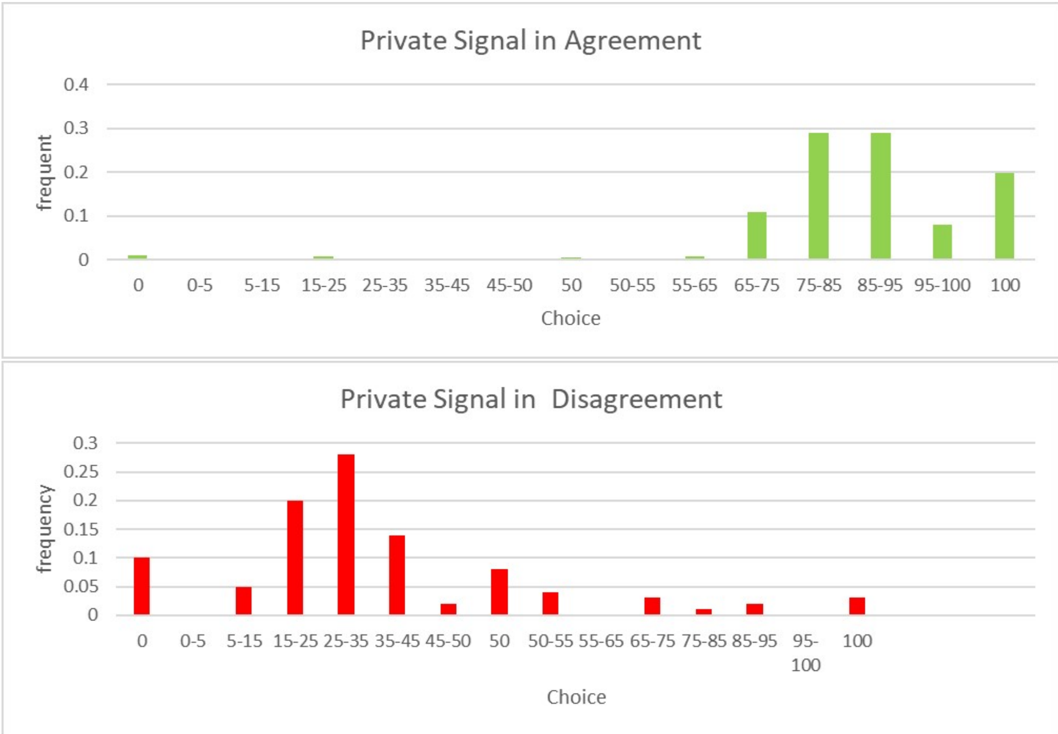

strongly to private signals that disagree with the public information compared to private signals in both the treatment and the control group.

Group-level Results

The analysis now turns to the group behavior observed throughout the experiment. When it comes to social learning, arguably the most important dimension with which to evaluate the performance of the group, is how many of its members are making the correct choice. The experimental setting allows us to easily quantify the effectiveness of the learning that occurs within each group, since we get a direct measure of the payoffs the subjects received for their choices, which is a perfect measure of welfare, both individually and for the group. I make a comparison of the treatment and control groups to answer the main research question of this paper. Additionally, I look into other features of group behavior that are important for social learning, namely, herd behavior, cascade behavior and the stability (or not) of those phenomena. To make definitions clear, we will particularly be looking for the following kinds of behavior.

Herd behavior is observed when the final choice of the subject (i.e. what he chooses in the 2nd stage of his play) is the same as the predecessor and

Table 4.5: Distribution of weight assigned on the own signal in the treatment group

	1st Quartile	**Median**	**3rd Quartile**
α_2^2	0.72	1.29	2.09
α_2^2 **with confirming signal**	0.61	0.92	1.69
α_2^1 **with contradicting signal**	0.9	1.70	2.80

Table 4.6: Kolmogorov-Smirnov test of equality of distribution functions. The significance level is denoted as follows: *** p <.001, ** p <.05, * p <.10.

Group	**D**	**p-value**	**Corrected p-value**
Control	0.1438	0.001	
Treatment	-0.0078	1.981	
Combined K-S	0.1438	0.07	0.05**

his private information agrees to the predecessors' action. This means that even though the subject appears to be mimicking its predecessor, it is also because of the fact that his private information is in confluence with what he observes from others, therefore, he still ends up communicating useful information for others, through this choice. *Cascade behavior,* is observed when the subject chooses to imitate the predecessors' action, and continues to do so at the 2nd stage of his decision round, despite the fact that his private information is contradicting this choice (for example, he will choose to play > 50 when he receives the message 'Good', despite his private information say 'Bad'). This means that he is significantly convinced that the others know better, and he will completely ignore his private information. Such behavior does not contribute anything towards the collective information of the group since it is an unconditional imitation of others.

It is clear that Herd and Cascade behavior are observationally equivalent, and therein lies the core challenge that social learning mechanisms face. These two being indistinguishable, can lead to the propagation of a sub-optimal action (for example the entire group choosing Good, when the true state is Bad and vice versa). The controlled environment of our experiment and the feature of two stage choice allows us to identify the two. Finally , I will look into the question of whether herding and cascade behavior is stable or can break once it starts. Early theoretical work (Bikhchandani et al., 1992) posited that once a herd started it was impossible to break, particularly in their setting of binary actions and bounded beliefs, however Smith and Sørensen (2000) showed that herds can be broken with positive probability, as action space gets richer. The experiment can offer some data to shed some light on this debate as well. Table 4.7 summarizes the findings at the group level.

Regarding herd behavior, one of over 5 subjects was observed at 23 out of 54 rounds (42.5%) for the control group and 12 out of 32 rounds (37.5%), indicating no statistically significant difference between groups according to the Mann-Whitney U test. A herd of at least 3 subjects was observed on all rounds, and all subjects herding on the same action was observed in 3 rounds in the control and 2 in the treatment group. Interestingly, only 4 herds turned out to be wrong in the control and 2 at the treatment group, which the difference between marginally significant (p=0.06). On the other hand, we observed a significant decrease in cascade behavior between the treatment and control groups (p=0.04) from a 35.1% to 25%.

This finding is consistent with what the hypothesized mechanism. Adding noise in the public signal may push subjects to assign greater weight to their privately held information than what they would in the case of perfect information about the actions of others. Moreover, the fact that herds occur often in our result, but are not necessarily cascades, is consistent with the Martingale

Table 4.7: Group-level Results
The significance level is denoted as follows: *** p < .001, ** p < .05, * p < .10.

	Control	Treatment	Mann-Whitney U Test for differences (p-values)
Average Earnings (in Euros)	14.15	15.21	0.008***
Herds (5 or more)	23	12	
% of Herds (out of rounds played, 54 and 38 respectively)	42.5	37.5	0.13
Incorrect Herds	4	2	0.10*
Cascades	19	8	0.04**
% of Cascades	35.1	25	
Overturns	120	88	0.03**
Overturns (378 and 224 decision points respectively)	31.7	39.3	

Convergence Theorem (Chamley, 2003) that states that in a bounded belief setting such as ours, public belief with eventually converge. This conclusion is further corroborated by the fact that the probability of an overturn, i.e. the breaking of a herd after it has formed is higher in the treatment group where we observed in 88 out of 224 decision points that occurred (39.3%) versus 120 out of 378 (31.7%) in the control.

Coming to the most crucial variable related to the efficiency of social learning, I find that the average payoff the subjects received is higher in the treatment, compared to the control group. Specifically, subjects on average received 15.21 euro of final payoff whereas those in the control group averaged 14.15 on comparable sessions. This increase of roughly 7% on average is statistically significant (p=0.008) and indicates that our Hypothesis 2 regarding efficiency is supported by the finding of this study. This is reflected by the fact that, given our payoff scheme, which is designed to make the true beliefs that subjects hold about the likelihood of the true state being either Good or Bad incentive compatible, higher payoff indicates a better informed decision on behalf of the subjects. The fact that the average payoff is higher in the treatment group despite the negative effect of introducing imperfect information about others, and the findings regarding the weighting of private information presented on a previous subsection lead us to the following result.

Result 2 (*Group Efficiency*): Under the treatment of imperfect information about the actions of others, we can observe an increase in the efficiency of social learning compared to the perfect information control, measured by an increase 7% in the average payoff of the subjects.

That being said, one caveat that I wish to point out is that the amount of noise introduced is relatively low (only 10% probability of misreporting of each subject's action), therefore it can be that further noise in the public signal might have an adverse effect in learning efficiency, hence the positive effect of subjects' individual signal being more informative, may not compensate the adverse effect of a noisy signal overall. Answering the question of the optimal level of noise on social learning efficiency is left for a subsequent study.

Survey Results

For the final part of the results, I now combine game behavior with reported answers from the questionnaire subjects were called to fill in after completing the experiment, to examine how a number of variables could affect the performance during the game in terms of profitability (which in term is a proxy for how well-informed decisions the subjects ultimately make), and how they update their beliefs, as measured here as the absolute difference between prior and posterior, meant to capture how strongly subjects change their mind between the two stages. The variables related to the subjects that I collected are the following: The subjects' risk-preferences, their perception about the rationality of others when making decisions, and their gender.

The risk-preferences allows us to ascertain whether the attitude towards risk of the subjects affect the way they perceive their private information vs. what others are doing, and whether it affects their likelihood of ultimately making correct guesses. It is measured on a 5-point Likert scale that we equate to a risk-level classification from risk-averse, to risk-neutral, to risk-seeker. We also asked the subjects to report how rational in general they find the choices other people make, on a similar 5-point Likert scale. This would allow us to spot whether this has any impact on how subjects weight their own information against the actions of others, and spot potential effects of ambiguity in forming posterior beliefs, that could explain deviations from Bayesian Rationality that we spotted in our experimental results. Finally, we also look into whether gender has an impact on the results, since a disparity in results between genders is often cited in the literature.

A series of Spearman rank-order correlations were conducted to determine if there were relationships between game behavior and survey answers. Table 4.8 provides correlations among the outcome variables, and the reported ones.

Regarding the effect of covariates on belief update, we find a statistically significant positive correlation (Spearman's $\rho = .199, p = .047$) with risk preferences. Note that the higher the risk preference score the more risk-seeker the subject reports himself to be. So, this positive coefficient can be interpreted as

Table 4.8: Correlations among outcome variables, and subject-specific variables.
Notes: Spearman's rank correlations. The table reports the correlation coefficients. The number of observations is presented in parentheses. The significance level is denoted as follows: *** p < .001, ** p < .05, * p < .10.

	1	2	3	4	5
1. Risk Preferences	-	0.03 (71)	-.106* (71)	.199** (63)	-.073 (71)
2. Rationality of Others	0.03 (71)	-	0.85 (71)	.125 (63)	.120** (71)
3. Gender	-.106* (71)	0.85 (71)	-	-.126** (63)	.062 (71)
4. Updating of Belief	.199** (63)	.125 (63)	-.126** (63)	-	.233** (71)
5. Total Profit	-.073 (71)	.120** (71)	.062 (71)	.233** (71)	-

the fact that more risk-seeking individuals will be more affected on average by the revelation of their private information when updating their beliefs, as compared to more risk-averse subjects. This is an interesting finding worth future investigation, since it could be one of the drivers for a common empirical finding of over-confidence with regards to private signals, that is observed in the literature. We do also find a significant negative correlation of gender on belief updating (Spearman's $\rho = -.126, p = .039$) meaning females tend to update their belief to a lesser degree on average compared to males between the two stages, meaning that they seem to assign less weight to their private information relative to males when making their decision. There should be noted though that we do find a negative correlation between being female and being a risk-seeker, which in turn is positively correlated with the weight put on private information, so potential for endogeneity can exist here.

With regards to total profits, the only statistically significant relation appears to be with the belief in rationality of others, where we have a positive relation between the degree of trust in the choices of others (Spearman's $\rho = .120, p = .041$) and how well subjects perform in the experiment. An interpretation that can be given to this result can be that those people with the higher degree of belief in others' rationality, can make the most effective inference from the public information, and probably understand better how Bayesian updating works, as well as the game in general, being able to play it better than those who do not view their fellow subjects as rational decision makers, hence assume a higher degree of error in their choices. We finally see a positive correlation between the two outcome variables (belief updating and total profits), which is in line with our initial hypothesis, and we discussed in a previous part of the result section.

4.6 Conclusion

In this chapter, I studied how learning from the actions of others, both at the individual and at the group level takes place in a control laboratory setting. This setup enables the quantification of the impact from two separate pieces of information for each subject, namely the privately held information of each individual, and the information he receives about the actions of others. This specific experimental setup also allows to discriminate between two types of behavior that are observationally equivalent under normal circumstances. The first being herding behavior, where a subject imitates the action of the predecessor, but does so because his private information agrees with that choice, in which case his action does carry some additional information, and the second being cascade behavior, where the subject completely disregards any information he might hold and unconditionally imitates others, contributing nothing towards the public information available to the group.

I find significant efficiency gains from introducing a small degree of noise in the public signal that each subject receives compared to the full information case in the control group, measured as a statistically significant, 7% increase in average payoff for each participant in the treatment group. This result is in line with the initial hypothesis, that despite the direct adverse effect of introducing noise, there is an indirect benefit, obtained by the subjects relying more on their private information, on average, hence making each individual action carry more information. This is corroborated further by the data, since I find a decrease in cascade behavior from 35.1% of the rounds in the control to 25% of the rounds in the treatment group, despite not observing changes in the degree of herding between the groups. The increased reliance on private information is something that I also observe at the individual level data when comparing the weight subjects assigns to their own vs. the information they see from others. Specifically, I observe a significant increase in the weight assigned to private information in the treatment group compared to the control with the full disclosure of public information.

The findings of this experiment relate to two different strands of literature: that of social learning, and of information design. I implement a novel design that better mimics the real-world challenge of inferring information from the actions of others, and allows me to disentangle learning from others, and learning from individuals' private signal and how subjects update their beliefs between these two stages. As I discussed in the introduction, the range of applications of a policy that optimally harnesses the privately held information of individuals and efficiently aggregates it to better inform the group is wide.

These results, in addition to testing the validity of mainstream theoretical models about how social learning works, where it fails and where it works, also provides some real-world policy implications. Primarily, it highlights the exploration/exploitation trade-off that any social learning mechanism faces, where an optimal balance between being informative and encouraging experimentation in order to make better recommendation is present, as I discussed in the introduction. There is a benefit to allowing a degree of uncertainty with regards to the actions of others, in order for the mechanism designer to incentivize individuals to reveal their private information through their action and hence provide a stronger signal back to the designer. I mentioned a number of social learning applications regarding internet platforms and user feedback, but our results can also apply to government policy. For example, think of a new innovation, that has potential to either increase, or decrease the productivity of firms. For the government to collect enough evidence to be able to either push for mass adoption or not, it requires a number of firms to experiment with such technology individually, in order to observe its effectiveness. How the government chooses to disclose this information might have important welfare implications, since a few initial false positives (false negatives) might lead to other firms to adopt (reject) a sub-optimal technology, based on insufficient data. Hence, a strategic obfuscation of information might have beneficial welfare implications, given specific circumstances.

In terms of future extensions of this line of research, this experimental setting allows for numerous directions. The most apparent one, (which unfortunately is not part of this article due to the disruption caused by the nation-wide lock-down during the COVID-19 pandemic), is an additional treatment group, to try how subjects respond to greater degrees of noise in the publicly available information. This would allow us to find the optimal spot, where informativeness vs. exploration inducing give the best result in terms of group efficiency. Additionally, this setup allows us to test a number of different models of communication and information disclosure, and measure how subjects respond them, and if their behavior is consistent to that which the theoretical literature prescribes. Particularly, the role of commitment and verifiability of information can be tested, with models ranging from Cheap Talk (Crawford and Sobel, 1982) where no commitment is implied to Bayesian Persuasion (Kamenica and Gentzkow, 2011), where the sender commits ex ante to a disclosure policy, and different intermediate settings.

4.7 Appendix

Instructions for Control Group (Translated from original version of Spanish)

Welcome

You are now participating in a scientific experiment that studies decision making and learning from the actions of others. During this experiment, you will play a game. Please read the following instructions about how to play the game very carefully, since, understanding and playing well may lead you to a significantly higher payoff. All the decisions you make during the experiment will be completely anonymous. You are welcome to ask the experimenters any questions regarding the experiment by raising your hand. Apart from these questions, any other sort of communication is prohibited and may lead to your immediate expulsion from the session.

Session Outline

The experiment will last for 12 decision rounds. The first 2 rounds are going to be practice ones, meaning that you will not be receiving payoffs, for them. These are meant for you to better familiarize yourself with the game and its rules. Then, 10 rounds of regular play will follow. The rules are exactly the same in all rounds, but each round is completely independent from any other. At the end of the experiment, you will answer a short questionnaire and then you will receive your payoff based on your performance in the last 10 rounds (details will follow), in addition to the 5-euro participation fee.

Description of a Decision Round

The scenario of the game is as follows: At each round, a new product comes. Your goal is to guess the quality of this product. It could either be of good quality (G) in which case its worth 100 points, or of Bad quality (B), and is worth 0 points. At the beginning of each round, the computer will randomly determine the quality of the product with each probability for each case (50% each). You will then be randomly matched with 7 other participants and take turns, one after the other to make your guess about the quality of the product. However, instead of simply deciding between a guess of Good or Bad quality, you will be asked to give your estimating about how likely the quality of the product is a good or a bad one. You will express your estimation on a scale of 0 to 100 (integers only), with 0 meaning that you are absolutely certain that the product is of bad quality, 100 that you are absolutely certain that the product is

good and any value in between show your degree or uncertainty. The number you report is the only determinant of your payoff in each round, so be very careful in what choice you will make. How the payoff is determined based on your choice will be explained with details below.
To assist you in your decision, you will be given two pieces of information:

1. Each participant will receive a signal regarding the product quality. This signal is private information for each participant. The signal is a message telling you either that the quality of the product is good or bad. The precision of this signal is 66.5%, meaning that with probability of 66.5 % the signal will be correct (2 out of three times), and with probability 33.5% it will be incorrect (1 out of 3 times). For example, say that the product in this round is of Good quality, then you will receive a message telling you the quality is good with 66.5% chance, and that the number is 0 with 33.5% chance.

2. In addition to your private signal, you will also be informed about what the participants that had played before you in the round had chosen. You will not be given their exact choice of number. Instead you will be getting the message "Good" if they choice was over 50, 'Bad' if they choice was under 50 and a random message of either 'Good' or 'Bad' with 50-50 chance if they chose exactly 50. The number of choices you will see depend on your position in the order of play. For example, the 1st participant to play will only have her private information available for her, whereas the 8th participant to play (the last of the round) will know all the choices of the previous 7 players. Remember that you will receive an imperfect signal of each participants choice and not their confidence in their guess.

Note that each round is independent from any other and so are all the information you might get in each round, as well as the outcomes and who you will be playing with are entirely unrelated to each other.

Each decision round goes as follows:
1st Step: You will be shown a message about the choices of the preceding players (Good or Bad) with the order they played, and you will be asked to choose a number between 0 and 100 (only integers are accepted), which should reflect your confidence about which of the two number is more likely, as explained above. You then click continue to go to the next step.
2nd Step: In the next step, more information will become available to you. Your private signal is revealed, and you will be given the opportunity to update your estimation by again picking a number between 0 and 100. You can choose

to not change your estimate by inputting the same number as before. The only exception to this rule is the first player, who will only make her choice based on her private signal, since no previous information will be available to her. You then confirm your choices are saved and wait for the round to finish. Once all the participants have made their choice, the true quality of the product as well as your payoff for this round will be revealed to you. You will then have to click to continue to the next round until all of them are played.

Payoff Calculation

The payoff for the round is determined using the following formula. :

$$\textbf{Round Payoff} = 100 - 0.01 \times (\theta - c)^2$$

Where θ is the true quality of the product. (Good Quality Good Quality =100 and Bad Quality =0) and c your choice in the given round, (from 0 to 100, integers only).

In each round you will be paid in experimental tokens, where 100 tokens correspond to 1.5 euro to be converted at the end of the experiment. Note that this formula is designed in a manner that the best strategy in order to maximize your payoff is to state your true belief about the quality of the product. Choosing numbers towards the two extremes (0 and 100) has the potential for more gains if your guess is closer to the true result, but also the risk of (near) zero payoff if your guess is far from it. More moderate guesses imply to lower gains but also lower risk. For example, if you guess 0 and the product is bad then you will receive the maximum payoff of 100 tokens and vice versa for the case that the product is good. Finally, note that the payoff in each round is randomly determined, either by the 1st or the 2nd decision you make with 50% chance each. So, make sure you make your best guess possible with the available information in each case.

Outline of a round (excluding the 1st player who skips steps 1 and 2, since no previous information is available)

1. A message about the choices of the previous players will be revealed to you
2. You will be called to make your estimation based on this information
3. Your private information is revealed to you

4. You will be called to update your prediction based on this additional piece of information

5. The round ends, when all players have made their decisions and a new round starts

5 Concluding Remarks

This dissertation is concerned with the study of social dilemma situations, where the individually optimal behavior does not align with the group-optimal behavior, hence creating externalities. This occurrence allows room for potential welfare improving mechanisms. The thesis is divided into three chapters, two of which are theoretical studies of two difference social dilemma situations, and the third a laboratory experiment. The second chapter studies the role of social norms in the equilibrium levels of cooperative behavior in a public goods game. The third chapter studies a social learning game, identifies a number of inherent inefficiencies, and proposes a mechanism to address them. Finally, the last chapter is a laboratory study to test the empirical validity of the findings in the third chapter, along with a number of other results related to social learning. In the following paragraphs, I shall briefly summarize the main results for each chapter and conclude with some implications and potential avenues of further research.

In the 2nd chapter of my thesis, I propose a model based to study the interplay of laws and norms in enforcing cooperation in a case of "social dilemma" and more specifically in a game of public good provision. I assume that free-riding is looked down upon, but that the more people free-ride the weaker the strength of the norm becomes. Individuals face two choices one economic, whether to free-ride or cooperate, and one political, how high the level of public good provision should be. I capture the idea that although many societies have very comparable law enforcement mechanisms to discourage free-riding (such as tax-evasion), large disparities in outcomes is observed, which can be attributed to more lax social norms against such behaviors. I conduct analysis on how the presence of norms affect cooperation and propose the optimal enforcement policy to achieve optimal cooperation. I find that a strict law enforcement policy may actually have an adverse effect on cooperation since it leads to break-down of the social norms effect, and instead advocate for a more moderate enforcement policy that complements the effects of pro-social norms. Finally, I find that an endogenously determined set of parameters (through a process of majority voting), lead to a lower level of aggregate contribution towards the public good, but also a reduced rate of free riding from what would be determined from a social planner. The immediate goal in terms of future extension, is to design an experiment that resembles the setting outlined in this paper. This is facilitated by the fact that public good games have been extensively tested experimentally, and the introduction of norms would be straightforward to implement and provide interesting evidence regarding cooperative behavior. Another planned extension of this model is to introduce a time aspect to it. The evolution of cooperative behavior over time along with the presence of a dynamically varying norm would provide a very interesting insight as to what set of policy parameters lead to more stable levels of cooperative behavior, where we might observe fluctuations, and how the two forms of enforcement mechanism interact with each other in a dynamic environment. The introduction of time in this model will allow to study more complicated settings in the sense that we can do away with the balanced budget constraint for each specific period, and make the policy analysis, especial in the case where it is endogenously determined, much more realistic and interesting.

In the third chapter, I analyze a game of social learning, where agents sequentially make a one-off choice whose payoff is a function of an unknown state variable. Each agent is endowed with an informative private signal, and based on her location in the sequence, the publicly available information up to this point, which is summarized by the public belief. Each agent makes a choice based on these two components, and public information is updated accordingly. An inherent inefficiency of such mechanism is identified, where agents tend to overly imitate others and ignore their own information, a situation referred to in the literature as an information cascade. A simple mechanism that relies on information design is proposed to address this inefficiency.

This mechanism entails the introduction of noise in the public signal. Despite the fact that the designer is benevolent, he still has to misinform a number of agents in order to induce them to act upon their own information and reveal it through their actions. To that end, the designer adds noise to the public signal, so he can learn the true state as quickly as possible. In fact, the amount of noise is shown to follow a hump-shaped trajectory, since initially herding is less likely and the designer has limited access to information hence limited capacity to obfuscate. But, as the likelihood for a sub-optimal herd increases in time the designer decreases the noise and eventually fully reveals his information hence inducing the herd . An additional result is that when the designer cannot personalize the disclosure policy to each individual agent, and communicate it, then only a second-best is possible. I refer to this case as the public disclosure case, where the policy needs to be incorporate a degree of randomization , to offset the fact that the same message may be received by all the agents at the same time period. An additional result is that when the designer cannot personalize the disclosure policy to each individual agent, and in private, then only a second-best is possible. I refer to this case as the public disclosure case, where the policy needs to be incorporate a degree of randomization, to offset the fact that the same message may be received by all the agents at the same time period.

Regarding plans for moving this research forward, information economics is a very fertile field for investigation with a potential for significant contributions. My immediate goal is to expand the social learning model to study situations where the state of nature constantly changes, to see how does the optimal mechanism change. This is very relevant in terms of real-life applications. For example, the quality of a product or a service tend to fluctuate over time, so it is interesting to see how learning reacts to that. Additionally, two other extensions are, first, to look into learning with heterogeneously informed agents and see if a mechanism exists to overstate the effect of the better informed over the less informed. Secondly, I plan to work on a model where instead of actions, the outcomes of agents are observable. This is a setting that is often the case in real life, where for example, successful businesses receive disproportionately high level of exposure and a large number of imitators, but their practices and strategy are not identifiable, nor is it possible to tell what part of their success is due to luck. The question that arises: Is it welfare-improving to try to imitate their observable characteristics? Does Information Design have a role to play in that setting?

In the last chapter, I study learning from the actions of others, both at the individual and at the group level from an empirical point of view. The experimental setting employed allows for the estimation of the effect in decision making from two separate pieces of information for each subject, namely the privately information that each individual hold, and the information he receives about the actions of others. This specific experimental setup also allows us to discriminate between two types of behavior that are observationally equivalent under normal circumstances. The first being herding behavior, where a subject imitates the action of the predecessor, but does so because his private information agrees with that choice, in which case his action does carry some additional information, and the second being cascade behavior, where the subject completely disregards any information he might.

I find significant efficiency gains from introducing a small degree of noise in the public signal that each subject receives compared to the full information case in the control group, measured as a statistically significant, 7% increase in average payoff for each participant in the treatment group. This result is in line with our initial hypothesis, that despite the direct adverse effect of introducing noise, there is an indirect benefit, obtained by the subjects relying more on their private information, on average, hence making each individual action carry more information. This is corroborated further by our data, since we find a decrease in cascade behavior from 35.1% of the rounds in the control to 25% of the rounds in the treatment group, despite not observing changes in the degree of herding between the groups. The increased reliance on private information is something that we also observe at the individual level data when comparing the weight subjects assigns to their own vs. the information they see from others. Specifically, we observe a statistically significant increase in the weight assigned to private information in the treatment group compared to the control with the full disclosure of public information.

In terms of future extensions of this line of research, this experiment is only the first part of a long-term research agenda that goes towards two directions (that often coincide as in this chapter). The first one, is to apply different levels of the same treatment (the noise of the public signal) which will allow us to find the optimal spot, where informativeness vs. exploration give the best result in terms of group efficiency. This experimental setting can also be applied to study additional instances of social learning, other than the observational one. Of equal relevance to the real world, is to study how learning occurs in settings where certain outcomes of others are observable, but not their choices (the opposite of the setting studied here). This is a relevant question, as long as the issue of imperfect inference persists, since often outcomes can be spurious, and can lead to inefficient outcomes.

Additionally, this experimental format allows us to test a number of different models of communication and information disclosure, and measure how subjects respond to them, and if their behavior is consistent to that which the theoretical literature prescribes. Particularly, the role of commitment and verifiability of information can be tested, with models ranging from Cheap Talk (Crawford and Sobel, 1982) where no commitment is implied to Bayesian Persuasion (Kamenica and Gentzkow, 2011), where the sender commits ex ante to a disclosure policy, and different intermediate settings

Bibliography

ACEMOGLU, D. AND M. O. JACKSON, "Social norms and the enforcement of laws," *Journal of the European Economic Association* 15 (apr 2017), 245–295.

AKERLOF, G. A., "A Theory of Social Custom, of Which Unemployment May be One Consequence," *The Quarterly Journal of Economics* 94 (1980), 749.

ALI, S. N. AND D. A. MILLER, "Ostracism and forgiveness," *American Economic Review* 106 (aug 2016), 2329–2348.

ALONSO, R. AND O. CÂMARA, "Persuading voters," *American Economic Review* 106 (nov 2016), 3590–3605.

ANDERSON, L. R. AND C. A. HOLT, "Information Cascades in the Laboratory," (1997).

ANDREONI, J., "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving," *The Economic Journal* 100 (1990), 464.

ASQUITH, P., T. R. COVERT AND P. PATHAK, "The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market," Technical Report, National Bureau of Economic Research, Cambridge, MA, sep 2013.

BANERJEE, A. V., "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics* 107 (aug 1992), 797–817.

BECKER, G. S., "A Theory of Social Interactions," *Journal of Political Economy* 82 (1974), 1063–1093.

BÉNABOU, R. AND J. TIROLE, "Incentives and prosocial behavior," *American Economic Review* 96 (dec 2006), 1652–1678.

BENABOU, R. AND J. TIROLE, "Laws and Norms," *NBER Working paper* (2011), 1–44.

BENOIT-SMULLYAN, E., E. DURKHEIM, S. A. SOLOVAY, J. H. MUELLER AND G. E. G. CATLIN, *The Rules of Sociological Method.*, volume 3 (The Free Press, 1938).

BENTHAM, J., "An introduction to the principles of morals and legislation," (1789).

BERGEMANN, D., B. BROOKS AND S. MORRIS, "The limits of price discrimination," *American Economic Review* 105 (mar 2015), 921–957.

BERGEMANN, D. AND S. MORRIS, "Information Design: A Unified Perspective," *Journal of Economic Literature* 57 (mar 2019), 44–95.

BERNHEIM, B. D., "A Theory of Conformity," *Journal of Political Economy* 102 (1994), 841–877.

BIKHCHANDANI, S., D. HIRSHLEIFER AND I. WELCH, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy* 100 (oct 1992), 992–1026.

BROCAS, I. AND J. D. CARRILLO, "Influence through ignorance," *RAND Journal of Economics* 38 (dec 2007), 931–947.

ÇELEN, B. AND S. KARIV, "Observational learning under imperfect information," *Games and Economic Behavior* 47 (apr 2004), 72–86.

———, "An experimental test of observational learning under imperfect information," *Economic Theory* 26 (2005), 677–699.

CHAMLEY, C. P., *Rational herds: Economic models of social learning* (Cambridge University Press, 2003).

CHANDRASEKHAR, A. G., H. LARREGUY AND J. P. XANDRI, "Testing Models of Social Learning on Networks: Evidence From Two Experiments," *Econometrica* 88 (jan 2020), 1–32.

CHE, Y. K. AND J. HÖRNER, "Recommender systems as mechanisms for social learning," *Quarterly Journal of Economics* 133 (may 2018), 871–925.

CRAWFORD, V. P. AND J. SOBEL, "Strategic Information Transmission," *Econometrica* 50 (1982), 1431.

DAL BÓ, P., A. FOSTER AND L. PUTTERMAN, "Institutions and behavior: Experimental evidence on the effects of democracy," *American Economic Review* 100 (2010), 2205–2229.

ELSTER, J., "Social Norms and Economic Theory," *Journal of Economic Perspectives* 3 (1989), 99–117.

FEHR, E., S. GACHTER AND G. KIRCHSTEIGER, "Reciprocity as a Contract Enforcement Device: Experimental Evidence," *Econometrica* 65 (1997), 833.

FEHR, E. AND K. M. SCHMIDT, "A theory of fairness, competition, and cooperation," *Quarterly Journal of Economics* 114 (1999), 817–868.

———, "Chapter 8 The Economics of Fairness, Reciprocity and Altruism - Experimental Evidence and New Theories," (2006).

FILIPPIS, R. D., A. GUARINO, P. JEHIEL AND T. KITAGAWA, "Updating ambiguous beliefs in a social learning experiment," *CeMMAP working papers* (2017).

FUDENBERG, D. AND D. K. LEVINE, "Whither Game Theory?," *Working paper* (2016).

GÄCHTER, S. AND E. FEHR, "Collective action as a social exchange," *Journal of Economic Behavior and Organization* 39 (jul 1999), 341–369.

GALTON, F., "Vox populi," *Nature* 75 (mar 1907), 450–451.

GOEREE, J. K., T. R. PALFREY, B. W. ROGERS AND R. D. MCKELVEY, "Self-correcting information cascades," *Review of Economic Studies* 74 (2007), 733–762.

GREINER, B., "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association* 1 (jul 2015), 114–125.

GUARINO, A. AND P. JEHIEL, "Social learning with coarse inference," Technical Report 1, 2013.

HAIGNER, S., M. KOCHER AND M. SUTTER, "Choosing the stick or the carrot? Endogenous institutional choice in social dilemma situations," *Available at SSRN 903911* 77 (2006), 1540–1566.

HANN-CARUTHERS, W., V. V. MARTYNOV AND O. TAMUZ, "The speed of sequential asymptotic learning," *Journal of Economic Theory* 173 (jan 2018), 383–409.

HAYEK, F. A., "The use of knowledge in society," Technical Report 4, 2009.

HENRICH, J., R. BOYD, S. BOWLES, C. CAMERER, E. FEHR, H. GINTIS AND R. MCELREATH, "In search of Homo economicus: Behavioral experiments in 15 small-scale societies," *American Economic Review* 91 (2001), 73–84.

HORNER, J. AND N. S. LAMBERT, "Motivational Ratings," *SSRN Electronic Journal* (2016).

KAMENICA, E. AND M. GENTZKOW, "Bayesian persuasion," *American Economic Review* 101 (oct 2011), 2590–2615.

KOLLOCK, P., "Social Dilemmas: The Anatomy of Cooperation," *Annual Review of Sociology* 24 (aug 1998), 183–214.

KOLOTILIN, A., T. MYLOVANOV, A. ZAPECHELNYUK AND M. LI, "Persuasion of a Privately Informed Receiver," *Econometrica* 85 (2017), 1949–1964.

KREMER, I., Y. MANSOUR AND M. PERRY, "Implementing the "Wisdom of the crowd," *Proceedings of the ACM Conference on Electronic Commerce* 50 (jul 2013), 605.

LEVINE, D. K., "Modeling Altruism and Spitefulness in Experiments," *Review of Economic Dynamics* 1 (1998), 593–622.

LEVINE, D. K. AND S. MODICA, "Peer discipline and incentives within groups," *Journal of Economic Behavior and Organization* 123 (2016), 19–30.

LINDBECK, A., S. NYBERG AND J. W. WEIBULL, "Social norms and economic incentives in the welfare state," *Quarterly Journal of Economics* 114 (1999), 1–35.

MELTZER, A. H. AND S. F. RICHARD, "Tests of a rational theory of the size of government," (1983).

MIDGLEY, L. AND M. OLSON, *The Logic of Collective Action: Public Goods and the Theory of Groups*, volume 22 of *Harvard economic studies*, 21. printi edition (Cambridge, Mass.: Harvard Univ. Press, 1969).

MYERSON, R. B., "Mechanism Design by an Informed Principal," *Econometrica* 51 (1983), 1767.

NÖTH, M. AND M. WEBER, "Information aggregation with random ordering: Cascades and overconfidence," *Economic Journal* 113 (2003), 166–189.

NYARKO, Y. AND A. SCHOTTER, "An experimental study of belief learning using elicited beliefs," *Econometrica* 70 (2002), 971–1005.

OKUNO-FUJIWARA, M. AND A. POSTLEWAITE, "Social norms and random matching games," *Games and Economic Behavior* 9 (1995), 79–109.

OSTROVSKY, M. AND M. SCHWARZ, "Information disclosure and unraveling in matching markets," *American Economic Journal: Microeconomics* 2 (aug 2009), 34–63.

RAYO, L. AND I. SEGAL, "Optimal information disclosure," *Journal of Political Economy* 118 (oct 2010), 949–987.

ROSENBERG, D. AND N. VIEILLE, "On the Efficiency of Social Learning," *SSRN Electronic Journal* (nov 2017).

SAMUELSON, P. A., "The Pure Theory of Public Expenditure," *The Review of Economics and Statistics* 36 (nov 1954), 387.

SCHOTTER, A. AND I. TREVINO, "Belief Elicitation in the Laboratory," *Annual Review of Economics* 6 (aug 2014), 103–128.

SMITH, L. AND P. SØRENSEN, "Pathological outcomes of observational learning," *Econometrica* 68 (mar 2000), 371–398.

SMITH, L. AND P. N. SORENSEN, "Rational Social Learning by Random Sampling," *SSRN Electronic Journal* (dec 2011).

SMITH, V. L., "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy* 70 (1962), 111–137.

TABELLINI, G., "The scope of cooperation: Values and incentives," *Quarterly Journal of Economics* 123 (2008), 905–950.

TAKAHASHI, S., "Community enforcement when players observe partners' past play," *Journal of Economic Theory* 145 (jul 2010), 42–62.

WEIZSÄCKER, G., "Do we follow others when we should? A simple test of rational expectations," *American Economic Review* 100 (2010), 2340–2360.

YOUNG, H. P., "The Evolution of Conventions," *Econometrica* 61 (1993), 57.