



**QUEEN'S
UNIVERSITY
BELFAST**

Altruistic observational learning

March, C., & Zieglmeier, A. (2020). Altruistic observational learning. *Journal of Economic Theory*, 190, Article 105123. <https://doi.org/10.1016/j.jet.2020.105123>

Published in:
Journal of Economic Theory

Document Version:
Early version, also known as pre-print

Queen's University Belfast - Research Portal:
[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights
© 2020 The Author.

General rights
Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Open Access
This research has been made openly available by Queen's academics and its Open Research team. We would love to hear how access to this research benefits you. – Share your feedback with us: <http://go.qub.ac.uk/oa-feedback>

Altruistic Observational Learning*

Christoph March[†] and Anthony Ziegelmeyer[‡]

Abstract

We report a laboratory experiment that tests the causal impact of altruism on observational learning behavior. Once endowed with a private signal, participants submit their guess about the payoff-relevant state in two parallel sequences. In the *observed* sequence, guesses are revealed publicly so that participants in both sequences can benefit from guesses that are informative. *Unobserved* guesses, on the other hand, never reveal any information to others as they remain private. We find that *observed* guesses are significantly more informative than *unobserved* guesses. The strong responses to private information benefit information aggregation as observational learning behavior is informationally more efficient in the *observed* sequence than in standard equilibrium outcomes. Once the incentives to make the empirically optimal guess are large enough, *observed* and *unobserved* behave quite similarly, and *observed* guesses are significantly less informative if fewer successors can benefit from the revelation of private signals. These findings are well in line with the qualitative predictions of an observational learning model where players have altruistic preferences.

Keywords: altruism, observational learning, experiments, herd behavior, overconfidence.

JEL Classification: C72, C92, D83.

1 Introduction

A substantial amount of evidence has accumulated in recent years indicating that, when they have limited private information and take public actions, people learn from observing the actions of others. Observational learning effects have been identified in settings as diverse as markets for experience goods like restaurant meals (Cai, Chen, and Fang, 2009) or movies (Moretti, 2011), U.S. presidential primaries (Knight and Schiff, 2010), agricultural technologies (Conley and Udry, 2010) and financial markets (Cipriani and Guarino, 2014). However, the more people act on the public information contained in others' actions, the less informative their own actions are. Information cascade models offer stylized illustrations

*For valuable comments and discussions we are grateful to Thomas Daske, Dirk Engelmann, Sebastian Krügel, Michael Kurschilgen, Birendra Rai, Marco Sahm, and Georg Weizsäcker. Helpful comments were also received from seminar audiences at Berlin, Belfast, Munich, Nottingham, and Strasbourg, as well as from conference participants at the 2015 EEA annual meeting and the 2015 Thurgau Experimental Economics Meeting. We thank Jörg Cyriax, Tobias Gschnaidtner, Florian Inderst, Leonard Przybilla, Alexander Schlimm, and Laura von Lekow for excellent research assistance. Finally, the first author gratefully acknowledges financial support from the European Research Council and hospitality of Paris School of Economics during his time as a postdoctoral researcher.

[†]TUM Munich University of Technology, Department of Business and Economics, Arcisstraße 21, 80333 Munich, Germany. Email: christoph.march@wi.tum.de.

[‡]Corresponding author: Queen's Management School, Queen's University Belfast, 185 Stranmillis Road, Belfast BT9 5EE, United Kingdom. Email: email@anthonyziegelmeyer.com.

of this self-defeating property of observational learning (Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992). In these models, as soon as public information supports one action by just enough to outweigh any one person’s private information, people rationally take uninformative imitative actions, implying a spectacular failure of information aggregation and significant chances of wrong herds. The key insight is that, by taking advantage of the public pool of information, rational people fail to exploit their private information in a socially optimal way.

If, instead of joining the cascade, a person would follow her private signal, and if this deviation would not be viewed as mere noise by others, her informative action would add to the public pool of information and benefit her followers. Such contrarian behavior by a number of people would then strongly increase the probability that actions converge to the truth. Seizing on this logic, several authors assert that dissenters, who are willing to disagree with the majority’s view and disclose their private views, may be surprisingly useful citizens. For example, Bernardo and Welch (2001) argue that entrepreneurs enrich society with extensive exploration of the economic environment as they tend to ignore the herd, and Sunstein (2019) emphasizes the value of judicial dissents to reduce conformity effects within the federal courts in the United States. Thus, in a given observational learning setting, people’s willingness to act altruistically by taking informative actions that are socially beneficial but individually suboptimal should matter greatly for the efficiency of the social outcomes. The present paper investigates, both theoretically and experimentally, the extent to which altruism increases the response to private information and whether it has the potential to improve the efficiency of the observational learning process.

We designed a laboratory experiment to test the causal impact of altruism on observational learning behavior. As in the seminal information cascade experiment (Anderson and Holt, 1997), one of two payoff-relevant states is randomly selected, participants are endowed with binary private signals about the state, they are asked, in sequence, to guess the state, and a monetary reward is paid when the guess is correct. The key innovation of our design is that guesses are made in two parallel sequences. Guesses in the *observed* sequence are revealed publicly whereas guesses in the *unobserved* sequence remain private. Hence, participants in both sequences learn from the same history of *observed* guesses and they may benefit when these guesses are informative while *unobserved* guesses never reveal any information to others. Altruistic participants should therefore respond more strongly to their private signals in the *observed* than in the *unobserved* sequence. We test this prediction by comparing participants’ behavior in the two sequences.

To clarify how altruism affects observational learning, our theoretical section analyzes an extension of Bikhchandani et al.’s (1992) canonical game where players have either standard or altruistic preferences. We begin by showing that sequential equilibria exist where altruistic responses to private information are stronger than standard ones, which in turn delays the onset of information cascades and enhances information aggregation. We then analyze the extended game through the lens of logit quantal response equilibrium. The analysis shows that, even with payoff-responsive decision errors, altruism has the potential to induce more informative observational learning. The section ends with detailed predictions of logit quantal response equilibrium in parametrized versions of the game. We illustrate these predictions by plotting, for each history of public guesses, the probability to contradict the private signal against the expected correctness of contradicting the signal. While standard probabilities to contradict private information lie on a S -shaped curve through $(0.5, 0.5)$, their altruistic counterparts lie below the S -shaped curve, and the distance between altruistic and standard probabilities increases with the level of altruism. Furthermore, the distance is smaller for later guesses than for guesses made early in the sequence and it

is negligible when the expected correctness of contradicting private information is either extremely low or extremely high. The same qualitative predictions hold in parametrized versions of the laboratory game, and these predictions form our testable hypotheses.

The experimental results confirm the main implications of altruistic observational learning. First, in the early periods of the sequences, *observed* guesses are significantly more informative than *unobserved* guesses. Indeed, when plotted against (an estimate of) the expected correctness of contradicting private information, the proportion of guesses that contradict private information is significantly lower in the *observed* than in the *unobserved* sequence. Second, the proportion of *observed* guesses that contradict private information lies significantly closer to a *S*-shaped curve through (0.5, 0.5) in the later than in the early periods of the sequence meaning that *observed* guesses are significantly less informative if fewer successors can benefit from the revelation of private signals. Third, we measure the amount of information aggregated by *observed* guesses and we find that they aggregate significantly more information than standard equilibrium guesses.

Beside altruism, however, other forces may drive stronger than standard responses to private information. As pointed out in the literature review below, various observational learning experiments have documented that participants are reluctant to contradict their private signals and have explained this behavior in terms of judgment or inferential biases, though in most of these experiments altruism remains an alternative explanation. Our experiment enables us to properly explore this auxiliary question since the absence of future informational benefits of guesses in the *unobserved* sequence implies that deviations from standard equilibrium predictions cannot be caused by altruistic considerations. We find that the proportion of *unobserved* guesses that contradict private information lie significantly below a *S*-shaped curve through (0.5, 0.5). Though no one benefits from their informative guesses, *unobserved* put too much weight on their private information relative to public information. In particular, they fail to learn successfully from public guesses by following their private signals more often than not when the monetary incentives to contradict private information are weakly stronger than the monetary incentives to follow private information.

Finally, we investigate whether observational learning behavior changes as participants gain experience playing the laboratory game. In the *observed* sequence, the reluctance of participants to contradict their private signals slightly increases and their guesses aggregate significantly more information over time. By contrast, the proportion of *unobserved* guesses that contradict private information lie significantly closer to a *S*-shaped curve over the course of the experiment. Thus, experience weakens the influence of judgment or inferential biases on observational learning behavior and altruism is the main reason why the guesses of experienced participants are more informative than standard guesses.

Related Literature. We review below the literature related to our paper, which helps us to further clarify its contributions.

First, our theoretical section relates to the few models of observational learning which allow for payoff interdependence. Smith, Sørensen, and Tian (2017) study the altruistic observational learning model with a general distribution of private information and an infinite number of players, and their analysis focuses on the constrained efficient equilibrium.¹ They show that with private information of bounded

¹Vives (1997) also presents a welfare analysis of observational learning, but in a market setting with Gaussian information. The author shows that, compared to the rational benchmark, players' response to private information is stronger and more public information is accumulated in the socially optimal benchmark.

strength higher degrees of altruism lead to smaller cascade sets which entails that uninformative actions are taken less readily. We prove related results in a simple model of observational learning with a finite sequence of altruistic players who receive symmetric binary private signals. In particular, Proposition 2 shows that equilibria exist in which the onset of information cascades is delayed and the likelihood of a herd on the less profitable action is reduced.

Ali and Kartik (2012) consider a simple setting of observational learning with *collective* preferences: a player’s payoff depends on a binary state of nature and on the profile of actions of any subset of all players, players may differ in how they care about the choices of others, and each player weakly prefers others to take the most profitable action.² They show that an equilibrium exists in which players behave as in the unique equilibrium of the standard model where payoffs are independent. By focusing on the specific case of homogeneous altruism in a similar observational learning setting, we provide detailed results about the impact of forward-looking incentives on players’ response to private information. Additionally, we show that, even with payoff-responsive decision errors, altruism has the potential to induce more informative observational learning.

Second, by shedding new light on the drivers of participants’ overweighting of private information, our experimental results complement the existing evidence on observational learning. Alongside the seminal work of Anderson and Holt (1997)—henceforth AH—, several experiments on observational learning report frequent deviations from the standard equilibrium in which participants follow their private signal but should herd on the guesses of their predecessors (e.g. Nöth and Weber, 2003; Çelen and Kariv, 2004; Goeree, Palfrey, Rogers, and McKelvey, 2007). Also, there is a positive correlation between the length and strength of laboratory cascades since participants are more likely to make uninformative imitative guesses after longer cascades (Kübler and Weizsäcker, 2005; Ziegelmeyer, Koessler, Bracht, and Winter, 2010). Combining data from 13 experiments based on Bikhchandani et al.’s (1992) canonical game, Weizsäcker (2010) shows that rational expectations together with standard responses to these expectations cannot account for these behavioral regularities. Altruistic observational learning, however, is a sound explanation for participants’ overemphasis on private information and its attenuation as laboratory cascades lengthen.

Indeed, participants are likely to understand that a stronger response to private information is an altruistic act as participants have collectively (almost) full information. In addition, the monetary cost for a participant who ignores a short laboratory cascade is rather trivial while the monetary benefits for her successors are potentially substantial, and there is ample evidence of altruistic behavior in laboratory games where own costs are low while others’ benefits are large (Andreoni, Harbaugh, and Vesterlund, 2008). But once the amount of information contained in preceding guesses is sufficiently large, a cascade-breaking and informative guess is rather costly and of little benefits which explains that participants are more willing to follow longer sequences of identical guesses. Despite its ability to account for the overweighting-of-private-signals phenomenon, the experimental literature has largely ignored the implications of altruistic observational learning.³ The reluctance of participants to contradict their private

²Thanks to the richness of the class of preferences considered, Ali and Kartik’s observational learning model encompasses many applications. For example, each player could be altruistic toward a different set of other players, i.e., the structure of altruism could be captured by a network. Another leading application is sequential voting with rich motivations for voters such as expressive-voting preferences or preferences about margins of victory.

³Alevy, Haigh, and List (2007) and Duffy, Hopkins, Kornienko, and Ma (2019) are notable exceptions. In the former study, the authors find that financial market professionals respond more strongly to private information and make better

information has been explained as resulting either from a commonly known non-Bayesian updating rule where the precision of private signals is exaggerated (Goeree et al., 2007) or from participants believing that others make more mistakes than they actually do (Nöth and Weber, 2003; Kübler and Weizsäcker, 2004). The experimental evidence presented in this paper shows that altruism and cognitive biases are complementary forces that shape observational learning behavior.⁴

Two more recent experimental papers have studied observational learning in a continuous action space where private signals can be perfectly inferred from rational guesses. The standard equilibrium outcome is therefore informationally efficient which implies that altruistic and standard responses to private information are equally strong. In Eyster, Rabin, and Weizsäcker’s (2018) single-file treatment, each private signal is the difference between the number of heads and tails from 100 coin flips, one participant submits a guess in each period after having observed her private signal and the history of preceding guesses, and participants are incentivized to guess the sum of their predecessors’ signals plus their own signal. The authors find that participants’ guesses are largely in line with the standard equilibrium predictions. In their social learning treatments, Angrisani, Guarino, Jehiel, and Kitagawa (2018) modify AH’s baseline setting by incentivizing participants to state their beliefs rather than to guess which of the two states is payoff-relevant. Their experimental results are at odds with the standard equilibrium predictions as participants tend to assign too much weight to their private signal relative to the public information contained in others’ beliefs. The authors attribute their findings to the overconfidence bias according to which participants tend to mistrust the ability of their predecessors to understand their private signals correctly.

To explain the conflicting results of the two studies, one may argue that participants view the task of adding up numbers as less challenging than the task of forming conditional probabilities and that they trust their predecessors to act rationally when completing the former but not the latter task. The behavior of participants in the *unobserved* sequence can be interpreted as giving some support to this explanation. Indeed, *unobserved* are reluctant to contradict their private information in the first part of the experiment. However, once participants in both sequences have gained extensive experience with the

use of available public information than do students. Consequently, the professionals are involved in weakly fewer overall laboratory cascades and significantly fewer cascades on the wrong action. The authors conclude that “...data reveal that the decisions of market professionals are consistent with behaviors that may mitigate informational externalities in market settings, and thus reduce the severity of price bubbles due to informational cascades.” In the latter study, before guessing the payoff-relevant state, participants choose between receiving a private signal or observing the history of previous guesses and participants are found to choose private information more frequently when they have successors than when they act last. The authors write: “By choosing and acting according to that private information, an individual can help followers by increasing the informativeness of social information. Such group effects could be one explanation for why subjects were found to be biased toward private information in standard social learning experiments relative to equilibrium predictions based on self-interest.”

⁴In their “majority rule institution” treatment, Hung and Plott (2001) incentivize participants to exhibit altruistic behavior by rewarding them according to whether a majority of guesses are right or wrong. They find that more information is revealed in this treatment compared to the usual one where each participant is rewarded according to whether her guess is right or wrong. Davis (2017) proposes an experimental test of Ali and Kartik’s theoretical result that the standard equilibrium outcome remains an equilibrium outcome of the observational learning game where players have collective preferences. To do so, the author modifies participants’ earnings in AH’s baseline setting as follows: the earnings of a participant who makes a correct guess increase with the number of correct guesses in the sequence whereas the earnings of a participant who makes an incorrect guess decrease with the number of incorrect guesses in the sequence. Note that this earnings function fails to induce participants to reveal their private information as long as they believe that such a guess is incorrect. Observational learning behavior does not significantly differ between the settings with original and modified earnings.

combination of private and public information, *unobserved* learn rather successfully from public guesses in the second part of the experiment.

The rest of the paper is organized as follows. Section 2 presents our theoretical analysis of a simple model of altruistic observational learning. Section 3 describes our experimental design and procedures, and it presents the hypotheses to be tested. Section 4 reports on the results of our experiment. Section 5 concludes. The online supplementary material contains a series of appendices with proofs, complementary theoretical and data analyses, as well as the experimental instructions of our experiment.

2 Theory

We analyze an extension of the specific observational learning game of Bikhchandani et al. (1992)—henceforth BHW—where players have either standard or altruistic preferences. In the absence of altruism, the equilibrium outcome exhibits a spectacular failure of information aggregation as the onset of information cascades is almost immediate. This corresponds to BHW’s standard equilibrium outcome. In the presence of altruism, we show that sequential equilibria exist where players increase their response to private information which leads to information cascades being delayed and more information being aggregated. We also establish that logit quantal response equilibria exist where altruism induces more informative observational learning.

2.1 A Simple Game of Altruistic Observational Learning

The finite set of players is $\{1, \dots, T\}$ with generic element t . Nature moves first and chooses a payoff-relevant state $\theta \in \Theta = \{\mathcal{B}, \mathcal{O}\}$ according to the common prior $p \equiv \Pr(\theta = \mathcal{B}) \in (0.5, 1)$.⁵ Each player is then endowed with a symmetric binary private signal $s_t \in S = \{b, o\}$ such that $\Pr(s_t = b \mid \theta = \mathcal{B}) = \Pr(o \mid \mathcal{O}) = 1 - \Pr(o \mid \mathcal{B}) = 1 - \Pr(b \mid \mathcal{O}) = q \in (p, 1)$. Conditional on the state, signals are independently distributed across players.

Time is discrete and, in period $t = 1, 2, \dots, T$, player t chooses action $x_t \in X = \{B, O\}$ where B stands for “guess state \mathcal{B} ” and O stands for “guess state \mathcal{O} ”. Before choosing her action, player t observes the history of previous actions $\mathbf{h}_t = (x_1, \dots, x_{t-1}) \in H_t = \{B, O\}^{t-1}$ where $\mathbf{h}_1 \equiv \emptyset$ and $H \equiv \bigcup_{t=1}^T H_t$.

Player t ’s preferences depend on the complete profile of actions $\mathbf{x} = (x_1, \dots, x_T)$ and the state θ , and they are represented by the von-Neumann Morgenstern utility function

$$u_t(\mathbf{x}, \theta) = \pi(x_t, \theta) + \alpha \sum_{\tau \neq t} \pi(x_\tau, \theta) \quad (1)$$

where $\pi(B, \mathcal{B}) = \pi(O, \mathcal{O}) = 1$, $\pi(B, \mathcal{O}) = \pi(O, \mathcal{B}) = 0$, $0 \leq \alpha \leq 1$ captures the level of altruism, and $\tau \in \{1, \dots, T\}$. If $\alpha > 0$ then player t values all players guessing the true state correctly. This concern for socially efficient outcomes implies that altruistic players take into account the informational benefits of their actions for others. Our functional form of prosocial preferences relates to the (utilitarian) social welfare function assumed by Smith et al. (2017) in their general welfare analysis of observational learning.

⁵We abstract from the non-generic case $p = 1/2$ to avoid the use of tie-breaking rules.

2.2 Beliefs, Strategies and Equilibrium Concepts

Denote by $\langle T, X, H, \Theta, p, S, q, \{u_t\}_{t=1}^T \rangle$ the simple game of altruistic observational learning (henceforth the AOL game). Without loss of generality, player t 's behavior is captured by the *behavioral strategy* $\sigma_t : S \times H_t \rightarrow \mathcal{D}(X)$ where (in a slight abuse of notation) $\sigma_t(s_t, \mathbf{h}_t)$ denotes the probability that she chooses action $x_t = B$.⁶ Call a behavioral strategy pure (completely mixed) if $\sigma_t(s_t, \mathbf{h}_t) \in \{0, 1\}$ ($0 < \sigma_t(s_t, \mathbf{h}_t) < 1$) for each $s_t \in S$ and $\mathbf{h}_t \in H_t$.

Player t forms her probability estimate of the payoff-relevant state, her *belief*, by updating the common prior p with her private and public information (signal s_t and history \mathbf{h}_t) in a Bayesian way. Let player t 's belief be given by the mapping $\mu_t : S \times H_t \rightarrow \mathcal{D}(\Theta)$ and let $\mu_t(s_t, \mathbf{h}_t)$ denote the probability player t assigns to state θ at history \mathbf{h}_t given signal s_t . Given her signal s_t , history \mathbf{h}_t , and the strategies $\sigma_{-t} = (\sigma_1, \dots, \sigma_{t-1}, \sigma_{t+1}, \dots, \sigma_T)$ of the other players, player t 's expected utility of choosing action x_t is given by

$$U_t(x_t | s_t, \mathbf{h}_t, \sigma_{-t}) = \sum_{\theta \in \Theta} \mu_t(\theta | s_t, \mathbf{h}_t) * \left[\pi(x_t, \theta) + \alpha \sum_{\tau < t} \pi(x_\tau, \theta) + \alpha C_t(x_t | \mathbf{h}_t, \theta, \sigma_{-t}) \right] \quad (2)$$

where

$$C_t(x_t | \mathbf{h}_t, \theta, \sigma_{-t}) = \sum_{(x_{t+1}, \dots, x_T)} \left[\prod_{\tau > t} \sum_{s_\tau \in S} \Pr(s_\tau | \theta) \sigma_\tau(x_\tau | s_\tau, \mathbf{h}_\tau) \right] \sum_{\tau > t} \pi(x_\tau, \theta)$$

with $\mathbf{h}_\tau \supseteq (\mathbf{h}_t, x_t)$ for each $\tau > t$. $C_t(x_t | \mathbf{h}_t, \theta, \sigma_{-t})$ is player t 's *continuation value* of action x_t at history \mathbf{h}_t and state θ given strategies σ_{-t} . Lemma A1 in Appendix A shows that continuation values can be defined recursively.

Since they care directly about the actions of other players, altruistic players are strategically motivated to be forward-looking. Forward-looking incentives render the complete characterization of all equilibria of the AOL game infeasible. However, the objective of our analysis is merely to show that, in the presence of altruism, the onset of cascades might be delayed and more information might be aggregated. To achieve this objective, we impose several restrictions on off-path beliefs. As a first restriction, we rely on the sequential equilibrium concept (Kreps and Wilson, 1982) to solve the AOL game, rather than on the perfect Bayesian equilibrium concept.⁷

Definition 1. A sequential equilibrium of the AOL game is a strategy profile σ^* and a system of beliefs μ^* such that

(i) strategies are sequentially rational, i.e., for each t , s_t , and \mathbf{h}_t ,

$$\sigma_t^*(s_t, \mathbf{h}_t) = \begin{cases} 1 & \text{if } U_t(B | s_t, \mathbf{h}_t, \sigma_{-t}^*) > U_t(O | s_t, \mathbf{h}_t, \sigma_{-t}^*) \\ 0 & \text{if } U_t(B | s_t, \mathbf{h}_t, \sigma_{-t}^*) < U_t(O | s_t, \mathbf{h}_t, \sigma_{-t}^*) \end{cases} ; \text{ and}$$

(ii) beliefs are consistent, i.e., $(\sigma^*, \mu^*) = \lim_{n \rightarrow \infty} (\sigma^{(n)}, \mu^{(n)})$ where, for each n , $\sigma^{(n)}$ is a profile of completely mixed behavioral strategies and $\mu^{(n)}$ is derived from $\sigma^{(n)}$ by Bayes rule:

$$\mu_t^{(n)}(s_t, \mathbf{h}_t) = \left[1 + \frac{1-p}{p} \frac{\Pr(s_t | \mathcal{O})}{\Pr(s_t | \mathcal{B})} \prod_{\tau < t} \frac{\sum_{s_\tau \in S} \Pr(s_\tau | \mathcal{O}) \sigma_\tau^{(n)}(x_\tau | s_\tau, \mathbf{h}_\tau)}{\sum_{s_\tau \in S} \Pr(s_\tau | \mathcal{B}) \sigma_\tau^{(n)}(x_\tau | s_\tau, \mathbf{h}_\tau)} \right]^{-1}$$

for each t , s_t , and \mathbf{h}_t where $\mathbf{h}_\tau \subset \mathbf{h}_t$ for each $\tau < t$.

⁶For a given (finite) set M , $\mathcal{D}(M)$ denotes the set of probability distributions over M .

⁷Sequential equilibrium beliefs satisfy $\frac{1-q}{q} \frac{\mu_t(s_t, \mathbf{h}_t)}{1-\mu_t(s_t, \mathbf{h}_t)} \leq \frac{\mu_t(s_t, (\mathbf{h}_t, x_t))}{1-\mu_t(s_t, (\mathbf{h}_t, x_t))} \leq \frac{q}{1-q} \frac{\mu_t(s_t, \mathbf{h}_t)}{1-\mu_t(s_t, \mathbf{h}_t)}$ for each $t < T$, \mathbf{h}_t , x_t and s_t .

The AOL game involves an informational externality as actions partially convey private signals, which might give rise to informational cascades. Given history \mathbf{h}_t , player t *cascades* on her predecessor's action if she acts uninformatively and imitates the action of her predecessor, i.e., $\sigma_t(b, \mathbf{h}_t) = \sigma_t(o, \mathbf{h}_t) = x_{t-1}$ (player t acts informatively if $\sigma_t(b, \mathbf{h}_t) \neq \sigma_t(o, \mathbf{h}_t)$). An *informational cascade* emerges after some history \mathbf{h}_t if every player $\tau \geq t$ cascades on x_{t-1} . Finally, a *herd* on action x_{t-1} emerges after some history \mathbf{h}_t if every player $\tau \geq t$ chooses action x_{t-1} .

Further restrictions on off-path beliefs and monotonic equilibria

We pin down off-path beliefs further by focusing on two special cases. First, with *error off-path beliefs* players treat all actions off the equilibrium path as uninformative about the state. Hence, player t 's belief given signal s_t at the off-path-history \mathbf{h}_t is equal to the belief of player $\tau < t$ given signal $s_\tau = s_t$ at the maximal subhistory $\mathbf{h}_\tau \subset \mathbf{h}_t$ that is on the equilibrium path. Second, with *signal revealing off-path beliefs* players treat off the equilibrium path action B (resp. action O) as revealing signal b (resp. signal o). Accordingly, player t 's belief given signal s_t at the off-path-history \mathbf{h}_t satisfies $\mu_t(s_t, \mathbf{h}_t) / [1 - \mu_t(s_t, \mathbf{h}_t)] = \mu_\tau(s_t, \mathbf{h}_\tau) / [1 - \mu_\tau(s_t, \mathbf{h}_\tau)] \cdot (q / (1 - q))^{n_B - n_O}$ where $\mathbf{h}_\tau \subset \mathbf{h}_t$ is the maximal subhistory of \mathbf{h}_t that is on the equilibrium path and n_B (resp. n_O) is the number of times action B (resp. action O) is chosen by the subset of players $\{\tau, \dots, t-1\}$. Whenever it applies, the off-path beliefs specification is assumed commonly known.

By restricting the analysis to pure strategies and either error or signal revealing off-path beliefs, we are able to capture the behavior of players by the simplified strategies $\hat{\sigma}_t(s_t, \Delta_t)$ where $\Delta_t \in \mathbb{Z}$ denotes the difference between the number of b and o signals that player t infers from history \mathbf{h}_t (see Lemma A3 in Appendix A). For the sake of clarity, our main analysis focuses on *monotonic* equilibria where players adopt such simplified strategies and which require that: i) strategies are weakly increasing in Δ ; and ii) the *information cascade set* weakly grows over time (no information cascade emerges as long as the public belief, the probability estimate of the state conditional on the current history, lies in an interval around one-half; the complement of that interval is called the information cascade set).

Definition 2. An equilibrium $\hat{\sigma}^*$ is **monotonic** if and only if

- (i) for each $t = 2, \dots, T$, each $\Delta_t \in \{2 - t, \dots, t - 1\}$ and each $s_t \in S$, $\hat{\sigma}_t^*(s_t, \Delta_t) \geq \hat{\sigma}_t^*(s_t, \Delta_t - 1)$; and
- (ii) for each $t < T$ and each $\Delta_t \in \{1 - t, \dots, t - 1\}$, $\hat{\sigma}_t^*(b, \Delta_t) \geq \hat{\sigma}_{t+1}^*(b, \Delta_t)$ and $\hat{\sigma}_t^*(o, \Delta_t) \leq \hat{\sigma}_{t+1}^*(o, \Delta_t)$.

The following properties hold in every monotonic equilibrium (see Lemma A4 in Appendix A): (i) players are weakly more likely to choose action B with signal b than with signal o ; (ii) players act informatively whenever $\Delta \in \{-1, 0\}$; (iii) for each $t = 1, \dots, T$, $\hat{\sigma}_t^*(b, \Delta_t) = 0$ only if $\Delta_t \leq -2$ and $\hat{\sigma}_t^*(o, \Delta_t) = 1$ only if $\Delta_t \geq 1$; (iv) if players cascade on action B when having inferred difference Δ then they also cascade on action B when the difference is $\Delta + 1$; and similarly (v) if players cascade on action O when having inferred difference Δ then they also cascade on action O when the difference is $\Delta - 1$.

2.3 Onset of Cascades and Information Aggregation in Monotonic Equilibria

There are multiple equilibrium outcomes in the AOL game. In the absence of altruism ($\alpha = 0$), the equilibrium outcome involves the spectacular failure of information aggregation as players cascade on action B (resp. action O) as soon as there is an imbalance of one B action (resp. two O actions) in

the history of previous actions. We refer to the informationally inefficient equilibrium outcome as the standard equilibrium outcome. In the presence of altruism ($0 < \alpha \leq 1$), equilibrium outcomes exist where cascades are delayed and the informativeness of public information is enhanced.

2.3.1 Immediate Cascades and Poor Information Aggregation

Ali and Kartik (2012) extend BHW's specific observational learning game by considering state-dependent signal precisions $\{q_B, q_O\}$ such that $q_B + q_O > 1$ and a general class of utility functions where each player has a type which identifies the strength of her preference for others to guess the true state correctly. They term this kind of payoff interdependence *collective preferences*. Ali and Kartik prove that the standard equilibrium outcome remains an equilibrium outcome of the game where players have *collective preferences* and the *error off-path beliefs* specification is assumed. Appendix A.1 shows that their game encompasses the AOL game which enables us to state the following proposition.

Proposition 1 (Ali and Kartik, 2012). *For any $0 < \alpha \leq 1$ there exists a monotonic equilibrium $\hat{\sigma}^*$ which for each $1 \leq t \leq T$ and each $\Delta_t \in \mathbb{Z}$ satisfies*

$$(\hat{\sigma}_t^*(b, \Delta_t), \hat{\sigma}_t^*(o, \Delta_t)) = \begin{cases} (1, 1) & \text{if } \Delta_t \geq 1 \\ (1, 0) & \text{if } -1 \leq \Delta_t \leq 0 \\ (0, 0) & \text{if } \Delta_t \leq -2 \end{cases}.$$

The equilibrium characterization of Ali and Kartik (2012) establishes that the core insights from standard economic models of observational learning, such as the swift onset of information cascades and their inherent fragility, can be relevant even when players are motivated by efficiency concerns.

2.3.2 Delayed Cascades and Improved Information Aggregation

For sufficiently long sequences of players, Proposition 2 characterizes a set of monotonic equilibria where compared to the standard equilibrium the onset of information cascades is delayed and more public information is accumulated.

Proposition 2. *Assume that the sequence of players is long enough so that the lower bound $0 < \underline{\alpha}(p, q) < 1$ exists.⁸ For each $\alpha > \underline{\alpha}(p, q)$ there exists a monotonic equilibrium $\hat{\sigma}^*$ which for each $1 \leq t \leq T$ and each $1 - t \leq \Delta_t \leq t - 1$ satisfies*

$$(\hat{\sigma}_t^*(b, \Delta_t), \hat{\sigma}_t^*(o, \Delta_t)) = \begin{cases} (1, 1) & \text{if } \Delta_t \geq \bar{\Delta}_t \\ (1, 0) & \text{if } \underline{\Delta}_t + 1 \leq \Delta_t \leq \bar{\Delta}_t - 1 \\ (0, 0) & \text{if } \Delta_t \leq \underline{\Delta}_t \end{cases}$$

with $\bar{\Delta}_t \geq 1$ and $\underline{\Delta}_t \leq -2$ for each $1 \leq t \leq T$, and either $\bar{\Delta}_t \geq 2$ or both $\bar{\Delta}_t \geq 2$ and $\underline{\Delta}_t \leq -3$ for some $t < T$. Moreover, $\underline{\Delta}_t \leq \underline{\Delta}_{t+1}$ and $\bar{\Delta}_t \geq \bar{\Delta}_{t+1}$ for each $1 \leq t < T$ with strict inequality for some $t < T$, i.e., the cascade set weakly grows over time.

In non-standard equilibria, players act informatively only if $|\Delta|$ is sufficiently small and sufficiently many successors can benefit from the revelation of their private information. For example, early players in the sequence reveal their private information when facing $\Delta \in \{-2, -1, 0, 1\}$ as long as several other players

⁸The restriction on T is weak. For instance, $T = 5$ is sufficient as long as $p < 0.74$ or $q < 0.87$, and $T = 6$ is sufficient as long as $p < 0.81$ or $q < 0.94$. See Appendix A.4 for details.

succeed them. On the other hand, players cascade if $|\Delta|$ is large or if they act late in the sequence. By inducing players to rely more on their private information when choosing their actions, altruistic observational learning accumulates more public information and it enhances the expected correctness of subsequent players' guesses. Said differently, the onset of information cascades is delayed and the likelihood of a herd on the ex-post wrong action is reduced. For non-negligible levels of altruism, such equilibrium play is intuitively more plausible as informative actions reflect players' concern for socially efficient outcomes.

Since our main objective has been to show that altruism may delay the onset of cascades and aggregate more information, we focused on monotonic equilibria. As expected, the combination of forward-looking incentives and information externalities implies the existence of many non-monotonic equilibria in the AOL game. Appendix B exhibits some of these additional equilibria and offers evidence on how rapidly the number of equilibria grows with the level of altruism.

2.4 Quantal Response Altruistic Observational Learning

We complement our previous theoretical analysis by investigating the behavioral implications of the homogeneous Logit Quantal Response Equilibrium (LQRE) in the AOL game. The motivation for this complementary investigation is twofold. First, there is no off-path behavior in LQRE as smoothed best responses ensure that all paths are reached with positive probability. We can therefore dispense with assuming specific off-path beliefs. Second, and most important, payoff-responsive decision errors alter the predictions of rational observational learning in a way similar to that of altruism. Indeed, the standard LQRE predicts that a herd-breaking action happens more frequently if the player received a private signal contradicting the herd choices and the herd is short which implies a positive relationship between the length and strength of herds as well as full information aggregation in the limit (Goeree, Palfrey, and Rogers, 2006). Thus, there is a legitimate concern that altruism has negligible influence on the response to private information and information aggregation in the presence of payoff-responsive decision errors. Our next proposition shows that this concern is unwarranted.

Let μ_t^\varnothing be shorthand for the public belief $\mu_t(\varnothing, \mathbf{h}_t) \in (0, 1)$ with $\mathbf{h}_t \in H_t$ and $1 \leq t \leq T$, let $\sigma_t^Q(s_t, \mu_t^\varnothing)$ denote player t 's LQRE probability to choose action B given signal $s_t \in S$ and μ_t^\varnothing , and let $\sigma^{Q_0} = (\sigma_1^{Q_0}, \dots, \sigma_T^{Q_0})$ denote the standard LQRE. In the next proposition we compare the action probabilities in σ^{Q_0} with the action probabilities in LQRE which are “*monotonic-within-periods*” for strictly positive levels of altruism. In a *monotonic-within-periods* LQRE the action probabilities satisfy the following two properties: i) $\sigma_t^Q(b, \mu_t^\varnothing) > \sigma_t^Q(o, \mu_t^\varnothing)$ for each $\mu_t^\varnothing \in (0, 1)$; and ii) $\partial \sigma_t^Q(s_t, \mu_t^\varnothing) / \partial \mu_t^\varnothing > 0$ for each $s_t \in S$ and each $\mu_t^\varnothing \in (0, 1)$. As clarified in Appendix C, large levels of altruism have to be assumed away to guarantee the existence of a *monotonic-within-periods* LQRE.⁹

Proposition 3. *For strictly positive but not excessively large levels of altruism, there exists a monotonic-within-periods LQRE σ^Q such that, for each $1 \leq t < T$, $\sigma_t^Q(b, \mu_t^\varnothing) > \sigma_t^{Q_0}(b, \mu_t^\varnothing)$ if $\mu_t^\varnothing \in [\underline{\mu}, 1/2)$ and $\sigma_t^Q(o, \mu_t^\varnothing) < \sigma_t^{Q_0}(o, \mu_t^\varnothing)$ if $\mu_t^\varnothing \in (1/2, \bar{\mu}]$ where $0 < \underline{\mu} < 1 - q$ and $1 > \bar{\mu} > q$.*

⁹Extensive numerical computations show that a *monotonic-within-periods* LQRE may fail to exist when the level of altruism exceeds one-half. The longer the sequence of players T and the higher the signal quality q the more likely is the restriction $\alpha < 0.5$ to be binding. Intuitively, if the second component of the utility function weights much larger than the first component—players mainly care that others guess the true state correctly—then the (almost) unique purpose of actions is to reveal private signals and the two properties of a *monotonic-within-periods* LQRE may not hold.

Proposition 3 shows that a LQRE exists where informative actions are more likely than in the standard LQRE if the public belief is not too distant from one-half. We therefore confirm that altruism has the potential to induce more informative observational learning even in the presence of payoff-responsive decision errors. Still, our characterization of the LQRE is only partial since we do not provide a closed-form expression of the interval $[\underline{\mu}, \bar{\mu}]$ for which informative actions are more likely in the presence of altruism. In particular, a characterization of the size of the interval *across* periods is unavailable though, obviously, $\sigma_T^Q(s_T, \mu_T^\emptyset) = \sigma_T^{Q_0}(s_T, \mu_T^\emptyset)$ for each $s_T \in S$ and each $\mu_T^\emptyset \in (0, 1)$. Note that if decision errors are negligible then $0 < \underline{\mu} \leq (1-q)^2 / (q^2 + (1-q)^2)$ and $1 > \bar{\mu} \geq q^2 / (q^2 + (1-q)^2)$ which implies that altruism increases the probability of informative actions when players face herds of size ≤ 2 .

To illustrate how the level of altruism impacts LQRE play in the AOL game, we plot LQRE responses to the value of contradicting private information in parametrized versions of the game. The value of contradicting private information, denoted by $vcpi$, is the probability of guessing correctly the true state when the guess contradicts the private signal. Thus, in guessing situation (s_t, \mathbf{h}_t) , $vcpi$ is $\Pr(\theta = \mathcal{B} \mid o, \mathbf{h}_t)$ and $\Pr(\theta = \mathcal{O} \mid b, \mathbf{h}_t)$ for $s_t = o$ and $s_t = b$, respectively. Players' responses to $vcpi$ are captured by their probabilities to contradict private information. For example, in guessing situation (o, \mathbf{h}_t) with public belief μ_t^\emptyset , player t 's response is $\sigma_t^Q(o, \mu_t^\emptyset)$ and $\sigma_t^{Q_0}(o, \mu_t^\emptyset)$ for $\alpha > 0$ and $\alpha = 0$, respectively. LQRE predictions are numerically computed for $p = 0.505$, $q = 0.7$, $T \in \{7, 12\}$, $\alpha \in \{0, 1/8, 1/4\}$ and logit precision parameter $\lambda = 5$ (Appendix C.3 outlines the algorithm used to derive these predictions). The upper and lower panel of Figure 1 plots the LQRE responses to $vcpi$ for $T = 7$ and $T = 12$, respectively. We distinguish between the early periods ($t \leq 3$) and the later periods ($t \geq 4$) of the games. White, light and dark grey bubbles coincide with the guessing situations of the parametrized games where altruism is absent ($\alpha = 0$), weak ($\alpha = 1/8$) and strong ($\alpha = 1/4$), respectively, and the size of a bubble is proportional to the likelihood with which the situation occurs. To aid visualization, we superimpose fitted curves from weighted linear regressions with a cubic polynomial in the value of contradicting private information.

Since standard LQRE responses are better responses to $vcpi$, white bubbles lie on a S -shaped curve through $(0.5, 0.5)$. Clearly, this observation applies in each panel of Figure 1 as standard LQRE probabilities to contradict private information are identical in the two games (in later periods there are more white bubbles in the lower than in the upper panel because games with longer sequences of players induce more guessing situations). Second, we observe that the grey fitted curves are located below the S -shaped curve through $(0.5, 0.5)$ which indicates that LQRE probabilities to contradict private information decrease in the presence of altruism. The stronger is altruism the more pronounced is the reduction as the dark grey curve ($\alpha = 1/4$) lies below the light grey curve ($\alpha = 1/8$). Third, the impact of altruism is smaller in later than in early periods and it becomes negligible for extreme values of contradicting private information. Thus, Figure 1 illustrates that altruistic LQRE actions are distinctly more informative than standard LQRE actions unless few successors can benefit from the revelation of private signals or $vcpi$ is extreme. This holds true even in the case of weak altruism. Indeed, in the early periods of the lower panel, the light grey curve is clearly below the dashed curve when $vcpi \in [0.3, 0.55]$. Fourth, Figure 1 illustrates that the stronger is altruism the more information is aggregated. For example, in guessing situation $(o, (B, B))$ of the lower panel, $vcpi$ is 0.537, 0.592 and 0.650 when $\alpha = 0, 1/8$ and $1/4$, respectively. We obtain the same qualitative results in parametrized versions of the AOL game where $q \in \{0.6, 0.8\}$ and $\lambda \in \{2.5, 7.5\}$ (in the extreme case where $T = 7$, $q = 0.6$ and $\lambda = 2.5$ the impact of altruism is hardly noticeable; details are available from the authors upon request).

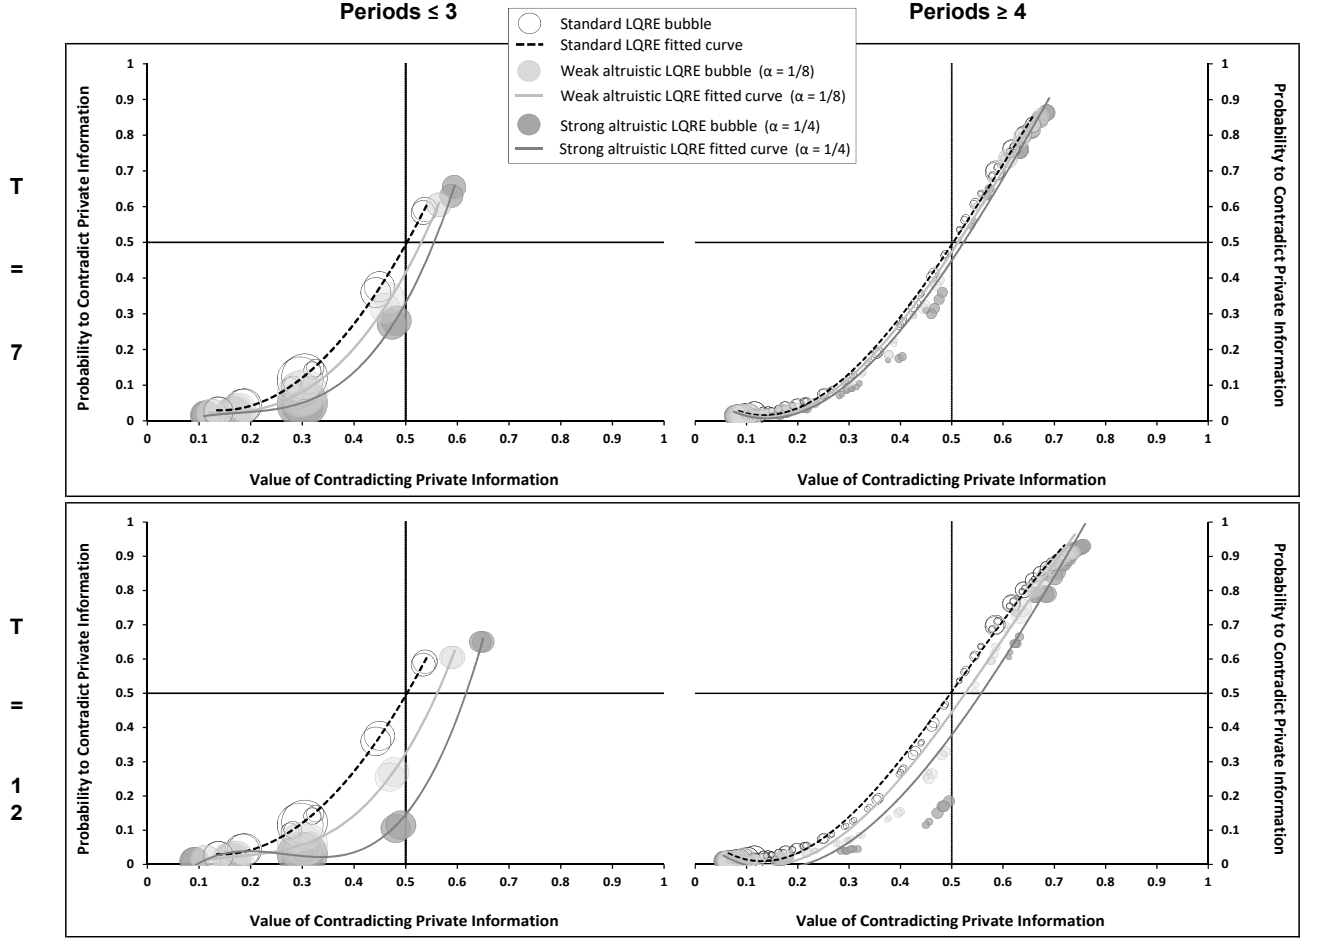


Figure 1: LQRE RESPONSES TO $vcpi$ IN EARLY AND LATER PERIODS
 $(p = 0.505, q = 0.7, T \in \{7, 12\}, \alpha \in \{0, 1/8, 1/4\}, \lambda = 5)$

3 Experimental Design, Procedures and Hypotheses

We designed a laboratory experiment to test the qualitative predictions laid out in the previous section. By comparing the behavior of participants whose actions are revealed publicly to the behavior of participants whose actions are concealed from others, we are able to test the causal impact of altruism on observational learning behavior. Below, we first describe the design and procedures of our experiment and then we present the hypotheses to be tested.

3.1 General Features of the Design

We implement a parameterized version of the observational learning setting described in Section 2.1, but with the essential modification that there are two parallel sequences of participants. Each repetition of the laboratory game begins with the random selection of one of two options and the selected option is not disclosed to participants until all decisions have been made. The two options are labeled ‘blue’ and ‘orange’ with the blue option having a 11/20 probability to be selected and the orange option having a 9/20 probability to be selected. Participants obtain independent private signals that reveal information about which of the two options has been randomly selected. Specifically, if the blue option has been randomly selected then each participant receives a blue signal with probability 2/3 and an orange signal

with probability $1/3$, while if the orange option has been randomly selected then each participant receives an orange signal with probability $2/3$ and a blue signal with probability $1/3$. Then, in each of the two sequences, participants are asked to guess which of the two options has been randomly selected at the beginning of the repetition and they are incentivized to guess correctly.

There are respectively 7 and 8 participants in the *observed* and *unobserved* sequence, and each repetition consists of 8 periods. Once all guesses have been submitted in a given period (but the last one), the guess of one *observed* is made public knowledge and this participant stops guessing. Guesses of *unobserved* remain private. Participants are randomly allocated to either the *observed* or the *unobserved* sequence at the beginning of the experimental session and they keep the same role during the entire session.

3.2 The Progress of a Session

Each experimental session is partitioned into three parts. In the practice part participants become familiar with the game during three of its repetitions and each participant submits only one guess per repetition. In the two non-practice parts many more guesses are collected in each of the six repetitions of the game since participants submit guesses in several situations distinguished only by the history of previous guesses they observe. Following Cipriani and Guarino (2009), this design feature allows participants to gain extensive experience with the combination of private and public information.

Practice part. In the first part of a session the procedures closely follow those used by AH in their baseline experiment except for the two parallel sequences of participants and the fact that guesses are collected and transmitted through computer terminals. Participants draw their private signals from a physical urn (with replacement) and they are randomly assigned to guessing periods. Each participant obtains a single draw from an urn which contains 14 balls indicative of the selected option and 7 balls indicative of the unselected option. Thus, if the blue (respectively orange) option has been randomly selected then the urn contains 14 blue balls and 7 orange balls (respectively 14 orange balls and 7 blue balls). In each of the first seven periods one *observed* and one *unobserved* simultaneously guess an option. The guess submitted in the *observed* sequence is then displayed on all participants' computer screens at the beginning of the next period. In the last period only the remaining *unobserved* submits a guess. From the second period on, participants may condition their guesses on the *observed* guesses submitted in previous periods. Each participant receives 1 Euro for a correct guess and nothing otherwise.

Non-practice parts. Repetitions in the last two parts of a session are identical to those in the practice part except that participants draw private signals from virtual urns displayed on their computer screens and submit multiple guesses. Concretely, each *unobserved* submits 8 guesses and each *observed* submits between 1 and 7 guesses. In the first period, all 15 participants guess one of the two options. The guess of one *observed* is then randomly selected to be made public at the beginning of the next period and this participant stops guessing. In the second period, each of the 14 remaining participants submits a guess. Again, the guess of one *observed* is randomly selected to be made public at the beginning of the next period and this participant stops guessing. This process continues until the last period where each of the *unobserved* submits a guess. For each participant, only one randomly selected guess is paid in each repetition (a correct guess is rewarded 1 Euro whereas an incorrect guess is not rewarded). Each *observed* is paid only for the last guess she submits, i.e., the guess which is made public. Each *unobserved* is randomly assigned to a period at the end of the repetition and paid for the guess made in that period. Exactly one *unobserved* is assigned to any given period.

Feedback screens. Due to the use of a strategy method-like procedure, the feedback that participants receive at the end of each repetition in the non-practice parts differs from the feedback they receive in the practice part. Each participant is reminded of her draw, her guessing period, the guess she made, the sequence of *observed* guesses, and she is informed about the selected option, the composition of the urn used and her earnings in the practice part. Feedback screens in the non-practice parts are identical to those in the practice part except that each participant is reminded only of the payoff-relevant guess she made and of the sequence of *observed* guesses which were made public. Note that a participant is never informed of the draws made by the other participants.

3.3 Experimental Procedures

Experimental sessions took place at the laboratory for experimental research in economics of the Technical University of Munich (*experimentTUM*) in July 2016. Students from the Technical University of Munich and the Ludwig-Maximilians University of Munich were invited using the ORSEE recruitment system (Greiner, 2015). We conducted nine sessions with sixteen participants in each session. One participant was randomly selected to serve as the laboratory assistant and the remaining participants were randomly assigned to computer terminals. The experiment was programmed in zTree (Fischbacher, 2007).

Each session started with short demonstrations of the option-selection procedure to small groups of participants. An experimenter shuffled a deck of 20 cards – 11 cards with a blue front and 9 cards with an orange front – and laid the cards face down on a table. The assistant then picked 1 card out of the 20 cards, and the front color of the picked card determined the randomly selected option.¹⁰ After the demonstrations, paper instructions for the practice part were distributed and participants were given time to read them at their own pace. Instructions were then read aloud and finally participants learned about their role (*observed* or *unobserved*).

Once the three repetitions of the practice part were over, paper instructions for (non-practice) part 1 were distributed and participants were given time to read them at their own pace. A summary of the instructions was then read aloud. The paper instructions were followed by a short on-screen-demonstration of the draws from the virtual urns. Again, one of the experimenters summarized aloud the main points of the demonstration. After that, the six repetitions of part 1 were run.

Part 1 was followed by a short break. Participants were offered soft drinks and water, and a paper questionnaire was distributed asking for gender, month and year of birth, academic major, mother tongue, and citizenship. Short paper instructions for (non-practice) part 2 were then distributed and the six repetitions of part 2 were conducted. Finally, participants privately retrieved their earnings.

In each session we collected 45 guesses from the three repetitions of the practice part and 552 guesses from the six repetitions of each following part. We collected a total of 3,213 *observed* and 7,128 *unobserved* guesses. On average, a participant in the role of *observed* earned 12.92 Euro whereas a participant in the role of *unobserved* earned 13.21 Euro, including a show-up fee of 3 Euro. A session lasted for about 85 minutes. In all parts of a session, participants only interacted through the computers and no other communication was permitted. Appendix E contains a translated version of the experimental instructions.

¹⁰The laboratory assistant randomly selected the option in each repetition of the game. The assistant also helped with the drawing of signal realizations from the physical urns in the practice part and she monitored the progress of the session on her own computer terminal.

3.4 Hypotheses

In our experiment, *observed* and *unobserved* face the same cognitive challenge of combining private and public information. They learn from the same history of public guesses after having been endowed with a private signal of identical quality. However, the possibility for participants in both sequences to benefit from informative guesses may induce *observed* to increase their response to private information. *Unobserved* guesses, on the other hand, never reveal any information to others. Since the only concern of *unobserved* is to guess the selected option correctly, they should act like non-altruistic players.

We derive our hypotheses about the impact of altruism on observational learning behavior from LQRE predictions in parametrized versions of the laboratory game. Allowing for decision errors in the considered benchmarks is appropriate as participants' behavior is inherently error prone and the standard LQRE fairly captures systematic deviations from Nash equilibrium in information cascade experiments, most notably the correlation between the length and strength of herds (Kübler and Weizsäcker, 2005; Ziegelmeyer et al., 2010). To account for the full range of behavioral predictions, we rely on LQRE responses to *vcpi*—the latter being simply the probability to receive €1 when the guess contradicts the private signal. LQRE predictions are numerically computed for $\alpha = 0$ in the *unobserved* sequence, $\alpha \in \{1/8, 1/4\}$ in the *observed* sequence, and $\lambda = 5$.¹¹ Based on these predictions, we construct plots of LQRE responses to *vcpi* in the early and later periods of the two sequences. These plots resemble those of Figure 1 that illustrate the predictions of LQRE in the AOL game. There is, however, a distinctive feature of the laboratory game to consider, namely that an informative public guess is potentially beneficial to successors in the *observed* sequence as well as to *unobserved*. Arguably, *observed* may value the welfare of their successors more strongly than they value the welfare of *unobserved* as they could identify more strongly with the former than with the latter.¹² We therefore compute two sets of predictions in the *observed* sequence where the ratio between the levels of altruism towards *unobserved* and *observed* is respectively 1/2 and 1. Below, our hypotheses describe the shapes of LQRE responses to *vcpi* in the two sequences (in the interests of brevity, the plots of these responses are relegated to Appendix C.4).

The first hypothesis is that participants in the *unobserved* sequence better respond to *vcpi*.

HYPOTHESIS 1. In the early and later periods of the sequence, *unobserved* responses to *vcpi* lie on a *S*-shaped curve through (0.5, 0.5).

For conciseness, we simply refer to the *S*-shaped curve on which the *unobserved* responses to *vcpi* lie as the *unobserved S*-shaped curve. The second hypothesis is that altruism causes guesses in the early periods of the *observed* sequence to be more informative than standard LQRE guesses.

HYPOTHESIS 2. In the early periods of the sequence, *observed* responses to *vcpi* lie below the *unobserved S*-shaped curve. The closer is *vcpi* to the lower or upper limit of its range, the smaller is the distance between these responses and the curve.

We expect a larger distance between the *observed* responses to *vcpi* and the *unobserved S*-shaped

¹¹Clearly, $p = 0.55$, $q = 2/3$, and there are respectively 7 and 8 players in the *observed* and *unobserved* sequence. The chosen value for the logit precision parameter is reasonably close to the values estimated in information cascade experiments.

¹²Chen and Li (2009) measure the effects of induced group identity on social preferences. They report that participants are significantly more likely to choose social-welfare-maximizing actions when matched with an ingroup member.

curve the more altruistic participants are or the more strongly *observed* identify with *unobserved*. The third hypothesis is that participants are less likely to submit informative guesses in the later than in the early periods of the *observed* sequence.

HYPOTHESIS 3. *Observed* responses to *vcpi* lie closer to the *unobserved* *S*-shaped curve in the later than in the early periods of the sequence.

As the three hypotheses make clear, our experiment mainly aims at ascertaining whether guesses are more informative when successors can reap these informational benefits than when they cannot and, if so, whether increases in participants' responses to private information are qualitatively captured by altruistic observational learning. However, we did not design the experiment to estimate participants' level of altruism or to determine the extent to which *observed* identify with successors in their sequence compared to successors in the *unobserved* sequence.

4 Results

First, we summarize the aggregate properties of our data. Second, we test the qualitative predictions of altruistic observational learning and we investigate whether participants' behavior changes as the session progresses. Third, we examine the informational efficiency of observational learning behavior. Note that our data analysis excludes the few guesses submitted during the practice part.

4.1 Descriptive Statistics

We first report the histories of public guesses in the different decision periods—that is, the histories of *observed* guesses that have been publicly revealed up to the (beginning of the) relevant period. Then, we assess the influence of public guesses on participants' propensity to contradict their private information.

Table 1 shows the distribution of public histories in each period inferred from the 108 repetitions of the laboratory game. To ease presentation, we shorten the notation of histories—e.g., histories *BBBB* and *OOOB* are shortened to *4B* and *3OB*—and from period 5 on we only report histories which occur at least 3 times.

Period 2	B 55%					O 45%						
Period 3	2B 33%		BO 22%			OB 18%			2O 27%			
Period 4	3B 25%	2BO 8%	BOB 12%		B2O 10%	O2B 9%	OBO 9%	2OB 3%	3O 24%			
Period 5	4B 24%	2BOB 4%	2B2O 4%	BO2B 8%	BOBO 4%	B3O 8%	O3B 5%	O2BO 4%	OBOB 6%	OBO2O 3%	4O 22%	
Period 6	5B 24%	2BO2B 4%	BO3B 8%		BOBOB 4%	B4O 8%	O4B 5%	O2BOB 3%	OBO2B 5%	4OB 3%		5O 19%
Period 7	6B 22%	2BO3B 4%	BO4B 7%			B5O 7%	O5B 5%	OBO3B 4%			4OBO 3%	6O 19%
Period 8	7B 20%	2BO4B 3%	BO5B 7%			B6O 7%	O6B 5%	OBO4B 4%			7O 19%	

Table 1. DISTRIBUTIONS OF PUBLIC HISTORIES

We observe that many histories start with *BO* (slightly more than a fifth of the histories materialized in period 3 or later) and that some other histories start either with *BBO* or *OOB* (together, slightly more than one tenth of the histories materialized in period 4 or later). These off-the-Nash-equilibrium-path guesses imply that empirical histories are more diverse than predicted. In particular, Bayesian rationality predicts that 75% of the final histories are full cascades, i.e., *7B* or *7O*, and that about 20% of final histories are either *O6B* or *OB5O*. The predicted distribution of final histories differs significantly from the empirical one (Chi-square test; p -value < 0.01) which has only 39% of its final histories that are full laboratory cascades. Note that these results are perfectly in line with those of past cascade experiments.

We now compare participants' propensity to contradict private information in the two sequences. In a given period, the information set of a participant is the couple (private signal, public guesses in preceding periods). As a convention, we denote the size of the majority of public guesses by $sm = \#blue - \#orange$ where $\#blue$ ($\#orange$) is the number of blue (orange) guesses in the public history and $sm \in \{-(t-1), \dots, t-1\}$ at the beginning of period $t \in \{1, \dots, 8\}$. We refer to the majority of public guesses as a contrary majority (resp. favoring majority) in cases where the participant's private signal and the majority of public guesses are conflicting (resp. concordant) pieces of information. Thus, participants face a contrary majority either when endowed with a blue signal and $sm < 0$ or when endowed with an orange signal and $sm > 0$. On the other hand, participants face a favoring majority either when endowed with a blue signal and $sm > 0$ or when endowed with an orange signal and $sm < 0$. If $sm = 0$, there is no majority in the history of public guesses. Table 2 reports the percentage of guesses that contradict private information by the signal of each role and for various majorities of public guesses. We show the percentages of guesses that contradict private information across the two non-practice parts as well as separately for each part. Note that *observed* and *unobserved* face majorities of size at most 6 and 7. We don't differentiate between favoring majorities since the percentages of guesses contradicting private information hardly change with the size of the favoring majority neither do we differentiate between large contrary majorities as fewer data are available for contrary majorities of size 5 or more.

Several observations can be made from Table 2. First, almost all guesses follow private information at favoring and no majorities, even more so in the *observed* than in the *unobserved* sequence. Second, the propensity to contradict private information increases (almost) systematically with the size of the contrary majority and few guesses follow private information once the public evidence is conclusive enough (contrary majorities of size 4 or more). Third, guesses account to some extent for the asymmetric prior as participants contradict more frequently their orange than their blue signals at contrary majorities of size 1 and the difference vanishes at larger contrary majorities. These three observations illustrate that participants have some success at learning from public guesses. Still, a fourth observation is that the propensity to contradict private information is low at contrary majorities of size 1 with an orange signal and it remains modest at contrary majorities of size 2 with either a blue or an orange signal. More importantly, *unobserved* tend to contradict their private information more often than *observed* at these short contrary majorities. Observations 1 and 4 therefore indicate that responses to private information are stronger in the *observed* than in the *unobserved* sequence unless participants face large contrary majorities. Finally, the behavioral differences between the two sequences tend to be more pronounced in the second than in the first non-practice part.

Type of the majority		Both parts				Part 1				Part 2			
		<i>Observed</i>		<i>Unobserved</i>		<i>Observed</i>		<i>Unobserved</i>		<i>Observed</i>		<i>Unobserved</i>	
		<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>	<i>b</i>	<i>o</i>
Favoring majority (of any size)		02% (551)	04% (492)	05% (1,604)	06% (1,372)	03% (291)	04% (252)	07% (733)	07% (704)	00% (260)	04% (240)	04% (871)	04% (668)
No majority		03% (524)	06% (528)	05% (688)	11% (768)	04% (273)	10% (245)	04% (323)	12% (381)	02% (251)	03% (283)	06% (365)	10% (387)
Contrary majority of size	1	14% (185)	26% (247)	25% (307)	38% (446)	11% (90)	30% (108)	17% (133)	36% (227)	16% (095)	24% (139)	31% (174)	40% (219)
	2	64% (086)	64% (120)	73% (180)	68% (247)	67% (039)	71% (052)	67% (069)	78% (128)	62% (047)	59% (068)	77% (111)	56% (119)
	3	84% (062)	91% (074)	87% (183)	83% (216)	82% (034)	95% (037)	84% (082)	89% (122)	86% (028)	86% (037)	90% (101)	76% (094)
	4	88% (043)	90% (048)	95% (133)	89% (167)	80% (025)	93% (028)	93% (075)	89% (109)	100% (018)	85% (020)	97% (058)	88% (058)
	≥ 5	100% (024)	85% (040)	92% (265)	93% (336)	100% (015)	87% (023)	89% (137)	93% (233)	100% (009)	82% (017)	96% (128)	92% (103)

Note: In each cell, the first row reports the percentage of guesses that contradict private information, either the blue signal **b** or the orange signal **o**, and the second row reports in brackets the number of guesses.

Table 2. PERCENTAGES OF GUESSES THAT CONTRADICT PRIVATE INFORMATION

4.2 Responses to the Empirical Value of Contradicting Private Information

To test our three behavioral hypotheses, we analyze *observed* and *unobserved* responses to the empirical value of contradicting private information, which is an estimate of $vcpi$ obtained using the procedure introduced in Weizsäcker (2010) and refined by Ziegelmeyer, March, and Krügel (2013). For each guessing situation, the empirical value of contradicting private information, denoted by \widehat{vcpi} , is the probability to receive €1 when the guess contradicts the private signal estimated across all observations with the same history and private signal. As the number of occurrences of the guessing situation increases in the dataset, the empirical value of contradicting private information approaches $vcpi$ (Appendix D.1 details the derivation of \widehat{vcpi}). For example, averaged across histories with no majority—including the empty history in period 1, \widehat{vcpi} equals 0.286 and 0.384 when the signal is respectively blue and orange. Also, as expected, it increases when histories induce contrary majorities: Averaged across observations where the size of the contrary majority is 1 (2), \widehat{vcpi} equals 0.441 and 0.533 (0.592 and 0.615) when the signal is respectively blue and orange.

Figure 2 plots the proportions of *observed* and *unobserved* guesses that contradict private information against \widehat{vcpi} where the latter is estimated separately in parts 1 and 2 as the number of occurrences of a given situation differs in the two parts. The left and right panel shows participants' responses to \widehat{vcpi} respectively in early periods (periods 1-3) and in later periods (up to period 7 for *observed* and up to period 8 for *unobserved*).¹³ The abscissae of bubbles are given by levels of \widehat{vcpi} and the size of a bubble reflects the number of occurrences of the situation. The ordinates of gray and white bubbles are given

¹³In the alternative decomposition where the early periods comprise period 4, later periods contain only 10% of the observations compared to almost 30% in the current decomposition.

by the proportions of contradictions for *observed* and *unobserved*, respectively. Moreover, each bubble corresponds to a guessing situation which occurs at least 10 times as \widehat{vcpi} is likely to be far away from $vcpi$ for rarely occurring situations. There are 101 distinct guessing situations depicted in Figure 2 for a total of 6,222 individual observations.

Figure 2 also superimposes fitted curves from a weighted linear regression that includes a cubic polynomial in \widehat{vcpi} fully interacted with indicator variables for *unobserved* and early periods. To correct for the fact that \widehat{vcpi} imperfectly measures $vcpi$, we follow the split-sample instrumental variable (IV) method described in Weizsäcker (2010) which obtains an instrument by partitioning the dataset in two subsamples. The gray and dotted black curve is the fitted curve for *observed* and *unobserved*, respectively. Appendix D describes the split-sample IV method, it reports the regression results, and it also contains robustness checks with different subsets of data. In all instances we find the same qualitative results.

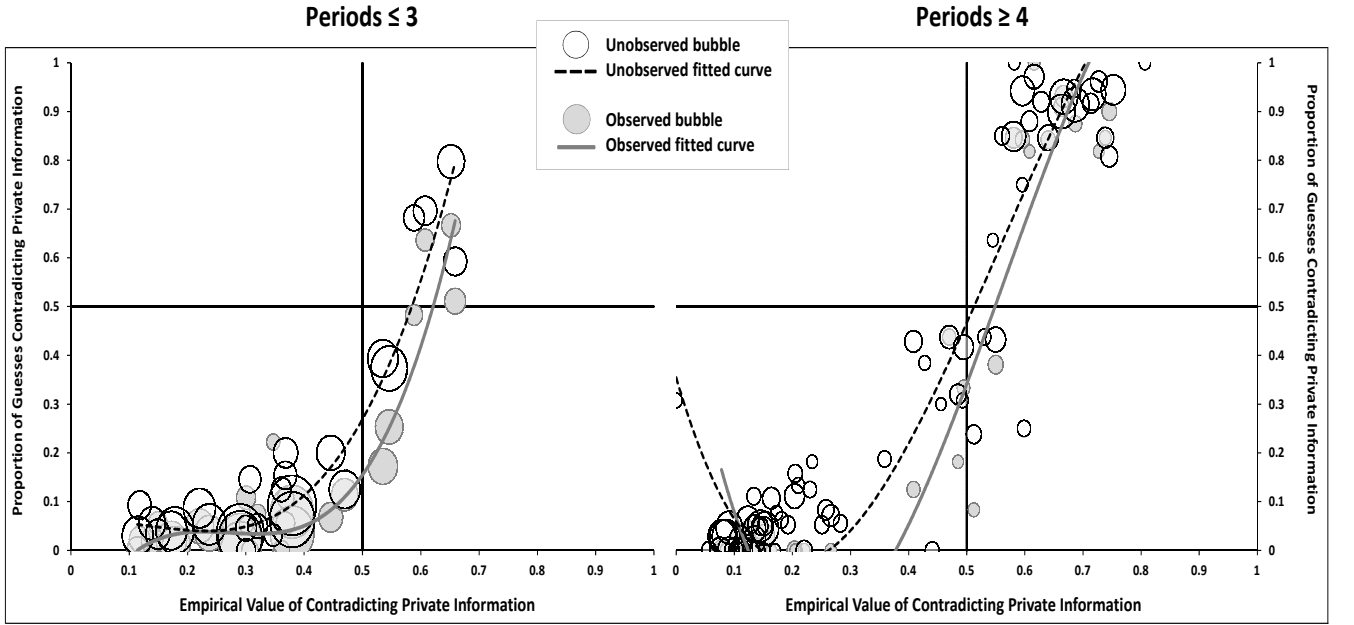


Figure 2: PARTICIPANTS' RESPONSES TO \widehat{vcpi} IN EARLY AND LATER PERIODS

In the remainder of this subsection, we first analyze *observed* and *unobserved* responses to \widehat{vcpi} in early and later periods of the two sequences. We then explore whether there are behavioral changes in our experiment by investigating whether these responses differ in the two non-practice parts.

Responses to the empirical value of contradicting private information in early periods

In situations where private information happens to support the empirically optimal guess *unobserved* largely follow their signal. Averaged across observations where $\widehat{vcpi} \leq 0.5$, the proportion of *unobserved* guesses that are optimal is 0.929. In the more challenging situations where they should contradict their private information, however, *unobserved* are much less likely to guess optimally. Averaged across observations where $\widehat{vcpi} > 0.5$, the proportion of *unobserved* guesses that are optimal is only 0.542.¹⁴

¹⁴Incentives to act optimally are stronger in the left than in the right half of the panel as $\widehat{vcpi} \in]0.1, 0.7[$. Nevertheless, the proportion of *unobserved* guesses that are optimal in the left half hardly changes when considering the same incentive levels as in the right half: Across observations where $0.3 \leq \widehat{vcpi} \leq 0.5$, the average proportion of *unobserved* guesses in line

Compared to standard LQRE predictions, *unobserved* responses to \widehat{vcpi} are too strong when the latter is slightly smaller than one-half and they are too weak when \widehat{vcpi} is slightly larger than one-half. Averaged across period 2-observations where *unobserved* face a contrary majority of size 1, the proportion of contradictions is 0.159 and 0.382 at $\widehat{vcpi} = 0.459$ and 0.541 with a blue and an orange signal, respectively. The fitted curve for *unobserved* goes through (0.5, 0.269) and (0.585, 0.5), and we reject our first hypothesis as the vertical distance between the dotted black curve and (0.5, 0.5) is strongly significant (two-tailed p -value = 6.41×10^{-6}).¹⁵ Of particular importance is the harmful reluctance of *unobserved* to contradict their private information when the monetary incentives to follow others are weakly stronger than the monetary incentives to follow the private signal ($0.50 < \widehat{vcpi} \leq 0.55$). Mistrust in the ability of *observed* to understand their private signals correctly (Angrisani et al., 2018) or biases in statistical reasoning (Goeree et al., 2007) may cause *unobserved* to learn unsuccessfully from others in these situations.

Still, once the monetary incentives to herd are strong enough, *unobserved* largely contradict their private information. Averaged across period 3-observations where *unobserved* face a contrary majority of size 2, the proportion of contradictions is 0.690 and 0.711 at $\widehat{vcpi} = 0.599$ and 0.655 with a blue and an orange signal, respectively.

RESULT 1. In the early periods of the sequence, *unobserved* responses to \widehat{vcpi} lie significantly below a S -shaped curve through (0.5, 0.5), contrary to the standard LQRE predictions. In particular, *unobserved* follow their private signal more often than not in situations where it is empirically optimal to contradict private information and the monetary incentives to do so are weak ($0.50 < \widehat{vcpi} \leq 0.55$).

Result 1 implies that, contrary to what our second hypothesis states, showing that *observed* guesses in early periods deviate significantly from standard LQRE guesses is no conclusive evidence that altruism significantly impacts observational learning behavior. To validate this prediction, we need to establish that, in the early periods of the two sequences, *observed* guesses are significantly more informative than *unobserved* guesses.

The left panel of Figure 2 shows that the fitted curve for *observed* lies below the one for *unobserved* over the entire range of \widehat{vcpi} . The gray curve goes through (0.5, 0.155) and (0.621, 0.5), and the proportion of *observed* contradictions is respectively 0.038, 0.093 and 0.379 when averaged across observations where $\widehat{vcpi} \in]0.1, 0.4]$, $]0.4, 0.5]$ and $]0.5, 0.7[$. Since the vertical distance between the dotted black curve and the gray curve at $\widehat{vcpi} = 0.5$ is strongly significant (one-tailed p -value = 0.009), we conclude that *observed* guesses are significantly more informative than *unobserved* guesses in early periods. Also in line with altruistic LQRE predictions, we find that *observed* guess more similarly to *unobserved* when \widehat{vcpi} is closer to the lower or upper limit of its range. Indeed, the difference between the proportions of *unobserved* and *observed* contradictions is respectively 0.024 and 0.035 when $\widehat{vcpi} \in]0.1, 0.3]$ and $]0.3, 0.5]$. Likewise, the difference between the proportions of *unobserved* and *observed* contradictions is respectively 0.179 and 0.110 when $\widehat{vcpi} \in [0.5, 0.6[$ and $[0.6, 0.7[$.

with private information is 0.897.

¹⁵Our regressions use robust standard errors that are (conservatively) clustered at the session level, and we apply a finite-cluster correction when computing the p -values of our hypothesis tests. See Appendix D for details.

RESULT 2. In line with the altruistic LQRE predictions, *observed* responses to \widehat{vcpi} lie significantly below those of *unobserved* in early periods, and the distance between these responses is smaller when \widehat{vcpi} is closer to the lower or upper limit of its range.

Responses to the empirical value of contradicting private information in later periods

The right panel of Figure 2 shows that *unobserved* responses to \widehat{vcpi} are comparable to those predicted by standard LQRE in later periods of the sequence. Indeed, the proportion of *unobserved* contradictions is respectively 0.046, 0.349, 0.684 and 0.920 when averaged across observations where $\widehat{vcpi} \in [0, 0.4[$, $[0.4, 0.5[$, $[0.5, 0.6[$ and $[0.6, 0.8]$. We find that the vertical distance between the dotted black curve and (0.5, 0.5) does not differ significantly from zero (two-tailed p -value = 0.504).¹⁶

As predicted by altruistic LQRE, *observed* guesses are less informative in later than in early periods. The gray curve goes through (0.5, 0.341) and (0.549, 0.5), and the proportion of *observed* contradictions is respectively 0.013, 0.276 and 0.793 when averaged across observations where $\widehat{vcpi} \in [0, 0.4[$, $[0.4, 0.5[$ and $[0.5, 0.8]$. We reject the null hypothesis that the vertical distance between (0.5, 0.5) and the fitted curve for *observed* is the same in early and later periods (one-tailed p -value = 0.013).

RESULT 3. In line with the altruistic LQRE predictions, *observed* responses to \widehat{vcpi} lie significantly closer to a S -shaped curve through (0.5, 0.5) in the later than in the early periods of the sequence.

4.2.1 Behavioral Dynamics

Thanks to its two non-practice parts, our experiment allows us to study whether observational learning behavior changes as participants gain experience with the guessing task. Of particular interest is the harmful reluctance of *unobserved* to contradict their private signals when the monetary incentives to do so are weakly stronger than the monetary incentives to follow private signals ($0.50 < \widehat{vcpi} \leq 0.55$). We therefore investigate whether *unobserved* responses to \widehat{vcpi} become more consistent with those predicted by standard LQRE as the session progresses. Additionally, we explore whether altruistic observational learning persists among participants well experienced with its payoff consequences by comparing *observed* responses to \widehat{vcpi} in parts 1 and 2.

Figure 3 plots the proportions of *observed* and *unobserved* contradictions against \widehat{vcpi} separately for the first and the second non-practice part. We superimpose fitted curves from a weighted linear regression that includes a cubic polynomial in \widehat{vcpi} fully interacted with indicator variables for *unobserved* and part 2. As in Figure 2, each bubble corresponds to a guessing situation which occurs at least 10 times and the regression uses the split-sample IV method to correct for the fact that \widehat{vcpi} imperfectly measures $vcpi$.

Unobserved adjust markedly their behavior as they gain experience with the guessing task. Indeed, over time, their responses to $\widehat{vcpi} \in [0.45, 0.50[$ weaken (the average proportion of contradictions is respectively 0.214 and 0.373 in part 1 and 2) while their responses to $\widehat{vcpi} \in]0.50, 0.55]$ strengthen (the average proportion of contradictions is respectively 0.355 and 0.424 in part 1 and 2). However, averaged across observations where $\widehat{vcpi} \in [0, 0.45[$ and $]0.55, 0.8]$, the proportions of *unobserved* contradictions are very similar in the two parts (0.063 and 0.861 in part 1 versus 0.064 and 0.849 in part 2). Consequently,

¹⁶Among the levels of \widehat{vcpi} , the one derived for guessing situation $(o, (B, O, B, O, B, O))$ is clearly an outlier. We estimate that there is no value of contradicting private information in this situation ($\widehat{vcpi} = 0$) while in 4 of the 13 occurrences of the situation *unobserved* contradict their orange signal when facing the no majority.

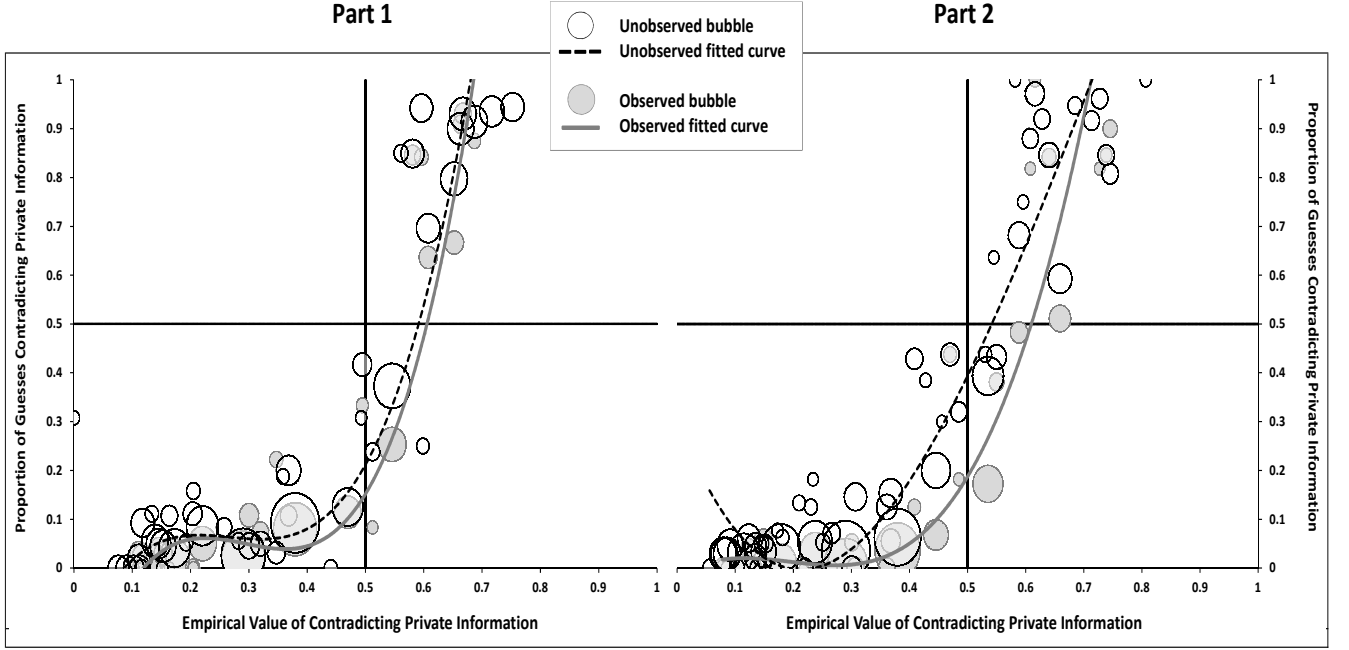


Figure 3: PARTICIPANTS' RESPONSES TO $\hat{v}cpi$ IN EACH PART

the fitted curve for *unobserved* is more comparable to a *S*-shaped curve through (0.5, 0.5) in part 2 than in part 1: The dotted black curve for *unobserved* goes through (0.5, 0.212) and (0.591, 0.5) in the left panel, and it goes through (0.5, 0.394) and (0.541, 0.5) in the right panel. Since the vertical distance between the dotted black curve and (0.5, 0.5) is significantly smaller in part 2 than in part 1 (two-tailed p -value = 0.004), we conclude that *unobserved* guesses are significantly more consistent with standard LQRE guesses in the second than in the first non-practice part.

Observed also adjust their behavior as the session progresses, though to a lower extent than *unobserved* and in the opposite direction. Indeed, *observed* slightly increase their responses to private information as they gain experience with the laboratory game. When averaged across observations where $\hat{v}cpi \in [0, 0.4[$, $[0.4, 0.5[$, $[0.5, 0.6[$ and $[0.6, 0.8]$, the proportion of *observed* contradictions is respectively 0.047, 0.153, 0.434 and 0.779 in part 1, and it is respectively 0.020, 0.147, 0.270 and 0.750 in part 2. Still, the gray curve hardly differs in the two panels, and we cannot reject the null hypothesis that the vertical distance between the fitted curve for *observed* and (0.5, 0.5) is the same in the two parts (two-tailed p -value = 0.596).

RESULT 4. *Unobserved* guesses are significantly more consistent with standard LQRE guesses in the second than in the first non-practice part. *Observed* guesses are slightly more informative in the second than in the first non-practice part, though the difference is not statistically significant.

Note that repeated game considerations do not provide a convincing explanation for the prevalence of information-revealing guesses in the *observed* sequence as the latter are more likely in part 2 than in part 1. For further insight on this matter, we ran a probit regression on the probability for the *observed* guess to be informative where the independent variable is a dummy variable which equals 0 for early rounds (rounds 1-3) and 1 for later rounds (rounds 4-6). We find that the probability of an informative guess is higher in later than in early rounds, which confirms that the behavior of *observed* is not driven by

repeated game considerations (we obtain the same finding in the regression restricted to the observations with $\widehat{vcpi} \in [0.5, 0.6]$; details are available from the authors upon request).

4.3 Information Aggregation

In this subsection we investigate whether, as predicted by altruistic observational learning, the strong responses to private information in the *observed* sequence benefit information aggregation. We assess the informational efficiency of observational learning behavior by comparing the amount of information aggregated in large majorities of *observed* guesses to the amount of information aggregated in standard equilibrium majorities. Also, since *observed* guesses become more informative as the session progresses, we examine whether large majorities aggregate more information in the second than in the first non-practice part.

The empirical value of contradicting private information is a natural measure of the information aggregated by a sequence of guesses. The more information guesses aggregate the lower the levels of \widehat{vcpi} at large favoring majorities and the higher the levels of \widehat{vcpi} at large contrary majorities. For example, when averaged over the blue and orange signals, a sequence of guesses contains no valuable information if $\widehat{vcpi} = 1/3$, a favoring majority aggregates one (resp. two) private signal(s) if $\widehat{vcpi} = 1/5$ (resp. $1/9$), and a contrary majority aggregates one (resp. two) private signal(s) if $\widehat{vcpi} = 1/2$ (resp. $2/3$).

We regress \widehat{vcpi} on an indicator variable for part 2 fully interacted with indicator variables for the type of majority. We distinguish between large favoring majorities, moderate majorities, and large contrary majorities where the size of a large majority belongs to $\{3, 4, 5, 6\}$. For conciseness, the two signals are bundled together. We use an OLS specification with robust standard errors clustered at the session level and we include every guessing situation which occurs at least 10 times in the *observed* sequence. Table 3 shows the predicted levels of \widehat{vcpi} by non-practice part and type of majorities. Appendix D.3 reports the regression results as well as robustness checks where the analysis is restricted to subsamples with a more precise measurement of \widehat{vcpi} and where the size of a large majority belongs to $\{4, 5, 6\}$. In all instances we find the same qualitative results.

	Part 1	Part 2
Large Favoring Majorities	0.129 (0.118, 0.140)	0.104 (0.093, 0.115)
Moderate Majorities	0.336 (0.327, 0.345)	0.350 (0.331, 0.370)
Large Contrary Majorities	0.637 (0.624, 0.650)	0.693 (0.664, 0.722)

Note: Every guessing situation which occurs at least 10 times in the *observed* sequence is included for a total of 2,541 individual observations. 95% robust confidence interval in brackets, clustered at the session level and constructed using the delta method.

Table 3. PREDICTED LEVELS OF \widehat{vcpi} IN THE *Observed* SEQUENCE

The regression results confirm the predictions of altruistic observational learning. First, in the standard equilibrium outcome, $vcpi = 51/89 \approx 0.573$ at any contrary majority of size larger than 2 and $vcpi = 19/121 \approx 0.157$ at any favoring majority of size larger than 2. Table 3 shows that in each non-practice part

large majorities of *observed* guesses aggregate significantly more information than any equilibrium history. Furthermore, large majorities of *observed* guesses in part 2 aggregate significantly more information than any standard LQRE history since $vcpi \in [0.124, 0.639]$ for every equilibrium history and every $\lambda \geq 0$ in our laboratory game.

Second, the amount of information aggregated in large majorities increases as the session progresses. At large favoring majorities the predicted level of \widehat{vcpi} is respectively 0.129 and 0.104 in part 1 and 2. The difference in the predicted levels of \widehat{vcpi} between the two parts is strongly significant (one-tailed p -value = 0.004). At large contrary majorities the predicted level of \widehat{vcpi} is respectively 0.637 and 0.693 in part 1 and 2. Once more, the difference in the predicted levels of \widehat{vcpi} between the two parts is strongly significant (one-tailed p -value = 0.001). On the other hand, the predicted levels of \widehat{vcpi} do not differ significantly between the two parts at moderate majorities. These predicted levels indicate that moderate majorities contain, on average, no valuable information, which is to be expected given that they comprise no majorities and about the same number of short contrary and favoring majorities.

RESULT 5. Observational learning behavior is informationally more efficient in the *observed* sequence than in standard equilibrium outcomes. Furthermore, *observed* guesses aggregate significantly more information in the second than in the first non-practice part.

5 Conclusion

The experimental evidence presented in this paper enriches our understanding of how people learn from the actions of others. We compare the behavior of participants who repeatedly play an observational learning game in two parallel sequences. In the *observed* sequence, the actions of participants are revealed publicly. In the *unobserved* sequence, however, the actions of participants remain private. Our results provide three new insights on the drivers of herding behavior.

First and foremost, we show that responses to private information are significantly stronger in the *observed* than in the *unobserved* sequence. Actions are therefore more informative when successors can reap these informational benefits than when they cannot. Still, *observed* and *unobserved* behave quite similarly once the incentives to take the empirically optimal action are large enough, and *observed* actions are significantly less informative if fewer successors can benefit from the revelation of private signals. These findings are well in line with the qualitative predictions of an observational learning model where players have altruistic preferences. As our theoretical section illustrates, even in the presence of payoff-responsive decision errors, altruistic observational learning has the potential to accurately describe participants' behavior in the *observed* sequence.

Second, we find that *observed* do not take less informative actions as they gain experience playing the laboratory game. On the contrary, responses to private information slightly increase over time in the *observed* sequence, though not significantly so. This finding corroborates the observation made in previous information cascade studies that behavioral changes over the course of the experiment are at most weak (Weizsäcker, 2010, footnote 13), and it casts doubt on the ability of repeated game considerations to convincingly explain the prevalence of these information-revealing actions. Based on our estimate of the value of actions, we also find that, as predicted by altruistic observational learning, the strong responses to private information benefit information aggregation. Indeed, observational learning behavior is informationally more efficient in the *observed* sequence than in standard equilibrium outcomes and

observed guesses aggregate significantly more information as the experimental session progresses.

Third, we show that *unobserved* fail to learn successfully from public actions. In particular, they follow their private signals more often than not when the monetary incentives to contradict private information are weakly stronger than the monetary incentives to follow private information. Thus, even when successors cannot benefit from informative actions, participants put too much weight on their private information relative to public information. Angrisani et al. (2018) observe the same suboptimal pattern in a continuous action space experiment where the standard equilibrium outcome is informationally efficient. The authors carefully test the ability of information redundancy neglect and overconfidence to capture participants' behavior, and their results are supportive of the overconfidence bias. Experimental evidence therefore suggests that altruism and overconfidence are two major forces that shape observational learning behavior. That said, we also find that experience weakens the influence of overconfidence on observational learning behavior as *unobserved* become significantly less reluctant to contradict their private information over the course of the experiment. This is the first evidence of strong changes in observational learning behavior over time, and it indicates that altruism shapes the behavior of experienced participants more heavily than overconfidence in our binary action space experiment.

Our findings are important from two perspectives. First, to properly assess the impact of cognitive biases on observational learning behavior, experimentalists need to “turn off” altruism. This can be achieved by implementing a *responsive* observational learning setting where rational actions perfectly reveal beliefs (Ali, 2018) or by concealing participants' actions from others. For example, in March and Ziegelmeyer (2018), we allow for multiple signal qualities in the *unobserved* sequence, and we find that, alongside overconfidence, informational misinferences and non-Bayesian updating undermine participants' success in learning from public actions. Second, future informational benefits of actions are a contextual factor which favors the aggregation of information and heightens efficiency levels relative to rational herding. Concerned by the underinvestment in public information of rational herders, economists have designed mechanisms which release additional public information or incentivize people to reveal their private information (Smith et al., 2017). Our results suggest that cheaper interventions which simply emphasize to people the value of signaling information to their successors might already improve economic welfare.

References

- ALEVY, J. E., M. S. HAIGH, AND J. A. LIST (2007): “Information Cascades: Evidence from a Field Experiment with Financial Market Professionals,” *Journal of Finance*, 62, 151–80.
- ALI, S. N. (2018): “On the Role of Responsiveness in Rational Herds,” *Economics Letters*, 163, 79–82.
- ALI, S. N. AND N. KARTIK (2012): “Herding with Collective Preferences,” *Economic Theory*, 51, 601–626.
- ANDERSON, L. AND C. HOLT (1997): “Information cascades in the laboratory,” *American Economic Review*, 87, 847–862.
- ANDREONI, J., W. T. HARBAUGH, AND L. VESTERLUND (2008): “Altruism in Experiments,” in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, Palgrave Macmillan.

- ANGRISANI, M., A. GUARINO, P. JEHL, AND T. KITAGAWA (2018): “Information Redundancy Neglect versus Overconfidence: A Social Learning Experiment,” Cemmap working paper 63/18.
- BANERJEE, A. V. (1992): “A simple model of herd behavior,” *The Quarterly Journal of Economics*, 107, 797–817.
- BERNARDO, A. AND I. WELCH (2001): “On the Evolution of Overconfidence and Entrepreneurs,” *Journal of Economics & Management Strategy*, 10, 301–30.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 100, 992–1026.
- CAI, H., Y. CHEN, AND H. FANG (2009): “Observational Learning: Evidence from a Randomized Natural Field Experiment,” *American Economic Review*, 99, 864–82.
- ÇELEN, B. AND S. KARIV (2004): “Distinguishing Informational Cascades from Herd Behavior in the Laboratory,” *American Economic Review*, 94, 484–498.
- CHEN, Y. AND S. X. LI (2009): “Group Identity and Social Preferences,” *American Economic Review*, 99, 431–57.
- CIPRIANI, M. AND A. GUARINO (2009): “Herd Behavior in Financial Markets: An Experiment with Financial Market Professionals,” *Journal of the European Economic Association*, 7, 206–233.
- (2014): “Estimating a Structural Model of Herd Behavior in Financial Markets,” *American Economic Review*, 104, 224–51.
- CONLEY, T. G. AND C. R. UDRY (2010): “Learning about a New Technology: Pineapple in Ghana,” *American Economic Review*, 100, 35–69.
- DAVIS, B. J. (2017): “An experiment on behavior in social learning games with collective preferences,” *Economics Letters*, 152, 93–5.
- DUFFY, J., E. HOPKINS, T. KORNIENKO, AND M. MA (2019): “Information Choice in a Social Learning Experiment,” *Games and economic behavior*, *forthcoming*.
- EYSTER, E., M. RABIN, AND G. WEIZSÄCKER (2018): “An Experiment on Social Mislearning,” Working paper.
- FISCHBACHER, U. (2007): “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, 10, 171–8.
- GOEREE, J., T. PALFREY, AND B. ROGERS (2006): “Social learning with private and common values,” *Economic Theory*, 28, 245–64.
- GOEREE, J., T. PALFREY, B. ROGERS, AND R. MCKELVEY (2007): “Self-correcting information cascades,” *Review of Economic Studies*, 74, 733–62.
- GREINER, B. (2015): “Subject pool recruitment procedures: Organizing experiments with ORSEE,” *Journal of the Economic Science Association*, 1, 114–25.

- HUNG, A. A. AND C. R. PLOTT (2001): “Information Cascades: Replication and an Extension to Majority Rule and Conformity–Rewarding Institutions,” *The American Economic Review*, 91, 1508–20.
- KNIGHT, B. AND N. SCHIFF (2010): “Momentum and Social Learning in Presidential Primaries,” *Journal of Political Economy*, 118, 1110–50.
- KREPS, D. M. AND R. WILSON (1982): “Sequential Equilibria,” *Econometrica*, 50, 863–894.
- KÜBLER, D. AND G. WEIZSÄCKER (2004): “Limited depth of reasoning and failure of cascade formation in the laboratory,” *Review of Economic Studies*, 71, 425–441.
- KÜBLER, D. AND G. WEIZSÄCKER (2005): “Are Longer Cascades More Stable?” *Journal of the European Economic Association*, 3, 330–9.
- MARCH, C. AND A. ZIEGELMEYER (2018): “Excessive Herding in the Laboratory: The Role of Intuitive Judgments,” Cesifo working paper series no. 6855.
- MORETTI, E. (2011): “Social Learning and Peer Effects in Consumption: Evidence from Movie Sales,” *Review of Economic Studies*, 78, 356–93.
- NÖTH, M. AND M. WEBER (2003): “Information aggregation with random ordering: Cascades and overconfidence,” *Economic Journal*, 113, 166–89.
- SMITH, L., P. N. SØRENSEN, AND J. TIAN (2017): “Informational Herding, Optimal Experimentation, and Contrarianism,” Working paper.
- SUNSTEIN, C. R. (2019): *Conformity: The Power of Social Influences*, NYU Press.
- VIVES, X. (1997): “Learning from Others: A Welfare Analysis,” *Games and Economic Behavior*, 20, 177–200.
- WEIZSÄCKER, G. (2010): “Do we follow others when we should? A simple test of rational expectations,” *American Economic Review*, 100, 2340–2360.
- ZIEGELMEYER, A., F. KOESSLER, J. BRACHT, AND E. WINTER (2010): “Fragility of information cascades: An experimental study using elicited beliefs,” *Experimental Economics*, 13, 121–145.
- ZIEGELMEYER, A., C. MARCH, AND S. KRÜGEL (2013): ““Do We Follow Others when We Should? A Simple Test of Rational Expectations”: Comment,” *American Economic Review*, 103, 2633–42.