
An Experiment On Social Mislearning

Erik Eyster (London School of Economics and Political Science)
Matthew Rabin (Harvard University)
Georg Weizsäcker (HU Berlin)

Discussion Paper No. 73

February 21, 2018

An Experiment on Social Mislearning

Erik Eyster, Matthew Rabin and Georg Weizsäcker*

February 13, 2018

Abstract

We investigate experimentally whether social learners appreciate the redundancy of information conveyed by their observed predecessors' actions. Each participant observes a private signal and enters an estimate of the sum of all earlier-moving participants' signals plus her own. In a first treatment, participants move single-file and observe all predecessors' entries; Bayesian Nash Equilibrium (BNE) predicts that each participant simply add her signal to her immediate predecessor's entry. Although 75% of participants do so, redundancy neglect by the other 25% generates excess imitation and mild inefficiencies. In a second treatment, participants move four per period; BNE predicts that most players anti-imitate some observed entries. Such anti-imitation occurs in 35% of the most transparent cases, and 16% overall. The remaining redundancy neglect creates dramatic excess imitation and inefficiencies: late-period entries are far too extreme, and on average participants would earn substantially more by ignoring their predecessors altogether. (JEL B49)

Keywords: social learning, redundancy neglect, experiments, higher-order beliefs

*Eyster: Department of Economics, London School of Economics; Rabin: Department of Economics and Business School, Harvard University; Weizsäcker: School of Economics and Business, Humboldt Universität zu Berlin & DIW Berlin. We are grateful for the constructive advice of Tom Palfrey, four anonymous referees, and Joel Sobel acting as co-editor, for comments by audiences at the ESA Meeting 2016 (Tucson), SMYE 2017 (Halle), THEEM 2017, CEU, Cologne, DIW Berlin, Essex, Maastricht, Pittsburgh, and Tel Aviv University, for excellent research assistance by Iuliia Grabova, Kevin He, Signe Moe and Alexander Volkmann, for the support of UCL decision laboratory's staff, and for financial support from the ERC (via Starting Grant 263412) and the German Science Foundation (via CRC TRR 190).

1 Introduction

The theory of how people learn by observing the actions and beliefs of others underlies an extensive and ongoing research program. Beginning with Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992), a literature on observational learning identifies how a rational person who observes the behavior of another person with private information and similar tastes may follow that person, even contrary to her own private information. Yet Eyster and Rabin (2014) show ways that the logic of social inference requires that rational agents greatly limit the scope of their imitation. If the actions a person observes are themselves influenced by social learning, then the person should recognize the redundancy inherent in prior actions and imitate only selectively. Hence extensive imitation is a mistake. Indeed, in most settings full rationality dictates that players should *anti*-imitate some of those they observe, as a way of subtracting off the correlation in the predecessors' beliefs. For example, if financial analysts write daily reports about an asset, and every day they read the previous day's reports in addition to receiving new private information, then if yesterday's reports were all great, a merely good report today signals bad news. Similarly, if multiple customers post online reviews of a restaurant, each of them aggregating new information with previously published reviews, a rational interpretation of reviews requires correcting for the influence of previous reviews.

Accounting for redundancy proves challenging even in settings devoid of rational anti-imitation. Experimental evidence, such as Kübler and Weizsäcker (2004), demonstrates such failure. Even pre-dating the experimental evidence, doubts about whether people fully account for redundancy motivated researchers to develop models of *redundancy neglect*. De-Marzo, Vayanos and Zwiebel (2003), for instance, model the idea that people may treat as independent repeated hearings of the same opinion, and show that this “persuasion bias” generates inefficiency. Eyster and Rabin (2010) and Eyster and Rabin (2014) explore implications of the assumption that people do not fully account for the redundancy in others'

actions, showing that in most settings this can lead to long-run incorrect and overconfident beliefs. By neglecting that those whom they imitate are themselves imitating, people end up being so over-influenced by potentially misleading early actions that long-run beliefs can converge to full confidence on the wrong state.

In this paper, we report on both the behavioral and efficiency properties of social learning in experiments designed to be conducive to efficient learning when rationality is common knowledge, yet amenable to the detection of redundancy neglect when extent. Our experiments are simple and—if it is common knowledge that all participants are strategically sophisticated—do not require the use of Bayes’ rule. In each of two treatments, the participants sum up integers. Each participant privately observes an integer “signal” as well as the public entries of all preceding participants. Her task is to aggregate the information: she makes an entry herself and gets paid for entering a number as close as possible to her “target”, which is equal to her own signal plus those of all earlier-moving participants. Given common knowledge of rationality, participants can recover the sum of their predecessors’ signals from the entries they observe through simple arithmetic. The experiment is designed to highlight the logic of real-world social learning tasks, for example in the interpretation of customer reviews, without demanding complicated math of those who understand the logic.

The lessons learned about rationality and redundancy neglect are varied. We observe very frequent Nash play and relatively mild effects of redundancy neglect in one treatment, and less Nash play and more redundancy neglect, including many extreme behavioral deviations from Nash, in the other treatment. Overall, the frequency of Nash play exceeds that of clear redundancy neglect, but the social consequences of the latter are severe.

In the first treatment—as in previous experiments—participants move single-file, and anti-imitation therefore plays no part of the Bayesian-Nash-Equilibrium (BNE) prediction. Empirically, 75% of participants employ their BNE strategies of simply adding their own signal to the previous entry. Most of the remaining 25% deviate in the direction of redundancy neglect, although some also deviate by ignoring their predecessors. Consequently, the 75%

who do the “right” thing wind up over-imitating; they would be better off by down-weighting the information that they glean from their predecessors. On the other hand, participants benefit from social learning—they do better than they would by ignoring their predecessors—and two notable markers of redundancy neglect that were identified by Eyster and Rabin (2014), overinfluence of initial movers and long-run extreme beliefs, appear neither strongly nor statistically significantly.

In our second treatment, “multi file”, even mild redundancy neglect can be very socially harmful. Here, participants move four at a time, which creates a large set of informational redundancies, causing BNE to predict frequent anti-imitation. Deviations from BNE strategies are far more common than in the single-file treatment. Early signals are heavily over-counted in multi file, pushing later entries towards costly extremities. The majority of games end with beliefs that are ten-fold too distant from the ex-ante belief. On average, participants lose from social learning in multi file—they would do better by ignoring others’ entries entirely and simply announcing their own signals.¹

Our findings complement the contemporaneous evidence of Enke and Zimmermann (2015) that people under-attend to correlation in signals, as well as several related findings in the experimental social-learning literature consistent with redundancy neglect. For example, Weizsäcker (2010) and March and Ziegelmeyer (2015) report on herding experiments in which participants are too prone to follow consensus by their predecessors. The survey of Choi, Gallo and Kariv (2015) reviews evidence that people communicating beliefs in a network employ the DeGroot rule (DeGroot (1974))—taking as their posterior the average belief of their neighbors—which, like redundancy neglect, ignores the observation (or network) structure. Chandresekhar, Larreguy and Xandri (2017) provide experimental evidence in a two-state setting with binary signals and actions that a model in which everyone uses a version of the DeGroot rule modified to their setting fits the data better than any model in

¹Nevertheless, we also uncover evidence of complex rational behavior: almost 16% of moves requiring anti-imitation conform to BNE, which is the first evidence we know of where people rationally anti-imitate.

which some agents are rational and others DeGroot.² Our work differs from these papers in a number of respects, especially in our main finding that people succumb to greater redundancy neglect in an environment that calls for anti-imitation than in one that does not. In addition, our experiments are the first to identify an environment in which the opportunity to learn from their predecessors harms people on average.

Whereas many prior experimental papers use the single-file observation structure, our four-file setting is novel and may appear arbitrary. As Eyster and Rabin (2014) discuss, most observation structures other than the single-file one necessitate some anti-imitation. We chose four-file because it is simple and requires anti-imitation. Real-world analogues of this structure include actions that are taken by many individuals periodically, either simultaneously or without the benefit of observing earlier movers in the same period. Airline passengers, for instance, rate airlines at the end of each flight; students rate courses at the end of each term. Every year, farmers choose how to plant or fertilize their fields, or doctors choose which way to treat their patients. The more realistic the set-up, the less natural it may be to assume common knowledge of the precise information structure—but even so the rational solution will almost surely involve anti-imitation.

We specified few precise statistical tests or hypotheses before running the experiments but predicted that redundancy neglect would lead to over-imitation, and hence work against the BNE prediction of limited imitation (in the single-file treatment) and anti-imitation (in the multi-file treatment). In our descriptive analysis of Sections 3 and 4, we formally investigate the presence of an extreme form of redundancy neglect modeled in Eyster and Rabin (2010), which they call “BRTNI play” (an acronym for “best response trailing naïve inference”), and whose presence we did hypothesize. BRTNI play requires that a player in period t add

²Gale and Kariv (2003) explore Bayesian learning theoretically in social networks in which agents get one dose of private information before repeatedly taking binary actions; after each period, each player observes the actions of her “neighbors”. Choi, Gale and Kariv (2012) estimate Quantal-Response Equilibrium (McKelvey and Palfrey 1995) from experimental data on that model. Golub and Jackson (2010) analyze DeGroot behavior in a closely related environment.

the entries of all players through period $t - 1$ to her own signal. This extreme prediction would quickly generate absurdly high (positive or negative) entries in our setting. Overall, the hit rate of BRTNI is modest. In the single-file treatment BRTNI play is observed much less often than BNE, whereas in the multi-file treatment BRTNI outperforms BNE, yet not by nearly the same margin.

Yet, Eyster and Rabin (2014) give a general definition of redundancy neglect, which encompasses BRTNI play as well as much milder forms of over-counting. In the present context, their general definition of redundancy neglect (roughly) requires that all players' entries equal their signals plus some non-negative linear combination of all previous entries whose coefficients sum to more than one. Whatever its exact form, redundancy neglect predicts that early signals exert undue influence on later moves. This general pattern is weakly confirmed in our first treatment and strongly confirmed in our second treatment.

Section 3 summarizes the results of the first, single-file treatment, where BNE predicts that every player should simply enter her signal plus her immediate predecessor's entry. When done correctly by all players, every action equals its target. For ease of reference, we shall refer to players who play this "naive Bayesian" strategy as *Nebi*. (Nebi is naive by using the BNE strategy even when the available observations make it unlikely or impossible that others are using the BNE strategy. A player whose predecessors are not Nebis does not maximize expected payoff by being Nebi.) As indicated above, the data reveal the presence of far more Nebi than BRTNI players: from $t = 3$ onwards, when the two behavioral rules make different predictions, Nebi outnumbers BRTNI 14:1. The even simpler rule of following one's own signal—the prediction of Eyster and Rabin's (2005) fully-cursed equilibrium as well as Stahl and Wilson (1994), Nagel (1995), and Crawford and Iriberri's (2007) (random) Level-1—appears approximately as frequently as BRTNI. About 88% of decisions accord to one of these three types of behavior. The remaining 12% either follow different rules or appear to make sign errors. In aggregate, these deviations produce over-imitation: 72% of participants who miss their target do so in the direction predicted by BRTNI and other

forms of redundancy neglect. Participants would earn more by adding their signal to a number 31% closer to zero than the number that they apparently use. Yet participants do not over-imitate strongly enough to negate the advantages of following others: they earn 76% of the maximum sum possible, whereas ignoring other players would deliver only 71% of the maximum payment.³

Section 4 summarizes the results in the second, multi-file treatment. Here too, BNE reasoning demands only simple arithmetic of the players, but now involves anti-imitating behavior. In period 1, optimal behavior remains trivial—a participant must just enter her signal. In period 2, she should add her signal to the four entries observed in period 1. But in period 3, because all four period-2 entries incorporate the signals of all four period-1 players, BNE dictates that each player must add her own signal to the sum of period-2 entries *minus* three times the sum of period-1 entries. Without such subtraction, period-3 players would inefficiently quadruple count each period-1 signal. BNE actions in periods 4 to 6 take more complicated and even less intuitive forms, but all involve a mix of adding and subtracting observed entries. Surprisingly, we estimate that participants anti-imitate upwards of 35% of the time in period 3 and 10% in periods 4 to 6. Nonetheless, Nebi play occurs much less frequently in this treatment, and less frequently than BRTNI play during periods for which the two models make different predictions: from $t = 3$ onwards, BRTNI outnumber Nebi 3:2. More generally, behavior corresponds less well to particular rules of thumb in this game. Even including initial periods, BRTNI and Nebi jointly account for only 55% of the data, and including the “cursed” or Level-1 follow-your-own-signal rule only brings explained behavior up to 58% of the data. Nevertheless, the cumulative effect of the non-BNE decisions is strong and clear. 78% of deviations from target go in the direction predicted by BRTNI. Participants make entries with much greater magnitude than those predicted by BNE, and on average would earn more by strongly shading their interpretation of prior entries towards

³Eyster and Rabin (2014) show how even mild over-counting of the sort of observed in this experiment can, when extrapolated to longer time horizons, produce severe long-term costs.

zero—optimally shading by 98%. Moreover, participants earn far less than they would by relying purely on their private signals. The latter result is similar to—albeit more dramatic than—subsequent evidence reported in March and Ziegelmeyer (2015) that experts with very precise private information tend to lose out in situations where a sufficiently large crowd of less-informed players contradicts the experts’ information.

Section 5 reports the results of two sets of regressions testing the BNE predictions. In the first set, we regress participants’ entries in the various periods on earlier signals. BNE predicts that all coefficients should equal one in both treatments, whereas redundancy neglect predicts that players in period t should implicitly weight signals of periods $t-2, t-3, \dots$ with coefficients larger than one. While the regression results in the single-file treatment paints a relatively positive picture of BNE, the results in the multi-file case confirm the presence of redundancy neglect. Nine out of the ten point estimates of these coefficients exceed 1, with estimates ranging from 3.0 to 19.6. In a second set of estimations, we regress entries on past *entries* (and current signals) to uncover players’ strategies. In the multi-file treatment, BNE predicts that players in period t assign negative coefficients to entries in periods $t-2$ and $t-4$. Although we discuss above and below evidence that some individuals engage in anti-imitation, there is no suggestion of average anti-imitation. While two of the six relevant point estimates are negative, they are very far from statistical significance.

Section 6 amends the previous analysis by allowing for probabilistic choice. Like in the the depth-of reasoning analysis of Kübler and Weizsäcker (2004), we estimate two models that encompass both Quantal Response Equilibrium and the Level- k family of models. The results confirm the qualitative features described above: the parameter estimates in multi file indicate lower best-response precision and lower belief in others’ rationality than in single file.

Putting our approach in the perspective of the literature, we note that our experiments differ in at least four ways from the standard experimental set-up developed by Anderson and Holt (1997) and subsequent papers studying herding. First, by giving participants targets

equal to the sums of signals received, our design isolates the redundancy-neglect error as much as possible from both statistical and computational errors. Second, the rich signal and action spaces allow us to finely identify the rules used by most participants; only rare happenstance leads different rules (especially probabilistic rules) to generate identical responses. Third, the rich action and signal spaces preclude the possibility that even rational social learners may fail to learn efficiently, in the sense that social beliefs may be less extreme than they would be if all private information were revealed, which is predicted in the coarser games of Anderson and Holt (1997). The errors of redundancy neglect, highlighted in this paper, give rise to a very different form of inefficiency: social beliefs are more extreme than they would be if all private information were revealed. Fourth, and most important from an economics perspective, we move away from the traditionally studied yet narrow band of herding settings where the BNE strategy closely resembles a less-rational tendency to imitate, and where welfare costs of over-imitation are limited. Our design thereby disentangles rational imitation from irrational over-imitation, improving our understanding of real-world social-learning tasks like interpreting financial reports or customer reviews.⁴

But there are several obvious limitations to our design. First, insofar as people *do* suffer from various statistical biases, neutralizing those biases does not enhance realism. Second, despite the strict BNE being “statistics-free” in this setting, once participants (rightly) start to doubt the common knowledge of rationality, the relative likelihood of signals matters. (Yet the data patterns and our probabilistic models do not suggest that the main departures from BNE predictions can be attributed to statistical issues.) Third, although our game is ‘logically equivalent to’ the task of social inference, that does not mean that we have tested responses to more naturalistic social inference. Perhaps people avoid neglecting redundancy when seeing groups of people reveal their beliefs about best behavior rather than adding

⁴In standard experiments, when redundancy neglect arises, imitating is almost always empirically optimal. The meta analysis of Weizsäcker (2010) confirms that success rates are higher—both with and without controls for incentives—in situations where all previous players agree. (See Table 9 of that paper.)

numbers. Or conversely, perhaps the experiment lays bare the logic of redundancy in a way that no real-world situation would, so that the experimental results under-estimate real-world redundancy neglect. As such, we view this experiment only as a first attempt to move herding experiments towards more realistic observation structures.

Callender and Hörner (2009) describe an altogether different reason for which rational people anti-imitate in the course of observational learning. In coarse-action environments such as the canonical example in which people learn about which of two restaurants A and B is better—choosing restaurants one at a time, and observing all past choices—Person t who observes Person $t - 1$ choose A cannot necessarily infer the intensity of $t - 1$'s belief in A. In this case, Person $t - 2$'s choice can provide additional information about Person $t - 1$'s posteriors, which Person t wishes to adopt as her own priors. Given certain signal structures, this can lead Person t to anti-imitate Person $t - 2$. When people either receive no information or perfect information, the history (A, B) reveals that Player 2 has a fully revealing B signal, whereas (B, B) does not; hence Player 3 is more apt to choose B after the former history than the latter, which amounts to anti-imitation. Callender and Hörner (2009) show that, given this information structure, when people cannot observe the order of their predecessors' moves, it can be optimal to follow the minority choice. Brunner and Goeree (2012) present lab evidence that people eschew such anti-imitation. Because both of our experimental treatments have fine action spaces, Callendar-Hörner-style anti-imitation plays no part in our experiments: our work is orthogonal to that of Callender and Hörner (2009) and Brunner and Goeree (2012).

We conclude the paper in Section 7, where we discuss further how our experiment fits into other research that studies a broader array of herding environments. We reiterate that the multi-file treatment is not designed to provide evidence that BNE fails in a consequential way. (But fail it does. And so too does BRTNI.) Rather, our aim is to shift focus from very special settings where BNE happens to be difficult to distinguish from intuitive imitative behavior. Our data show that consideration of different and seemingly more realistic social-

learning environments may lead to very different conclusions about BNE’s fit as well as about the efficiency of social outcomes. We conclude Section 7 and the paper by discussing how some of the traditional models of limited rationality proposed in the behavioral game theory literature have difficulty accounting for our results.

2 Experimental Design

In each of the two experimental treatments, “single file” and “multi file”, twelve participants interact. In each period $t = 1, \dots, T$, one participant in the single-file treatment, and four participants in the multi-file treatment, receive private information. $T = 24$ in the single-file treatment, and $T = 6$ in the multi-file treatment, in order that each participant in each treatment receive private information exactly twice. Each participant’s information is generated by simulating 100 coin flips that are mutually independent as well as independent of all other random draws in the experiment. The signal of participant i in period t , $s_{i,t}$, comes from the difference between the number of heads and the number of tails of this participant’s current set of coin flips. Upon receiving his or her signal, the participant makes an “entry” $e_{i,t}$, whose payoff

$$\pi_{i,t}(e_{i,t}, tar_{i,t}) = \max\{0, 24 - 0.25 \times |e_{i,t} - tar_{i,t}|\}, \quad (1)$$

depends upon the target $tar_{i,t}$, given by the sum of all signals in periods 1 through $t - 1$ plus the participant’s own signal. That is, in the single-file treatment, the target is simply the sum of signals up to the current period. In the multi-file treatment, participant i ’s target is also the sum of signals up to the present period but excludes signals of the other three participants moving concurrently.⁵ The payoff function penalizes deviations from the target in a linear fashion up to the point at which a participant’s entry lies 96 away from the target,

⁵Including these three signals would not change the optimal strategy in the game.

beyond which there is no punishment for further error.⁶ Upon completion of each period, all participants receive an updated list of all previous entries. Figure 1 shows transcripts of the decision screens for single file (a) and multi file (b), illustrating the similarity between the two treatments.

This social-learning environment corresponds to the logical structure of three different types of observational-learning settings. First, it approximates the standard model of two-state social learning when the action space is the continuum, such that actions reveal posteriors.⁷ Expressed in log-likelihood-ratio terms, each player in such a model would optimally add her private belief (the log likelihood ratio of her signal) to her predecessor’s posterior. Second, it approximates a situation with a binary state about which each person receives a signal corresponding to 100 coin flips, where each flip is extremely weakly correlated with the true state. For example, flips might land heads 51% of the time in State 1 and 49% of the time in State 2. With nearly equal beliefs over the likelihood of two states, Bayesians would update in a manner that is approximately linear in the difference between heads and tails realizations. In this sense, the experimental design also encapsulates the salient features of social learning under weak private signals. The signal structure also lends itself to a third, direct interpretation under which the value of some asset is literally the sum of the signals.⁸ This “wallet-game” signal structure has been tested in other experimental settings (as a

⁶Despite facing flat incentives for entries very far from target, participants whose guesses veer way off target lack any obvious alternative strategy to simply entering their best guesses. Missing the target by more than 96 is a substantial error: because the standard deviation of a player’s signal is 10, and that of the target in period t is $2\sqrt{25t}$ in single-file and $2\sqrt{(t-1)100+25}$ in multi-file, both of which lie below 49 for each t , a participant who simply entered her signal would make a less substantial error more than 95% of the time.

⁷Lee (1993) first analyzed that model.

⁸For example, consider a group of fliers learning about the average meal quality on an airline that uses different catering companies; let $M \subset \mathbb{R}$ be the finite set of all company-branded meals. The quality of the meals also depends upon the flight crew; let $C \subset \mathbb{R}$ be the finite set of all flight crews. For simplicity, meal $m \in M$ prepared by crew $c \in C$ has quality $m + c$. Person i observes own experience $m_i + c_i$. Assume that the cardinality of $M \times C$ is large enough relative to the number of learners that all observations are distinct. To estimate the average experience from eating the $|M \times C|$ possible meals, Person i optimally averages her own experience with any information available about other customers’ experiences. Let $\mu = E[m + c]$ be the expected quality of each type of meal; we can interpret $m_i + c_i - \mu$, namely the “surprise” in i ’s experience,

(a) Single file

You are Participant 5.
Entries in this sequence – Periods 1 to 4:

-8
-10
-6
6

Of your 100 coin flips in this period, 51 came up Heads and 49 came up Tails.

Your entry in Period 5: _____

(b) Multi file

You are in Group 2.
Entries in this sequence – Period 1:

-8 -10 -6 6

Of your 100 coin flips in this period, 51 came up Heads and 49 came up Tails.

Your entry in Period 2: _____

Figure 1: Decision screens

as her signal. Then person 1 estimates the average experience to be

$$b_1 = \frac{m_1 + c_1 - \mu}{|M \times C|} + \mu.$$

In single file, person 2 to combines that with his own experience to judge the average experience as

$$b_2 = \frac{|M \times C|b_1 + m_2 + c_2 - \mu}{|M \times C|} = b_1 + \frac{1}{|M \times C|}(m_2 + c_2 - \mu).$$

In general, for $t \geq 2$,

$$b_t = b_{t-1} + \frac{1}{|M \times C|}(m_t + c_t - \mu).$$

form of common-values auctions (Avery and Kagel (1997)) and corresponds to situations in which separate people observe the value of separate components of an asset—e.g., people care about the sum of everyone’s money but only know the contents of their own wallets. Since the best guess for the final sum coincides with the running total, this interpretation works equally well for a target equal to the sum of all signals as it does for the sum of all present and past signals.⁹ Finally, a potential attraction of using our simple arithmetic set-up is that the underlying model speaks to a new set of applications—social learning games where information about the target is fully revealed by the sequence of consecutive actions.

The 168 participants are students at University College London. Seated at visually separated computer terminals, they first receive and read the experimental instructions and complete a brief understanding test before beginning the computerized games.¹⁰ In seven of the 14 sessions, participants play the single-file game, and the remaining seven sessions they play the multi-file game. Each session includes 12 participants who play either single-file or multi-file three times in a row, resulting in a total of 21 repetitions of each of the two games. For each participant, one of her six choices (chosen at random after the experiment) gets paid out. Because participants receive no feedback about the true value of the target until after all decision-making, the experimental design does little to promote learning across the three games per player and largely prevents us from addressing learning. Nevertheless, Appendix A includes additional results that separate the data between the three games per session.¹¹

To test the instructions, and in order to ascertain whether participants could complete the desired number of games in the ninety minutes allotted to the experiment, we initially

From person 2 onwards, this setting differs from our experimental setup only by the factor $|M \times C|$ on signals, namely only in the sense that it concerns an average rather than a sum.

⁹In the context of social learning, Çelen and Kariv (2005) model the state in this way—as the sum of all signals—yet employ a binary action space and payoff functions with the property that players care only about the state’s sign rather than its magnitude.

¹⁰The experiment uses z-Tree (Fischbacher (2007)). The instructions are available in Online Appendix 2.

¹¹The design choice of excluding feedback about the target was made in order to prevent contamination by giving very heterogeneous feedback and experiences.

piloted the experiment. The first pilot sessions revealed a lack of sufficient time for our desired four games per treatment, but sufficient time for three games. We therefore changed our experiment to include only three games, and made some minor alterations to instructions, post-experimental questionnaire and the payment procedure. We then ran one more pilot on each treatment, after which we changed no other facets of the experiment. The data of the pilot sessions are not, and were not meant to be, included in the data analysis.

Our primary hypothesis was that participants in both experiments would neglect redundancy by explicitly overcounting early entries and thereby implicitly overcounting early signals. Despite our lack of a formal specification of redundancy neglect more general than BRTNI when designing the experiments, we hypothesized that participants' entries would drift above or below their targets in a manner predictable from first-period signals. BRTNI predicts such "momentum", as do many other types of overcounting. Specifically, we hypothesized that positive (negative) first-period signals would be predictive of the event that later entries lie above (below) their targets. This would violate BNE and other rational-expectation predictions. In addition, we hypothesized that participants in the multi-file treatment would not anti-imitate as per BNE, leading them to implicitly overcount early signals.¹² Because BNE in the single-file game lacks anti-imitation, we anticipated that deviations from BNE would be stronger in the multi-file game. If so, then players would earn especially meager payoffs in the multi-file treatment.

Because extreme redundancy neglect might generate entries so large that they could not plausibly be near the target, we expected that at least some participants would recognize that something was amiss and employ some form of correcting behavior. In addition, BRTNI players would understand that entries with magnitudes greatly in excess of 100 could not possibly reflect private signals alone. For neither case did we formulate hypotheses on how

¹²The precise BNE prediction of a player in period t of multi-file is to add one's own signal to the following sum (if applicable): once the sum of $t - 1$ entries minus three times the sum of $t - 2$ entries plus nine times the sum of $t - 3$ entries minus 27 times the sum of $t - 4$ entries plus 81 times the sum of $t - 5$ entries.

participants would adjust their behavior. Rather than to study such mechanisms, our aim was to explore the presence of redundancy neglect, in a setting designed to be inhospitable to it. Our statistical analysis thus sticks closely to the a priori formulated empirical questions.

3 Descriptive Analysis of the Single-File Treatment

Figures 2 and 3 show the evolution of entries across periods in the single-file game through the mean and median, respectively, of the absolute values of entries and targets across the 21 single-file games.¹³




Figure 2: Single-file mean absolute entries and targets

Figure 2 shows that after a few periods during which the mean entry nearly coincides with the mean target, a large disparity emerges before swiftly vanishing. The underlying fluctuation of entries derives solely from one game, Game 34. In it, several participants

¹³By using absolute values of each entry, the figures treat positive and negative entries symmetrically. One may worry about a potential bias towards making positive-valued entries, but such a bias cannot be discerned in our data. In the single-file treatment, the random signals happen to be strictly negative (49%) more often than strictly positive (43%), leading to an overall tendency towards negative entries (56% negative, versus 41% positive). The asymmetry in signals is particularly strong in $t = 1$, where 15 out of 21 (71%) of signals happen to be strictly negative.




Figure 3: Single-file median absolute entries and targets

(those in periods 3, 4, 7, 9, 10) act like BRTNI by adding up all previous entries, running entries up above 1800 before subsequent players make corrections by choosing entries near zero. Although this episode strongly affects the mean across all 21 games, Game 34 is unusual: Figure 3 shows that across games the median entry closely tracks the median target. Overall, mean and median entries in the single-file game depart only mildly from Nebi play (i.e., from best responding to BNE play from everyone else). The tables in Online Appendix 1 provide a full account of all raw data in each game, showing that in three of 21 games all entries nearly coincide with their targets.¹⁴

But, as Game 34 indicates, there are some systematic deviations from optimal behavior, including redundancy neglect. In Table 1, we organize the individual entries by classifying them according to their *exact* consistency with the Level- k family of models. This family was developed for games of complete information by Stahl and Wilson (1994) and Nagel (1995) and subsequently applied to games of incomplete information (e.g., Crawford and Iriberri (2007)). We follow previous applications to social learning games (e.g. Kübler and

¹⁴In one game, all entries match targets exactly; in another, the same would be true but for someone who flips the sign of his or her private signal; in the third game, someone appears to have made the mildest of arithmetic mistakes (mis-summing -36 and -8 to -46).

Weizsäcker (2004)) by maintaining the assumption that Level-0 play randomizes uniformly over available actions, independent of private signal. In our games, several members of this family correspond to natural prototypical behaviors: “fully cursed” behavior is equivalent to Level 1 (following only one’s own signal); BRTNI (full redundancy neglect) coincides with Level 2; best responding to BRTNI/Level 2 is Level 3; and Nebi in period t agrees with Level k for $k > t - 1$. For each model, the table reports the number of decisions consistent with the model’s specified strategy.¹⁵

¹⁵Because participants move twice per game, in the second half of the game they may treat their own past actions differently than those of other players. Due to the subtlety of this asymmetry, we adopt expansive definitions of cursed, BRTNI, Level-3 and Nebi play, coding an action as consistent with Level- k regardless of whether players treat their past actions differently than those of other players. Likewise, we code an entry as consistent with Level- k if the player best responds to the beliefs that others play a Level- $(k - 1)$ strategy and that these others treat their own prior actions the same as they would treat entries by other players. Altogether, this necessitates allowing for two different versions of fully cursed and Nebi strategies (accounting for own previous signal versus not), three different versions of BRTNI (ignoring multiple entries per player, accounting only for own previous entries, and accounting for own and others’ previous entries) and four different versions of Level-3 (ignoring multiple entries per player, accounting only for own previous entries, accounting for own and others’ previous entries but ignoring that others account for their predecessors’ previous entries, and full accounting for all previous entries). In all cases the simplest version of the Level- k model has the highest consistency rate; allowing for the more sophisticated versions makes only a minor difference, as Table 11 (and Table 16 for multi file) illustrates.

Table 1: Single-file entries consistent with different behavioral models

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	NA
t=1	21	21	21	21	0
t=2	1	18	18	18	2
t=3	1	6	14	14	2
t=4	1	5	0	15	2
t=5	1	2	0	17	1
t=6	2	1	1	15	3
t=7	1	1	0	17	2
t=8	1	1	0	17	3
t=9	1	2	0	18	1
t=10	0	1	0	14	6
t=11	3	1	0	11	6
t=12	1	0	0	14	6
t=13	4	0	0	16	3
t=14	4	1	0	14	5
t=15	0	0	0	15	6
t=16	3	0	0	16	3
t=17	1	1	0	20	0
t=18	2	0	0	15	5
t=19	1	1	0	19	1
t=20	0	0	0	18	3
t=21	0	0	0	12	9
t=22	1	0	0	15	6
t=23	3	0	0	17	4
t=24	2	2	0	17	2
Total in $t \geq 3$	33	25	15	246	79
Total in $t \geq 4$	32	19	1	332	77
Total	55	64	54	385	81

In each of the first three periods, whenever two or more models make the same prediction, Table 1 codes participants as following more than one type. From period 4 onwards, however, the different models are identified by different predictions except in a small number of cases (20 of 441, or 4.5%, of classifications) where serendipity produces signals and previous entries that align just so. The behavior of 16 of the 81 participants classified as “NA” can be attributed to one of the models in the table by allowing for the possibility that the participant inadvertently flips the sign of his or her private signal, the model’s predicted action minus the signal (namely his or her inference from others according to the model), or both.

A striking feature of the table is that the large majority of entries is consistent with Nebi play, i.e. naïve BNE play, or equivalently with a Level- k model where $k > t - 1$. Of the 504 entries in our data set, 385 are consistent with this prediction, and the proportion is constant in t even in the latter half of games. Such sophistication stands in stark contrast to all other estimates of Level k of which we are aware.¹⁶

The same result appears if we consider jointly the six choices that a single participant makes. Appendix Table 15 summarizes the individual participants’ decision patterns, showing that 57 of the 84 participants playing the single-file game (68%) choose the Nebi action in at least five out of six opportunities. No other model (among Level 1, Level 2/BRTNI or Level 3) has even a single such adherent. A separate analysis also finds that few participants learn to play higher-level models over the course of the experiment’s three games. Only three of 84 participants make both entries consistent with Level m in a first game before making both entries consistent with Level n , for $n > m$, in a later game.¹⁷

The single-file game appears to be so simple that participants understand the logic of BNE—they simply add their own signal to the previous period’s entry—and can apply it even

¹⁶Nebi’s high hit rate reveals that the vast majority of participants do not round their entries to multiples of 10, or similar. To the extent that a minority of participants do round, we would under-estimate Nebi’s consistency with the data.

¹⁷No participant does the opposite. In multi file, the numbers are the same: three participants consistently show a higher level later in the experiment than earlier, and no participant shows the reverse.

in late periods in the game. This behavioral rule does not require anti-imitation and, in the single-file game, is optimal if and only if everyone else follows it. Because not everyone does adhere to Nebi play, the Nebi strategy is not empirically optimal for late movers. Indeed, in periods $t > 3$, while 75% follow their Nebi strategy, only 17% hit their target—because entries have previously departed from Nebi play. While playing Nebi may be arithmetically simple, best-responding to predecessors who play differently presents greater challenges. Level 3, in contrast, has a very low hit rate, in periods where it does not coincide with Level 2/BRTNI and/or Nebi play.

Table 2 also shows that BRTNI’s precise prediction of strong redundancy neglect attracts support only in the earlier rounds, and only in a relatively small proportion of cases. Moreover, Appendix A shows that BRTNI’s hit rate tends to decrease in the participants’ experience of playing the game.

This, however, still leaves open the possibility that other forms of redundancy neglect occur—as we discussed above, a large set of behaviors would lead to overcounting. We return to the issue of overcounting in Section 5 and here merely document that BRTNI’s point prediction often predicts the direction of participants’ departure from BNE: of the 377 entries off target, 270 (72%) err in the direction of BRTNI.¹⁸ Figure 19 of Appendix A shows the distribution of entries relative to their targets.

Altogether, the evidence in the single-file game suggests that most participants employ the Nebi strategy. They are able to use the actions of others to their own benefit. Average earnings are GBP 18.25, whereas simply relying on one’s own signal(s) would pay GBP 16.99 on average. Nevertheless, the presence of a small minority of participants who do not follow BNE—and tend to neglect redundancy—drags overall behavior away from the target. Because errors by a minority have long-lived effects, people moving later in the game

¹⁸As described in Footnote 15, we frequently report (as we do here) a simplified variant of BRTNI that does not fully match BRTNI’s proper definition, which would have her assume that any predecessors’ second move is the sum of that predecessor’s two signals, rather than simply her second signal. It matters little for our analysis here or elsewhere whether we use the full or simplified definition.

benefit less from social learning. If the games were even longer, this effect would increase. For illustration, consider the following simulation. Of the 305 decisions classified as cursed, BRTNI and Nebi for periods $t \geq 3$ in Table 1, 8% were cursed, 11% were BRTNI and 81% Nebi. Assuming that each player behaves as one of these three types, with likelihoods given by the classification, we simulate a hypothetical 48-period game, in which each player moves four times, 100,000 times. The average and median simulated payoffs from the game's last entry (in $t = 48$) are 8.26 and 3.50. A player in $t = 48$ who instead ignored the other players and played the sum of his four private signals would have average and median simulated payoffs of 11.87 and 13.

Finally, we ask what would be the ex-post optimal behavior, given other's choices. Under the maintained hypothesis that participants use their own signals correctly, we can decompose the entry $e_{i,t}$ of a participant i with signal $s_{i,t}$ into $e_{i,t} = s_{i,t} + (e_{i,t} - s_{i,t})$; the term $e_{i,t} - s_{i,t}$ measures what the participant infers about the target from predecessors. We consider a class of alternative rules $e'_{i,t}(\gamma) := s_{i,t} + \gamma(e_{i,t} - s_{i,t})$, where $\gamma \geq 0$ represents a shading factor, and identify the value of γ that maximizes the participant's payoff. A value of $\gamma < 1$ indicates that the participant overshoots the target by over-inferring from predecessors; a value of $\gamma > 1$ indicates that the participant undershoots the target by under-inferring from predecessors. For all participants, we find that $\gamma = 0.69$ maximizes payoffs: the average participant over-infers, and would have earned GBP 18.51 (instead of GBP 18.25) by shading her inference by 31%. Yet this disguises very substantial heterogeneity between those who exhibit Nebi play and those who do not. For non-Nebi players, we find that $\gamma = 0.28$ would have earned GBP 17.16 (instead of GBP 15.12). For Nebi players, $\gamma = 0.90$ would have earned them GBP 19.33 (instead of GBP 19.28).

4 Descriptive Analysis of the Multi-File Treatment

Figures 4 and 5 are the multi-file analogs to Figures 2 and 3, respectively. They depict the mean and median entries relative to targets across periods. The median in period t of Figure 5 corresponds to the median across the 21 games of the mean of the four period- t entries, while the mean in period t of Figure 4 represents the mean across the 21 games of the mean of the four period- t entries.




Figure 4: Multi-file mean absolute entries and targets

Players in the multi-file games deviate much more from their targets than players in the single-file games, and later players do not correct earlier players' errors. Like in most single-file games, in most multi-file games the (absolute) target lies between 10 and 70 in the final period. Yet participants make entries whose absolute values are higher by an order of magnitude. The average absolute $t = 6$ entry surpasses 600 and this is not driven by outliers: in a majority of games, the final-period average exceeds 500. The deviations from target start accumulating in $t = 3$, the first period in which redundancy neglect can have an impact, and by $t = 5$ most games have mean entries that outstrip their targets by tenfold.

Although on average participants make choices that are too extreme, in several games




Figure 5: Multi-file median absolute entries and targets

a subset of them appear to recognize that entries are too extreme and take corrective action. Typically, however, they do not influence the crowd’s belief enough to prevent later participants from making even more extreme entries. Table 2 gives an example in the form of Game 17.¹⁹ It shows for each period the sum of previous signals ($\sum_{i=1}^4 \sum_{t'=1}^{t-1} s_{i,t'}$) as well as each player’s signal $s_{i,t}$ (in brackets)—summing the two gives the target—as well as the player’s entry.

¹⁹We selected this game as typical in its variability of behavior; Online Appendix 1 provides a full account of the data.

Table 2: Signals and entries of Game 17

	$\sum_{i=1}^4 \sum_{t'=1}^{t-1} s_{i,t'}$	$s_{1,t}$	$e_{1,t}$	$s_{2,t}$	$e_{2,t}$	$s_{3,t}$	$e_{3,t}$	$s_{4,t}$	$e_{4,t}$
t=1	[0]	[-8]	-8	[-10]	-10	[-6]	-6	[6]	6
t=2	[-18]	[-2]	-20	[-6]	-36	[16]	2	[0]	-18
t=3	[-10]	[-4]	-24	[-8]	-98	[6]	-12	[-6]	-96
t=4	[-22]	[8]	-24	[-8]	-150	[-6]	-34	[6]	-26
t=5	[-22]	[4]	-16	[-2]	-40	[6]	-584	[2]	-534
t=6	[-12]	[-18]	-34	[-16]	-1654	[10]	-2046	[-4]	-1732

Entries in the first two periods equal or approximate targets, as most reasonable models would predict. From period 3 onwards, however, Nebi play prescribes anti-imitation: the players should realize that the negative entries in $t = 2$ share a common source in the form of $t = 1$ entries. Accounting for this redundancy, while at the same time gleaning information about $t = 2$ signals from $t = 2$ play, requires $t = 3$ players to imitate entries in $t = 2$ and anti-imitate those in $t = 1$. Yet two of the four players in $t = 3$ of Game 17 do not follow this logic and report entries consistent with BRTNI: they simply add their signal to the sum of previous entries. In $t = 4$, three of the four players behave in ways more moderate than BRTNI, and one player chooses an extreme entry of -150 , an instance of strong redundancy neglect. In $t = 5$ and $t = 6$, several entries are even more extreme. One of them, the entry of -2046 in $t = 6$, is actually Nebi play from a participant who, while rather smart, makes the game's most severe prediction error! Overall, the example shows that in spite of some players' attempts to moderate behavior along the way, the significant number of strong redundancy neglecters propagates extreme beliefs.²⁰

²⁰The extreme behavior of this and many other participants raises the questions as to whether subjects understand the target's ex-ante distribution, which is concentrated around zero, and whether such understanding helps to avoid extreme beliefs. In our debriefing questionnaire, we asked the participants the following "If a coin is flipped 2500 times, what is the probability that ($\#$ Heads $-$ $\#$ Tails) lies between -100

For comparison with single file, the following table reports the consistency of the data with the various members of the Level- k family of models. It too indicates that these models fit the data very differently in the multi-file treatment than they do in the single-file treatment.

Table 3: Multi-file entries consistent with different behavioral models

Period	Cursed	BRTNI	Level 3	Nebi	NA
	($k = 1$)	($k = 2$)	($k = 3$)	($k > t - 1$)	
t=1	81	81	81	81	3
t=2	3	67	67	67	14
t=3	3	18	29	29	37
t=4	4	22	0	16	44
t=5	1	29	0	6	49
t=6	1	8	0	2	73
Total in $t \geq 3$	9	77	29	53	203
Total in $t \geq 4$	6	59	0	24	166
Total	93	225	177	201	220

Table 3 shows that Nebi fits the data well in $t = 1$ and $t = 2$, periods in which its prediction coincides with BRTNI and Level 3, and even in $t = 3$, where it makes a different prediction than BRTNI. Overall, 63% of entries in the first three periods hit their targets. In periods 4, 5 and 6, the success rate falls to 6%. For these periods, Nebi involves intricate imitation as well as anti-imitation. The fact that $\frac{24}{252} \approx 10\%$ of decisions in the second half of the experiment match Nebi demonstrates a high degree of sophistication amongst some participants. However, in each of $t = 4, 5, 6$, BRTNI fits a higher proportion of entries than Nebi or the other models.²¹ As Appendix A Table 20 shows, the same pattern appears if we and 100 for these coin flips?" Surprisingly, the accuracy of their answers does not predict their earnings. The correlations are -0.05 (p -value 0.63, Pearson product-moment test) for single file and -0.16 (p -value 0.18) for multi file.

²¹Of the 220 unexplained observations, 11 can be explained by enriching one of the proposed models by

consider all six choices of a single subject. Only four participants make five or six choices consistent with Nebi, whereas 13 participants make five or six choices consistent with BRTNI. Figure 20 in the Appendix depicts the deviations from target for all periods. Just like in the single-file treatment, BRTNI predicts the systematic direction of the deviations: 78% of deviations lie on BRTNI's side of zero.

Also similar to the single-file treatment, BRTNI behavior diminishes with experience in multi file; see Tables 17 to 19 in Appendix A.²² But regardless of whether participants follow BRTNI or another form of redundancy neglect, they make far too many extreme entries, which Online Appendix 1 documents for nearly all games.

Altogether, the multi-file treatment induces strong herding that leads participants severely astray. 44% of the data are consistent with BRTNI, and more are consistent with a general propensity to neglect redundancy, which leads to increasingly extreme and off-target predictions. Many participants, especially late movers, estimate the targets to lie in extremely unlikely regions. Such misestimation comes at a price: participants in the last three periods earn an average of GBP 8.90, whereas the simple strategy of reporting one's own signal would have earned GBP 16.60. Across all periods in the multi-file treatment, participants earn an average of GBP 14.93, whereas reporting their signals would have earned them GBP 18.32. Not only do participants learn sub-optimally, but they mislearn so acutely that they would be better off without the possibility of learning! To our knowledge, this is the first experiment to document that such an effect occurs on average across the entire population of players.

Finally, Figure 6 illustrates how early signals come to excessively influence later play by establishing a relationship between the sign of $t = 1$ signals and later deviations from target. It depicts the average entries and targets (not their absolute values) in multi-file games whose

allowing participants to flip signs of their private signals or flip the signs of what they infer from their predecessors, or both.

²²We especially note that in the third game of each session, Nebi play, which includes anti-imitation, occurs more frequently than the simpler BRTNI play.

first four signals sum to positive values versus the average entries and targets in multi-file games whose first four signals sum to negative values. Because the signals are *i.i.d.*, the sign of the first four signals does not predict later signals, and, hence, the targets (dotted line) remain stable on average after $t = 1$. Participants' entries, however, differ dramatically depending upon the sign of the sum of first-period signals. Early positive signals generate positive momentum whereby later entries tend to exceed their targets, increasingly so over the course of the game. Figure 21 in Appendix A gives an analogous picture for single file that shows no significant momentum.

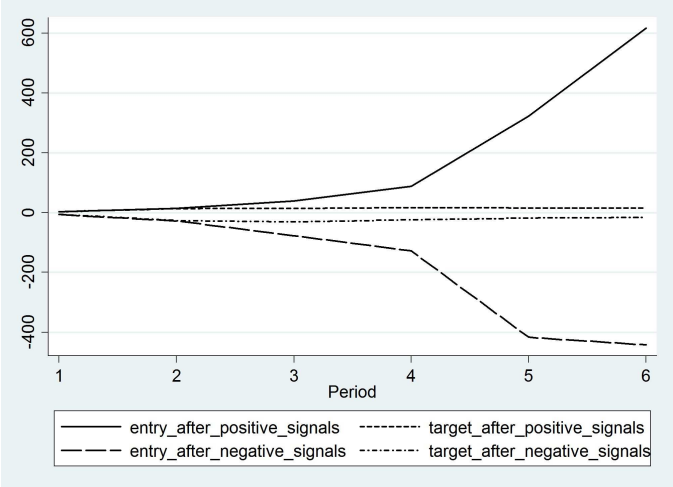


Figure 6: Average entry and target in multi file, separated out by the sign of the sum of the game's first-period signals

Overall, whereas entries approximate targets fairly well in the single-file treatment whose observation structure is standard in the literature, we find much stronger evidence for redundancy neglect in the multi-file treatment. This discrepancy suggests that our experimental setup per se does not induce strong deviations from target; rather, only in the multi-file treatment in which sequential rationality predicts anti-imitation do participants strongly and reliably fail at rational inference.

5 Regression Analysis

In this section, we present linear regressions that test for both the presence of over-counting and the lack of anti-imitation. These tests were anticipated in the design of our experiment: redundancy neglect predicts that entries react too strongly to earlier signals (over-counting) and that regression coefficients for certain earlier entries are not sufficiently negative (lack of anti-imitation). The analysis first focusses on the multi-file treatment, before turning to the single-file treatment for comparison.

We first investigate over-counting in multi file, describing the connection between early signals and later entries. This parallels Figure 6’s nonparametric description. We regress the participants’ period- t entries, $e_{i,t}$, on the period- t signals they receive, $s_{i,t}$, as well as on the sum of all signals in every prior period t' , $\bar{s}_{t'} = \sum_{i=1}^4 s_{i,t'}$. Nebi predicts that all coefficients equal one, since all signals are correctly accounted for in equilibrium.²³ BRTNI makes the same predictions for $t = 1$ and $t = 2$. For periods $t \geq 3$, BRNTI makes two distinct qualitative predictions.²⁴ First, it predicts that for each i and $t' < t - 1$, the estimated effect of $\bar{s}_{t'}$ on $e_{i,t}$ should exceed one. Second, it predicts that for each i and $t' < t$, the effect of $\bar{s}_{t'}$ on $e_{i,t+1}$ should exceed the effect of $\bar{s}_{t'}$ on $e_{i,t}$. Moreover, Eyster and Rabin (2014) show that *any* rule whereby players neglect redundancy makes the first prediction above. In addition, any rule like BRTNI whereby players correctly weight their immediate predecessors, and fail to anti-imitate, makes the second prediction. Since signals are exogenous and mutually independent in our design, these hypotheses can be well tested with a regression; all coefficients have causal interpretations. Table 4 presents the regression results, where player indexes are omitted from the dependent variables e_t for conciseness

²³Because Nebi/BNE also predicts that a constant regressor has a zero coefficient, we omit the constant. Empirically, the inclusion of a constant regressor leaves the results essentially unchanged.

²⁴Precisely, BRTNI predicts that for each i and t , $e_{i,t} = s_{i,t} + \sum_{t' < t} \bar{e}_{t'}$, where $\bar{e}_{t'} = \sum_{i=1}^4 e_{i,t'}$. This prediction implies that $e_{i,t} = s_{i,t} + \sum_{t' < t} 5^{t-1-t'} \bar{s}_{t'}$, so that third-period entries correctly weight \bar{s}_2 but quintuple-count \bar{s}_1 ; fourth-period entries correctly weight \bar{s}_3 , quintuple-count \bar{s}_2 , and overcount \bar{s}_1 twenty-five fold; etc.

(similarly in the subsequent tables).

Table 4: Multi-file regressions of period- t entries on current and past signals.
(Standard errors in parentheses are clustered by session.)

	e_1	e_2	e_3	e_4	e_5	e_6
\bar{s}_1	-	1.025 (0.056)	2.983 (0.426)	3.073 (0.832)	13.513 (3.113)	19.609 (5.447)
\bar{s}_2	-	-	0.999 (0.231)	3.951 (1.307)	9.112 (2.459)	10.101 (10.426)
\bar{s}_3	-	-	-	1.066 (0.914)	6.711 (3.912)	10.362 (7.240)
\bar{s}_4	-	-	-	-	- 4.786 (3.750)	-1.553 (9.511)
\bar{s}_5	-	-	-	-	-	3.683 (5.912)
s_t	0.912 (.070)	1.067 (.106)	1.359 (0.812)	-1.555 (2.358)	-0.702 (7.238)	1.195 (15.480)
R^2	0.83	0.90	0.69	0.25	0.44	0.19
obs.	84	84	84	84	84	84

For $t' = 1, 2$, signals in t' affect the entries in $t' + 1$ with weights of approximately one, as predicted by BNE and BRTNI. This suggests that on average, players in early rounds correctly imitate their immediate predecessors. However, these same signals attract far larger coefficients in periods $t' + 2, t' + 3, \dots$ (the smallest point estimate being 3.0 and the largest 19.6). In addition, along each of the first three rows in the table, the estimated coefficients increase monotonically, as predicted by BRTNI and other redundancy-neglect models. The fact that participants implicitly over-count early signals so dramatically illustrates how far behavior deviates from Nebi. Late actions implicitly weight the early signals very heavily, consistent with a substantial degree of redundancy neglect. For $t' = 3, 4$, the regressions' large standard errors render all of the estimated coefficients statistically insignificant. Indeed, in these periods, even the participants' own signals do not have significant effects on their actions.

We now turn to tests of anti-imitation. When players move multi-file, BNE calls for them to anti-imitate some of their predecessors in order to avoid inefficiently over-counting early signals. We regress participants' period- t entries $e_{i,t}$ on their period- t signals $s_{i,t}$ and

on lagged entries $\bar{e}_{t'}$, with $t' < t$ and $\bar{e}_{t'} = \sum_{i=1}^4 e_{i,t'}$. BNE/Nebi predicts a coefficient on $s_{i,t}$ equal to one and coefficients on $\bar{e}_{t'}$ that oscillate and diverge, with predicted levels of 1, -3, 9, -27, 81 for $t' = t-1, t-2, t-3, t-4, t-5$, respectively. BRTNI predicts that every entry gets the same coefficient of 1. Other rules embedding different forms of redundancy neglect predict that all of the coefficients are non-negative, and that their sum exceeds one. Table 5 shows the regression results.

Table 5: Multi-file regressions of period- t entries on current signals and past entries. (Standard errors in parentheses are clustered by session.)

	e_1	e_2	e_3	e_4	e_5	e_6
\bar{e}_1	-	0.925 (0.041)	0.088 (1.057)	- 3.908 (7.569)	-11.623 (8.605)	38.315 (16.924)
\bar{e}_2	-	-	0.700 (0.300)	2.037 (1.918)	3.102 (2.214)	-12.636 (4.502)
\bar{e}_3	-	-	-	-0.011 (.307)	0.456 (.403)	0.928 (.857)
\bar{e}_4	-	-	-	-	0.519 (.149)	0.299 (.339)
\bar{e}_5	-	-	-	-	-	0.254 (.122)
s_t	0.912 (.070)	0.956 (.104)	1.245 (0.663)	-0.611 (2.515)	-1.277 (3.817)	-1.460 (13.209)
R^2	0.83	0.95	0.70	0.24	0.54	0.28
obs	84	84	84	84	84	84

The coefficients differ from the anti-imitation pattern predicted by BNE/Nebi play: of the six predicted negative coefficients, four have estimated positive signs. Altogether, most coefficients in the table are estimated to be insignificantly different from zero, and most differ significantly from Nebi’s prediction.²⁵ Anti-imitation should appear most simply in the third period, where BNE/Nebi call for it for first time. Given behavior in $t = 1$ and $t = 2$, participants in $t = 3$ should anti-imitate the actions in $t = 1$. However, the coefficient of \bar{e}_1 in the regression of e_3 is close to zero and differs significantly from the Nebi-predicted value of -3.

²⁵Note, however, that the property of redundancy-neglect models that the sum of previous entries’ coefficients exceeds one is unconfirmed—directed null hypotheses in either direction cannot be rejected.

Overall, the regressions of Table 5 confirm that the central tendency of the data patterns does not include anti-imitation. This can also be tested via likelihood-ratio test of the joint hypothesis that all coefficients are non-negative. In each of the six regressions of Table 5, the null hypothesis that all coefficients are non-negative cannot be rejected at any reasonable significance level, with p -values of 1.00, 1.00, 1.00, 0.82, 0.80 and 0.72, respectively.

Rather than confirm BNE/Nebi’s prediction for multi file, the estimated decision weights of Table 5 are somewhat more reminiscent of BNE/Nebi’s prediction for the *single*-file treatment, in which all weights of previous entries should be non-negative. To explore whether behavior in both treatments shows the same pattern, we turn to analogous regressions for single file. To readily compare coefficients with those of the multi-file treatment, we group periods $t = 1, \dots, 4$ of single file into “super-period” $\tilde{t} = 1$, periods $t = 5, \dots, 8$ into super-period $\tilde{t} = 2$, and so forth. This coarse time structure suppresses the sequencing of moves within super-periods in order to facilitate comparison of the coefficients to those of Table 4: we regress the single-file participants’ super-period- \tilde{t} entries, $e_{i,\tilde{t}}$, on their period- t signals, $s_{i,t}$, as well as on the sum of all signals in every prior super-period \tilde{t}' , $\bar{s}_{\tilde{t}'} = \sum_{i=1}^4 s_{i,\tilde{t}'}$. Table 6 presents the regression results.

Table 6: Single-file regressions of super-period- \tilde{t} entries on current and past signals. (Standard errors in parentheses are clustered by session.)

	e_1	e_2	e_3	e_4	e_5	e_6
\bar{s}_1	-	1.302 (0.697)	5.751 (4.849)	1.672 (1.027)	1.664 (.701)	0.680 (1.897)
\bar{s}_2	-	-	1.036 (0.875)	1.211 (1.166)	1.376 (.956)	-1.942 (.911)
\bar{s}_3	-	-	-	1.502 (.558)	1.584 (.289)	0.017 (1.112)
\bar{s}_4	-	-	-	-	0.496 (.477)	2.553 (.589)
\bar{s}_5	-	-	-	-	-	-0.382 (1.145)
s_t	0.782 (.155)	1.380 (.437)	0.017 (2.933)	0.500 (.569)	1.228 (.388)	-1.337 (2.358)
R^2	0.18	0.25	0.16	0.41	0.64	0.20
obs.	84	84	84	84	84	84

BNE/Nebi play predicts that all coefficients equal one, since all signals are correctly

accounted for in later periods. Despite most point estimates exceeding one, only three of the 15 coefficients that describe the influence of past signals differ statistically significantly from one (and one of these three has the opposite sign). This contrasts starkly with the results from multi file. In particular, the regressions provide no indication that early signals exert increasingly strong influence on later and later actions.

One may speculate that the reaction to earlier entries, not signals, is comparable between the two treatments. Appendix A Table 21 refutes this and shows significant differences from Table 5, which reports many coefficients much further from zero.²⁶

Overall, the comparison of regressions for single file and multi file shows strong dissimilarities between the treatments. An important caveat for this observation is that the theoretically predicted coefficients (of BNE, BRNTI, and most other reasonable models) also differ between treatments. In the next section, we turn to a structural model of probabilistic behavior that implements the respective rules of the two games. There, estimates will indicate that even the structural patterns of behavior differ between the treatments.

6 Structural Model Estimation

This section complements the previous analyses by estimating a structural model of decision-making that encompasses BRTNI, Nebi and other behaviors; it also generalizes both Quantal-Response Equilibrium (QRE) (McKelvey and Palfrey 1995) and Level k . As in Kübler and Weizsäcker (2004), players noisily best respond to their beliefs, and regard other players as making imprecise choices too. In particular, player j “better responds” to her beliefs about her target by applying a probabilistic best-response function characterized by the precision parameter λ^1 . A higher value of λ^1 indicates greater precision; the limiting case of $\lambda^1 = \infty$ corresponds to the true best response to player j ’s beliefs. Player j ’s probabilistic beliefs

²⁶We include the table for completeness but relegate it to Appendix A because the BNE/Nebi and BRTNI predictions do not harmonize with the super-period presentation.

about her target depend naturally upon her beliefs about her predecessors' behavior. We assume that player j believes that each prior mover i also acts probabilistically, but with a potentially different precision λ^2 . (The superscripts refer to the level of reasoning.) Player j 's beliefs also depend upon her beliefs about i 's beliefs about his predecessors, etc. We assume that j thinks i thinks that his predecessors act with precision λ^3 , and so on, for all levels of reasoning.

The following subsection describes the model's precise updating process where, like in Kübler and Weizsäcker (2004), the noise in best responses $(\lambda^1, \lambda^2, \dots)$ can result from logistic disturbances to players' utilities (or their beliefs about other players' disturbances, beliefs about other players' beliefs about still other players' disturbances, etc.). For reasons of complexity, we cannot directly estimate this model in the context of our experiment. However, Subsection 6.2 describes an estimable approximation that uses coarser action and signal grids.²⁷ Subsection 6.3 describes a different variant of the model, where we make a normality assumption on the probabilistic choice itself. Both model estimations largely confirm the conclusions drawn from the analysis of the previous sections.

6.1 The model

Let $F_{i,t}(x; h_{i,t}, s_{i,t}, (\lambda^1, \lambda^2, \dots))$ denote the probabilistic belief of player i in period t that her target $tar_{i,t}$ does not exceed x , following the history $h_{i,t} = (e_{l,\tau})_{l,\tau}$ of observed entries, the player's private signal $s_{i,t}$, and for a given vector of precision parameters $(\lambda^1, \lambda^2, \dots)$. The logistic best response to the beliefs $F_{i,t}$, characterized by precision λ , is a probability distribution that assigns entry $e_{i,t}$ the probability

$$\sigma(e_{i,t}; F_{i,t}, \lambda) = \frac{\exp(\lambda(\int \pi_{i,t}(e_{i,t}, tar_{i,t})dF_{i,t}))}{\sum_e \exp(\lambda(\int \pi_{i,t}(e, tar_{i,t})dF_{i,t}))}, \quad (2)$$

²⁷We are indebted to Tom Palfrey for suggesting this approach.

where the payoff of an entry for a given target is given by (1), and the denominator sums over all available entries e . We assume that all players' choices have common precision λ^1 .

Any player i acting in $t = 1$ is assumed to have correct beliefs about her target, namely $F_{i,1}(x; s_{i,1}, (\lambda^1, \lambda^2, \dots)) = \mathbf{1}_{x \geq s_{i,1}}$ for any $(\lambda^1, \lambda^2, \dots)$. The player applies $\sigma(\cdot; F_{i,1}, \lambda^1)$ to her belief and thus chooses an entry equal to her signal (and, hence, target) with greatest probability, and increasingly distant entries with lower yet positive probabilities. Subsequent players' beliefs depend on the history and on $(\lambda^2, \lambda^3, \dots)$, as players attempt to back out their predecessors' signals.²⁸ The main assumption of our analysis is that j reasons about her predecessors' decision-making through the truncated precision vector $(\lambda^2, \lambda^3, \dots)$ as described above.

For a simple example of the approach, consider player j acting in $t = 2$ of the single-file game. Player j , firstly, knows that player i received a private signal that he interprets correctly, and, secondly, believes that i better responds to his beliefs with precision λ^2 . Player j 's beliefs about her own target are thus:

$$F_{j,2}(x; h_{j,2}, (\lambda^1, \lambda^2, \dots)) = \Pr [s_{i,1} \leq x - s_{j,2} | e_{i,1}, (\lambda^1, \lambda^2, \dots)]$$

$$= \int_{\tilde{s}_{i,1} \leq x - s_{j,2}} \frac{\sigma(e_{i,1}; \mathbf{1}_{x \geq \tilde{s}_{i,1}}, \lambda^2) \phi(\tilde{s}_{i,1})}{\int_{\tilde{s}'_{i,1} \leq 100} \sigma(e_{i,1}; \mathbf{1}_{x \geq \tilde{s}'_{i,1}}, \lambda^2) \phi(\tilde{s}'_{i,1}) d\tilde{s}'_{i,1}} d\tilde{s}_{i,1},$$

where $\phi(\hat{s})$ denotes the probability of an individual signal taking on value \hat{s} .

Likewise, for general t and $s < t$, j 's beliefs about the probability that i 's signal $s_{i,\tau}$ lies

²⁸For simplicity, we ignore the fact that players choose twice and therefore have different information about an earlier signal. See Table 11 and Table 16 for evidence that our participants ignore this, too.

below some value y , conditional upon i 's entry $e_{i,\tau}$ as well as the rest of $h_{j,t}$, is

$$G_{i,\tau}^{j,t}(y; h_{j,t}, (\lambda^1, \lambda^2, \dots)) = \Pr [s_{i,\tau} \leq y | h_{j,t}, (\lambda^1, \lambda^2, \dots)]$$

$$= \int_{\tilde{s}_{i,\tau} \leq y} \frac{\sigma(e_{i,\tau}; F_{i,\tau}(x, \tilde{s}_{i,\tau}, (\lambda^2, \lambda^3, \dots)), \lambda^2) \phi(\tilde{s}_{i,\tau})}{\int_{\tilde{s}'_{i,\tau} \leq 100} \sigma(e_{i,\tau}; F_{i,\tau}(x, \tilde{s}'_{i,\tau}, (\lambda^2, \lambda^3, \dots)), \lambda^2) \phi(\tilde{s}'_{i,\tau}) d\tilde{s}'_{i,\tau}} d\tilde{s}_{i,\tau},$$

where the hidden information $\tilde{s}_{i,\tau}$ and $\tilde{s}'_{i,\tau}$ represent randomness from j 's perspective. Using $G_{i,\tau}^{j,t}(\cdot)$, player j 's belief about her own target $F_{j,t}$ can be calculated recursively.²⁹ For each level of reasoning about other players, one more element of the precision vector is truncated to determine higher-order beliefs.

Notice that in the construction of $F_{j,t}$, players from $t = 3$ onward consider others' entries multiple times and with differently truncated precision vectors. A player acting in period 3, for instance, may interpret the entry $e_{i,1}$ in a different way than he believes players in period 2 do. The model thus allows for redundancy neglect when λ^3 is relatively low, and similarly for many other deviations from rational expectations about others' play. BRTNI corresponds to the parameter values $\lambda^1 = \lambda^2 = \infty$ and $\lambda^3 = \lambda^4 = \dots = 0$.

6.2 Approximation on a coarse grid

The model presented in the previous subsection defies a one-to-one implementation due to the enormous number of possible targets and histories, which renders the calculation of all relevant beliefs and likelihoods infeasible. To estimate our limited-depth-of-reasoning model, we simplify the game by constraining all signals, targets and entries to lie on a much smaller

²⁹Letting $I_{j,t} = \{(i, s) : s < t\}$ be the set of player-entry pairs prior to (j, t) , the belief $F_{j,t}$ is the $|I_{j,t}|$ -fold convolution of the random variables with distributions $\{G_{i,s}^{j,t}(\cdot)\}_{(i,s) \in I_{j,t}}$. These variables are independent, conditional on $h_{j,t}$: for players k, i, j acting in periods $\tau < s < t$, respectively, knowledge of player k 's signal does not affect the distribution $G_{i,s}^{j,t}$ because $G_{i,s}^{j,t}$ conditions on k 's entry (making knowledge of $s_{k,\tau}$ superfluous in the construction of $G_{i,s}^{j,t}$), and conversely, knowledge of i 's signal does not affect the distribution $G_{k,\tau}^{j,t}$ because i 's entry does not inform it.

grid of possible values. The central simplifying assumption is that each participant’s private information derives from only two coin flips, rather than from the 100 in the experiment.

In our “grid game”, we take a player’s signal to be the sum of two independent and equiprobable random variables that can take on one of two values $\{-\delta, \delta\}$. The level of δ is fixed at $\sqrt{50} \approx 7.07$, chosen to equate the target’s variance in the grid game to that in the original game.³⁰ Just as the original game, the grid game becomes richer as more information gets aggregated. By analogy to the coarse signal structure, we force the players’ targets and entries to belong to V_t , whose values are integer multiples of 2δ with minima and maxima given by the smallest and largest possible targets in the period. For example, in single file, $V_1 = \{-2\delta, 0, 2\delta\}$ and $V_2 = \{-4\delta, -2\delta, 0, 2\delta, 4\delta\}$; in multi file, $V_1 = \{-2\delta, 0, 2\delta\}$ and $V_2 = \{-10\delta, -8\delta, \dots, 8\delta, 10\delta\}$.³¹

To estimate the model empirically, we code each signal $s_{i,t}$ to the closest value of $V_1 = \{-2\delta, 0, 2\delta\}$ denoted by $\hat{s}_{i,t}$, each entry $e_{i,t}$ to the closest value of V_t denoted by $\hat{e}_{i,t}$, and each target $tar_{i,t}$ to $\hat{tar}_{i,t} := \sum_{\tau < t} \sum_i \hat{s}_{i,\tau} + s_{i,t}$. Analogous to the formulation of the logistic response in (2), the likelihood of observing entry $\hat{e}_{i,t} \in V_t$ is therefore described as

$$\hat{\sigma}(\hat{e}_{i,t}; \hat{F}_{i,t}, \lambda) = \frac{\exp(\lambda(\int \pi_{i,t}(\hat{e}_{i,t}, \hat{tar}_{i,t})d\hat{F}_{i,t})}{\sum_{\hat{e} \in V_t} \exp(\lambda(\int \pi_{i,t}(\hat{e}, \hat{tar}_{i,t})d\hat{F}_{i,t})},$$

where $\hat{F}_{i,t}$ describes the belief about the target $\hat{tar}_{i,t} \in V_t$. This belief is constructed recursively as in the previous subsection’s model, such that players employ different precision levels $(\lambda^1, \lambda^2, \dots)$ at different levels of the belief hierarchy.

³⁰We ran simulations to show that the large majority of empirical entries and targets can be suitably approximated on the grid, in the sense that they do not lie outside of the extremal grid values. The proportion of observations that lie outside of the grid’s range is, however, far bigger in the multi-file treatment than in the single-file treatment. The simulations also show that allowing for more than two coin flips per person would improve the approximation only very slowly, yet at a significant cost in terms of complexity.

³¹Formally, in single file, $V_t = \{-2t\delta, -2(t-1)\delta, \dots, -2\delta, 0, 2\delta, \dots, 2(t-1)\delta, 2t\delta\}$, whereas in multi file, $V_t = \{-(4(t-1)+1)2\delta, -4(t-1)2\delta, \dots, -2\delta, 0, 2\delta, \dots, 4(t-1)2\delta, (4(t-1)+1)2\delta\}$.

Tables 7 and 8 present the estimations for single file and multi file, respectively. The first column of each table shows the parameter estimates in a QRE, where all λ values are constant. The second column shows the benchmark case of full randomness, where $\lambda^1 = 0$. Subsequent columns report low-level parameter estimates when setting higher-level parameters to zero. These models generalize Level- k models by permitting probabilistic reasoning at every level of reasoning. For example, the column labelled “[$\lambda^3 = 0$]” gives the estimate of a probabilistic version of BRTNI/Level 2. Up to five parameters can be reliably estimated in both datasets.

Table 7: Single-file estimates of grid-based approximation (standard errors in parentheses)

	Restrictions						
	QRE [$\lambda^1 = \lambda^2 = \dots$]	[$\lambda^1 = 0$]	[$\lambda^2 = 0$]	[$\lambda^3 = 0$]	[$\lambda^4 = 0$]	[$\lambda^5 = 0$]	[$\lambda^6 = 0$]
λ^1	0.33 (.01)	0 (-)	0.19 (.43)	0.24 (.12)	0.26 (.03)	0.27 (.01)	0.26 (.01)
λ^2	0.33 (.01)	-	0 (-)	0.21 (.55)	0.39 (.09)	0.67 (.08)	1.15 (.11)
λ^3	0.33 (.01)	-	-	0 (-)	0.20 (.01)	0.42 (.02)	0.62 (.02)
λ^4	0.33 (.01)	-	-	-	0 (-)	0.20 (.01)	0.37 (.01)
λ^5	0.33 (.01)	-	-	-	-	0 (-)	0.19 (.01)
Log likelihood	-1110.4	-1536.1	-1288.2	-1173.6	-1108.3	-1056.7	- 1020.1

Table 8: Multi-file estimates of grid-based approximation (standard errors in parentheses)

	Restrictions						
	QRE [$\lambda^1 = \lambda^2 = \dots$]	$[\lambda^1 = 0]$	$[\lambda^2 = 0]$	$[\lambda^3 = 0]$	$[\lambda^4 = 0]$	$[\lambda^5 = 0]$	$[\lambda^6 = 0]$
λ^1	0.13 (.01)	0 (-)	0.09 (.17)	0.11 (.14)	0.11 (.05)	0.11 (.01)	0.11 (.01)
λ^2	0.13 (.01)	-	0 (-)	0.31 (.15)	0.54 (.26)	0.68 (.02)	0.74 (.14)
λ^3	0.13 (.01)	-	-	0 (-)	0.26 (.23)	0.51 (.03)	0.62 (.10)
λ^4	0.13 (.01)	-	-	-	0 (-)	0.39 (.02)	0.63 (.12)
λ^5	0.13 (.01)	-	-	-	-	0 (-)	0.52 (.19)
Log likelihood	-1361.9	-1432.5	-1373.8	-1328.5	-1315.7	-1311.3	-1310.3

The estimates confirm the general results of the descriptive analysis and the regressions: in single file, behavior demonstrates a higher depth of reasoning than in multi file. Not only are the parameter estimates larger in single file than in multi file—indicating that participants play closer to best responses, believe that their predecessors play closer to best responses, etc.—higher-level parameters have greater statistical significance in single file than in multi file. Comparing across the columns of Table 8, only λ^1 , λ^2 and λ^3 improve the models’ goodness of fit at high levels of significance, whereas λ^4 and λ^5 are less relevant.³² Moreover, whereas the point estimate $\lambda^3 = 0.26$ in multi file provides evidence of reasoning at a higher level than that employed by BRTNI/Level 2—participants reckon that their predecessors believe their own predecessors play more rationally than pure noise, and hence expect their predecessors’ actions to embed inferences from their own predecessors—its statistical insignificance (and limited magnitude relative to λ^2) suggests a limited role for

³²The standard errors of λ^1 and λ^2 are fairly large, owing to the fact that significant shares of decisions follow BNE and other predictions exactly. We test for the significance of higher levels of reasoning using likelihood-ratio tests. The significance levels of rejecting the indicated model restrictions (against the next more flexible restriction) are $p < 0.01$ for each of the restrictions in single file as well as for $\{[\lambda^1 = \lambda^2 = \dots], [\lambda^2 = 0], [\lambda^3 = 0], [\lambda^4 = 0]\}$ in multi file. For $[\lambda^4 = 0]$ in multi file, the critical level is $p = 0.16$.

such sophistication. Estimating higher-level parameters improves the goodness of fit of the parameter estimates of single file in Table 7 much more substantially. Overall, in multi file, a model with two parameters λ^1 and λ^2 (i.e. with restriction $[\lambda^3 = 0]$) outperforms QRE, which can be interpreted as saying that the noisy version of BRTNI outperforms the noisy version of BNE. In single file, the opposite is true.

6.3 Do people recognize others' mistakes?

In this subsection we focus on a qualitative question that underlies much of the previous analysis: do participants take other people's mistakes—and their collective tendency to overshoot—into account rationally?

One quick answer lies in how participants shade immediate predecessors' entries. To explore this question, we compare the ex-post optimal reaction shading rule to participants' actual shading rule. The quantity $\frac{tar_t - s_t}{e_{t-1}}$ gives the ratio of own target (net of own signal) to last period's entry in the same game. (In the case of multi-file, we take e_{t-1} to denote the average of the four immediate predecessors' entries.) The quantity $\frac{e_t - s_t}{e_{t-1}}$ describes the participant's actual reaction to the previous entry. The two measures vary widely across our participants. However, in both treatments, the median value of the optimal shading factor lies below the median of the actual shading factor: for single file, $\frac{tar_t - s_t}{e_{t-1}} = 0.79$ and $\frac{e_t - s_t}{e_{t-1}} = 1.00$; for multi file, $\frac{tar_t - s_t}{e_{t-1}} = 0.63$ and $\frac{e_t - s_t}{e_{t-1}} = 2.82$. Much like in related analysis at the end of Section 3, we see modest under-shading for single file but vast under-shading for multi file.

These last calculations use only information about how participants responds to their immediate predecessors, and require knowledge of signal realizations unavailable to the participants. We therefore turn to an exploration of whether people's behavior is better described by a Nebi-style model in which they do not heed others' errors or a QRE-style model in which people take others' errors into account. As explained earlier, the immense size

of the signal and action spaces in our model prevent us from estimating logit QRE. One alternative would be to simplify the problem using the grid-based approximation of Section 6.2. For the present question, this has the drawback that it rounds the most extreme entries in the dataset towards zero, penalizing models that correctly predict overshooting of the predecessors' entries.

We offer here a different approach through a variant of QRE, dubbed “pseudo-QRE”, that uses all information in the entries and models each player’s entry as a function of private signal and the history $h_{i,t}$, without truncation. We keep the analysis tractable by assuming that each player deviates from best responding to her predecessors’ actions via a normally distributed error term. Because the players’ private signals in our experiment are approximately normally distributed, this formulation allows us to use standard formulas for Gaussian inference to provide closed-form expressions for behavior.³³

For player i in period $t = 1$, we assume that $e_{i,1}$ is generated by the random choice rule

$$\tilde{e}_{i,1}(h_{i,1}; \tau_\epsilon) = s_{i,1} + \epsilon_{i,1},$$

where $\epsilon_{i,1} \sim N(0, 1/\tau_\epsilon)$, with τ_ϵ being the precision of the error term $\epsilon_{i,1}$ (the inverse of its variance). Like in logit QRE, Player i in pseudo-QRE chooses all actions with positive probability, and better actions with higher probability.³⁴ The parameter τ_ϵ orders sophistication in a similar manner to the QRE parameter λ : as $\tau_\epsilon \rightarrow \infty$, Player i ’s choice converges to the best response $s_{i,1}$; as $\tau_\epsilon \rightarrow 0$, Player i ’s choice becomes an improper uniform distribution on the real line.

³³Let X be the number of heads from 100 flips. Our signals are $S = 2X - 100$, which has $\text{var}(2X - 100) = 4\text{var}(X) = 4(0.5)(0.5)100 = 100$. The random variable S is approximately $N(0, 100)$.

³⁴Payoffs in our experiment are isomorphic to the maximum of the absolute distance between entry and target and a constant. It is well known that choosing the median of a distribution minimizes the expected absolute distance between estimator and target. It is straightforward to show that for normally distributed random variables (or, more generally, symmetric and unimodal densities), choosing the median (and, hence, mean) maximizes expected payoff in our experiment, and that payoff decreases in distance from the mean. As a consequence, normally distributed entries around the mean satisfy the better-reply property of QRE: better actions get played with higher probabilities.

Player j acting in $t = 2$ recognizes that player i 's choice process includes a random component and chooses according to

$$\tilde{e}_{j,2}(h_{j,2}; \tau_\epsilon) = E[s_{i,1}|e_{i,1} = \tilde{e}_{i,1}(h_{i,1}; \tau_\epsilon)] + s_{j,2} + \epsilon_{j,2} = \gamma_\epsilon e_{i,1} + s_{j,2} + \epsilon_{j,2},$$

where $\epsilon_{j,2} \sim N(0, 1/\tau_\epsilon)$, and $\gamma_\epsilon := \frac{\tau_\epsilon}{\tau_\epsilon + 1/100}$ is the formula for normal inference under the assumption that $s_{i,1} \sim N(0, 100)$ (which is approximately true) and $\epsilon_{i,1}$ has precision τ_ϵ . Just as in logit QRE, in a pseudo-QRE Player j correctly perceives the randomness in Player i 's choice. The fact that $\gamma_\epsilon \in (0, 1)$ expresses that players who recognize that their opponents make errors shade their interpretation of their predecessors' entries towards zero.

Player k in $t = 3$ of the single-file game views both of his predecessors as choosing with precision τ_ϵ and views Player j as taking i 's precision to be τ_ϵ :

$$\begin{aligned} & \tilde{e}_{k,3}(h_{k,3}; \tau_\epsilon) \\ &= E[s_{i,1}|e_{i,1} = \tilde{e}_{i,1}(h_{i,1}; \tau_\epsilon)] + E[s_{j,2}|e_{j,2} = \tilde{e}_{j,2}(h_{j,2}; \tau_\epsilon)] + s_{k,3} + \epsilon_{k,3} \\ &= \gamma_\epsilon(1 - \gamma_\epsilon)e_{i,1} + \gamma_\epsilon e_{j,2} + s_{k,3} + \epsilon_{k,3}, \end{aligned}$$

where $\epsilon_{k,3} \sim N(0, 1/\tau_\epsilon)$.

Iterating gives

$$\tilde{e}_{i,t}(h_{i,t}; \tau_\epsilon) = \gamma_\epsilon \sum_{j=1}^{t-1} (1 - \gamma_\epsilon)^{t-1-j} e_{i,j} + s_{i,t} + \epsilon_{i,t}, \quad (3)$$

where again $\gamma_\epsilon = \frac{\tau_\epsilon}{\tau_\epsilon + 1/100}$. The first row of Table 9 describes the maximum-likelihood estimation of this model. The estimated precision corresponds to the inference coefficient $\gamma_\epsilon = 0.022$, which demonstrates that players choose actions very noisily: in a BNE, the parameter would be 1.

Table 9: Sophisticated versus naive Nebi in single file.
(Standard error in parentheses.)

Model	$\sqrt{\frac{1}{\tau_\epsilon}}$	Log likelihood
Pseudo-QRE	65.37 (2.25)	-2707
Noisy Nebi	45.62 (1.47)	-2515

We can now test pseudo-QRE against our model of naïve Bayesians (Nebi). Nebi predicts that $\gamma_\epsilon = 1$ and $\tau_\epsilon = \infty$: Nebis do not err. Allowing for finite values of τ_ϵ —whilst restricting $\gamma_\epsilon = 1$ —enriches Nebi by possible errors in best response but retains the property that other players are believed to play optimally. Both this “noisy Nebi” model and the pseudo-QRE model have one degree of freedom. Consistent with our earlier finding in Table 1 that 81% of subjects behave exactly as Nebi, Table 9 shows that noisy-Nebi fits the data far better than pseudo-QRE. Participants in single file come closer to completely overlooking the limits to their predecessors’ sophistication than they do to fully accounting for those limitations.

For the multi-file treatment, the formula for pseudo-QRE is almost identical:

$$\tilde{e}_{i,t}(H_{i,t}; \tau_\epsilon) = \gamma_\epsilon \sum_{j=1}^{t-1} (1 - 4\gamma_\epsilon)^{t-1-j} \bar{e}_j + s_{i,t} + \epsilon_{i,t},$$

where $\bar{e}_j = \sum_{i=1}^4 e_{ij}$ and $\gamma_\epsilon = \frac{\tau_\epsilon}{\tau_\epsilon + 1/100}$.

Once more we can compare the goodness of fit of pseudo-QRE to that of noisy Nebi. The first row in Table 10 indicates that the maximum-likelihood estimate of pseudo-QRE corresponds to the inference parameter $\gamma_\epsilon = 0.0060$, less than one-third its value in single file. Just like in all previous analyses, we observe a treatment difference: the comparison of the two models’ likelihoods shows that in multi file, pseudo-QRE fits behavior better than does noisy Nebi, contrary to the corresponding result of single file. In multi file, accounting

for others’ errors is far more widespread than believing others to be Bayes-rational.

Table 10: Sophisticated versus naive Nebi in multi file.
(Standard error in parentheses.)

Model	$\sqrt{\frac{1}{\tau_\epsilon}}$	Log likelihood
Pseudo-QRE	128.86 (0.58)	-3679
Noisy Nebi	245.73 (7.74)	-4004

7 Conclusion

Our experiments are designed to separate possible errors in inference that one may make when observing others’ actions from possible unrelated errors in Bayesian updating. We find considerable amounts of inference errors, but their prevalence and importance differ between our two treatments. In the single-file treatment, most participants behave in a manner consistent with BNE, and they benefit from learning from others. Nevertheless, collectively they exhibit statistically significant degrees of excessive imitation. In the multi-file treatment, participants engage in substantial over-imitation that produces outcomes dramatically different from BNE predictions. Participants commit inference errors in such abundance that the average participant would earn more if she did not have the opportunity to learn from others’ behavior and simply entered her signal.³⁵

We attribute the deviations from optimality mainly to redundancy neglect, through which people fail to appreciate that their predecessors already incorporate prior observations. As perhaps redundantly discussed in several earlier sections, this type of behavior can take many different precise forms. Eyster and Rabin (2010) model the extreme version of BRTNI players

³⁵Similar to our finding of treatment variation, experiments by Esponda and Vespa (2014), Albert, Costa-Gomes and Weizsäcker (2017) and Enke (2017) show that increasing the cognitive burden of their choice tasks reduces people’s ability to account for selection in Bayesian Games.

who fully neglect that their predecessors' actions incorporate inferences made from their own predecessors; BRTNIs interpret every predecessor's action at face value, as reflecting that player's private information alone. This prediction dovetails with prior experimental literature on social-learning games, in which taking others' play at face value features as one of the most discussed behavioral patterns, alongside a pattern of too frequently following one's own signal. (See, *inter alia*, discussions by Kübler and Weizsäcker (2004) and March and Ziegelmeyer (2015) on the connections and interplay of the two effects.) The finding that people fail to think through how others think through still others' behavior also relates to large body of evidence from many different contexts on higher-order reasoning (see, e.g., the early contributions of Stahl and Wilson (1994) and Nagel (1995)).

Our experimental evidence adds to the discussion in two ways. First, our experiments illustrate how errors in higher-order reasoning can lead people to neglect certain correlations. Enke and Zimmermann (2015) provide evidence that experimental participants neglect the correlation in signals when these signals draw upon a common source. By ignoring the commonality of the underlying source of information, participants in their experiment double count that source, similar to how our experimental participants, by neglecting the redundancy in their predecessors' behavior, double count earlier participants' actions.

Second, our experiment sheds light on several solution concepts in the behavioral/experimental literature that it was not designed to explore. For this reason, our experiment might prove especially informative: we designed the game because of its economic importance and the importance of a type of error in reasoning that only partly corresponds to these general solution concepts, rather than to vindicate or bash any one of them. Variants of redundancy neglect that are weaker than BRTNI/Level 2 may clearly better explain the behavior—and its consequences—than BNE, yet it is noteworthy that overall BNE clearly outperforms any variant of cognitive-hierarchy or Level- k models that we are aware of. Out of hundreds of choices observed, only once did behavior match Level 3 when failing to match Level 1 or Level 2. Indeed, far more behavior matched Level 2 than Level 1 or Level 2, although of

course we think that Level 22 is not the right conceptualization. The data patterns and the analysis in Section 6 also suggest that in the single-file game, models which incorporate decision noise such as Quantal Response Equilibrium or noise-enhanced versions of cognitive-hierarchy models like Camerer, Ho and Chong (2004) offer little improvement over BNE.³⁶ Not only do these models miss systematic deviations made out of proportion to the errors' low costliness, but their very *raison d'être*—the (generally quite compelling) notion that players who make mistakes may optimize with respect to other players' mistakes—turns out backwards here because our many Nebi participants very often fail to take into account their predecessors' over-counting (or other mistakes). As much as all of these models deserve credit for improving fit in many games, it may also be worth noting examples such as our single-file game in which the enhanced solution concepts offer worse predictions than traditional solution concepts.

³⁶In a related discussion, Goeree, Palfrey, Rogers and McKelvey (2007) show how random noise can help overcome informational externalities in social-learning games and how, therefore, QRE leads to greater efficiency than BNE in traditional coarse-action settings. In our settings, because BNE yields the first best, QRE can only lessen efficiency.

A Additional Tables and Graphs

The following figures correspond to Figures 1, 2, 4 and 5, but depict behaviour in the first, second, and third games of each session separately. In the single-file treatment, they show that entries are more extreme in the first two games of each session than in the final game; in the multi-file treatment, they show no discernible pattern.

The first three figures in this appendix should be compared to Figure 2.




Figure 7: Single-file mean absolute entries and targets for first games




Figure 8: Single-file mean absolute entries and targets for second games




Figure 9: Single-file mean absolute entries and targets in third games

The next three figures in this appendix should be compared to Figure 3.




Figure 10: Single-file median absolute entries and targets for first games




Figure 11: Single-file median absolute entries and targets for second games




Figure 12: Single-file median absolute entries and targets in third games

The next three figures should be compared to Figure 4.




Figure 13: Multi-file mean absolute entries and targets for first games

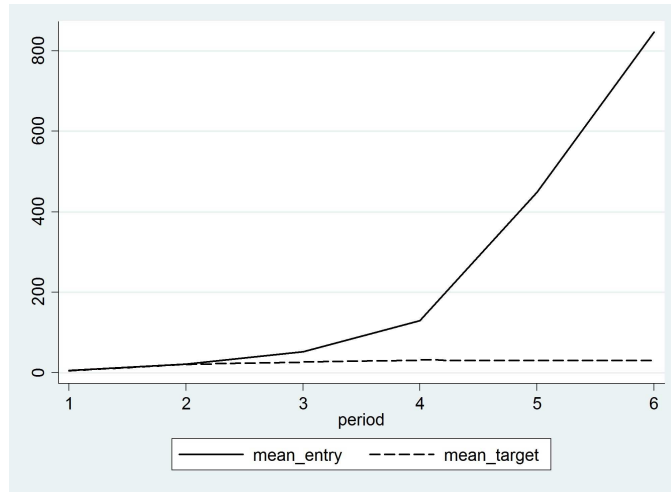


Figure 14: Multi-file mean absolute entries and targets for second games




Figure 15: Multi-file mean absolute entries and targets in third games

The next three figures should be compared to Figure 5.




Figure 16: Multi-file median absolute entries and targets for first games




Figure 17: Multi-file median absolute entries and targets for second games




Figure 18: Multi-file median absolute entries and targets in third games

Figures 19 and 20 show the distributions of the deviations from target, such that higher values indicate higher deviations in the same direction as BRTNI's deviation from target.




Figure 19: Single-file entry minus target. Positive values indicate deviations whose sign matches that of BRTNI's deviation. For scaling purposes, the 127 entries on target (zero deviation), and six outliers in the right tail, are not depicted.




Figure 20: Multi-file entries off target. Positive values indicate deviations whose sign matches that of BRTNI's deviation. For scaling purposes, the 172 entries on target (zero deviation), 56 outliers in the right tail, and 4 outliers in the left tail are not depicted.

Table 11 lists the frequencies of cases where a model's simple version is inconsistent with an observation but Table 1 lists the observation as consistent (cf Footnote 15).

Table 11: Single-file entries inconsistent with a model's simple version but consistent with an expanded version

Period	Cursed	BRTNI	Level 3	Nebi	NA
	($k = 1$)	($k = 2$)	($k = 3$)	($k > t - 1$)	
t=1	0	0	0	0	0
t=2	0	0	0	0	0
t=3	0	0	0	0	0
t=4	0	0	0	0	0
t=5	0	0	0	0	0
t=6	0	0	0	0	0
t=7	0	0	0	0	0
t=8	0	0	0	0	0
t=9	0	0	0	0	0
t=10	0	0	0	0	0
t=11	0	0	0	0	0
t=12	0	0	0	0	0
t=13	0	1	0	0	0
t=14	0	1	1	0	0
t=15	0	0	0	0	0
t=16	0	0	0	0	1
t=17	0	1	0	0	0
t=18	0	1	0	0	1
t=19	0	0	1	0	0
t=20	0	0	0	0	0
t=21	0	0	0	0	0
t=22	0	1	0	0	1
t=23	0	1	0	0	1
t=24	0	0	1	0	1
Total	0	6	3	0	5

The following tables are the analogues of Table 1, broken down by game.

Table 12: Single-file entries consistent with different behavioral models for first games

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	NA
t=1	7	7	7	7	0
t=2	1	5	5	5	1
t=3	0	4	3	3	0
t=4	0	4	0	3	0
t=5	1	0	0	5	1
t=6	1	1	1	4	1
t=7	1	1	0	5	0
t=8	0	1	0	5	1
t=9	1	1	0	5	0
t=10	0	1	0	4	2
t=11	2	1	0	2	2
t=12	1	0	0	3	3
t=13	2	0	0	5	1
t=14	1	0	0	4	3
t=15	0	0	0	6	1
t=16	2	0	0	6	0
t=17	0	0	0	7	0
t=18	1	0	0	3	3
t=19	1	1	0	6	0
t=20	0	0	0	5	2
t=21	0	0	0	4	3
t=22	1	0	0	5	2
t=23	1	0	0	6	1
t=24	1	1	0	5	1
Total in $t \geq 3$	17	16	4	101	27
Total in $t \geq 4$	17	12	1	98	27
Total	25	28	16	113	28

Table 13: Single-file entries consistent with different behavioral models for second games

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	NA
t=1	7	7	7	7	0
t=2	0	6	6	6	1
t=3	1	2	6	6	0
t=4	0	1	0	5	2
t=5	0	1	0	6	0
t=6	1	0	0	6	0
t=7	0	0	0	6	1
t=8	1	0	0	7	0
t=9	0	1	0	7	0
t=10	0	0	0	4	3
t=11	1	0	0	4	2
t=12	0	0	0	6	1
t=13	1	0	0	6	0
t=14	1	0	0	5	2
t=15	0	0	0	6	1
t=16	0	0	0	4	3
t=17	0	0	0	7	0
t=18	1	0	0	6	1
t=19	0	0	0	7	0
t=20	0	0	0	6	1
t=21	0	0	0	5	2
t=22	0	0	0	4	3
t=23	2	0	0	5	2
t=24	1	1	0	6	0
Total in $t \geq 3$	10	6	6	124	24
Total in $t \geq 4$	9	4	0	118	24
Total	17	19	19	137	25

Table 14: Single-file entries consistent with different behavioral models in third games

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	NA
t=1	7	7	7	7	0
t=2	0	7	7	7	0
t=3	0	0	5	5	2
t=4	1	0	0	7	0
t=5	0	1	0	6	0
t=6	0	0	0	5	2
t=7	0	0	0	6	1
t=8	0	0	0	5	2
t=9	0	0	0	6	1
t=10	0	0	0	6	1
t=11	0	0	0	5	2
t=12	0	0	0	5	2
t=13	1	0	0	5	2
t=14	2	1	0	5	0
t=15	0	0	0	3	4
t=16	1	0	0	6	0
t=17	1	1	0	6	0
t=18	0	0	0	6	1
t=19	0	0	0	6	1
t=20	0	0	0	7	0
t=21	0	0	0	3	4
t=22	0	0	0	6	1
t=23	0	0	0	6	1
t=24	0	0	0	6	1
Total in $t \geq 3$	6	3	5	121	28
Total in $t \geq 4$	6	3	0	116	26
Total	13	17	19	135	28

Table 15 complements Table 1 and the previous three tables by listing the numbers of subjects for whom n of 6 entries coincide exactly with the prediction of the listed models, where $n = 0, 1, \dots, 6$.

Table 15: Number of single-file subjects whose entries coincide exactly with predictions of different behavioural models for exactly n of 6 games

	Cursed	BRTNI	Level 3	Nebi
	$(k = 1)$	$(k = 2)$	$(k = 3)$	$(k > t - 1)$
$n = 6$	0	0	0	32
$n = 5$	0	0	0	25
$n = 4$	0	1	0	9
$n = 3$	3	1	0	5
$n = 2$	9	9	10	6
$n = 1$	28	39	34	5
$n = 0$	44	34	40	2
Total	84	84	84	84

Table 16 lists the frequencies of cases where a model's simple version is inconsistent with an observation but Table 3 lists the observation as consistent (cf Footnote 15).

Table 16: Multi-file entries inconsistent with a model's simple version but consistent with an expanded version

Period	Cursed	BRTNI	Level 3	Nebi	NA
	$(k = 1)$	$(k = 2)$	$(k = 3)$	$(k > t - 1)$	
t=1	0	0	0	0	0
t=2	0	0	0	0	0
t=3	0	0	0	0	0
t=4	0	1	1	0	0
t=5	0	1	3	0	0
t=6	0	0	1	0	0
Total	0	2	4	0	0

The following tables are the analogues of Table 3, broken down by first, second, and third games.

Table 17: Multi-file entries consistent with different behavioral models in first games

Period	Cursed	BRTNI	Level 3	Nebi	NA
	($k = 1$)	($k = 2$)	($k = 3$)	($k > t - 1$)	
t=1	28	28	28	28	0
t=2	3	21	21	21	4
t=3	3	9	1	1	18
t=4	4	9	0	1	15
t=5	0	7	0	2	19
t=6	0	3	0	0	25
Total in $t \geq 3$	7	28	1	4	77
Total in $t \geq 4$	4	19	0	3	59
Total	38	77	50	53	81

Table 18: Multi-file entries consistent with different behavioral models in second games

Period	Cursed	BRTNI	Level 3	Nebi	NA
	($k = 1$)	($k = 2$)	($k = 3$)	($k > t - 1$)	
t=1	27	27	27	27	1
t=2	0	20	20	20	8
t=3	0	4	12	12	12
t=4	0	7	0	5	17
t=5	1	13	0	2	13
t=6	0	4	0	1	23
Total in $t \geq 3$	1	28	12	20	65
Total in $t \geq 4$	1	24	0	8	53
Total	28	75	59	67	74

Table 19: Multi-file entries consistent with different behavioral models for third games

Period	Cursed	BRTNI	Level 3	Nebi	NA
	$(k = 1)$	$(k = 2)$	$(k = 3)$	$(k > t - 1)$	
t=1	26	26	26	26	2
t=2	0	26	26	26	2
t=3	0	5	16	16	7
t=4	0	6	0	10	12
t=5	0	9	0	2	17
t=6	1	1	0	1	25
Total in $t \geq 3$	1	21	16	29	61
Total in $t \geq 4$	1	16	0	13	54
Total	27	73	68	81	65

Table 20 complements Table 3 and the previous three tables by listing the numbers of subjects of whose six entries exactly x are consistent with each of the listed models, where $n = 0, 1, \dots, 6$.

Table 20: Number of multi-file subjects whose entries are consistent with different behavioral models for exactly n games

	Cursed	BRTNI	Level 3	Nebi
	$(k = 1)$	$(k = 2)$	$(k = 3)$	$(k > t - 1)$
$n = 6$	0	7	0	0
$n = 5$	0	6	0	4
$n = 4$	0	10	0	9
$n = 3$	4	14	32	24
$n = 2$	22	27	31	27
$n = 1$	37	18	19	18
$n = 0$	21	2	2	2
Total	84	84	84	84

Figure 21 is analogous to Figure 6 but uses the single-file data. It shows the between-game average entries and targets in single file and by period, separated out according to the sign of the sum of the game's first four signals. Apart from the episode in Game 34, entries do not discernibly deviate from target. Figure 21 shows that the sign of early signals does not cause systematic deviations from target.




Figure 21: Average entry and target in single file, separated out by the sum of the game's first four signals.

The following table shows results analogous to those of Table 5, for single file.

Table 21: Single-file regressions of period- t entries on current signals and entries in past super-periods. Standard errors in parentheses are clustered by session.

	$e_{\bar{1}}$	$e_{\bar{2}}$	$e_{\bar{3}}$	$e_{\bar{4}}$	$e_{\bar{5}}$	$e_{\bar{6}}$
$\bar{e}_{\bar{1}}$	-	0.659 (.228)	-0.187 (.484)	0.401 (.329)	-0.109 (.143)	0.364 (.291)
$\bar{e}_{\bar{2}}$	-	-	1.068 (0.493)	-0.155 (.119)	0.198 (.511)	-0.175 (.136)
$\bar{e}_{\bar{3}}$	-	-	-	0.043 (.026)	-0.030 (.007)	0.122 (.020)
$\bar{e}_{\bar{4}}$	-	-	-	-	0.213 (.012)	-0.297 (.129)
$\bar{e}_{\bar{5}}$	-	-	-	-	-	0.105 (.166)
s_t	0.782 (.155)	1.531 (0.441)	-3.611 (3.594)	0.741 (.474)	0.710 (.125)	0.427 (.642)
R^2	0.18	0.51	0.48	0.19	.89	0.67
obs	84	84	84	84	84	84

BRTNI predicts that all coefficients in the table should exceed one. For example, a BRTNI player at $t = 5$ should play

$$e_5 = \sum_{\tau < 5} e_\tau + s_5 = e_{\bar{1}} + s_5.$$

A BRTNI player at $t = 6$ should play

$$e_6 = \sum_{\tau < 6} e_\tau + s_6 = e_{\bar{1}} + e_5 + s_6 = 2e_{\bar{1}} + s_5 + s_6,$$

and so forth. Since s_t , for $t = 5, \dots, 8$, is uncorrelated with $e_{\bar{1}}$, regressing the observations $\{e_t\}_{t=5}^8$ on $e_{\bar{1}}$ and own signal yields a predicted coefficient larger than one; the same pattern holds for later entries. Hence, Table 21 illustrates once again that any redundancy neglect in single file takes a far milder form than that expressed by BRTNI.

References

- Albert, Philipp, Miguel A. Costa-Gomes, and Georg Weizsäcker**, “Cursed bets on markets,” Mimeo 2017.
- Anderson, Lisa R. and Charles A. Holt**, “Information Cascades in the Laboratory,” *American Economic Review*, 1997, 87 (5), 847–862.
- Avery, Christopher and John Kagel**, “Second-Price Auctions with Asymmetric Payoffs: An Experimental Investigation,” *Journal of Economics & Management Strategy*, 1997, 6 (3), 573–603.
- Banerjee, Abhijit**, “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 1992, 107 (3), 797–817.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch**, “A Theory of Fads, Fashion, Custom and Cultural Change as Information Cascades,” *Journal of Political Economy*, 1992, 100 (5), 992–1026.
- Brunner, Christoph and Jacob K. Goeree**, “Wise Crowds or Wise Minorities?,” Mimeo 2012.
- Callender, Steven and Johannes Hörner**, “The Wisdom of the Minority,” *Journal of Economic Theory*, 2009, 144 (4), 1421–1439.
- Camerer, Colin, Teck-Hua Ho, and Juin-Kuan Chong**, “A Cognitive Hierarchy Model of Games,” *Quarterly Journal of Economics*, 2004, 119 (3), 861–898.
- Çelen, Bogaçhan and Shachar Kariv**, “An Experimental Test of Observational Learning under,” *Economic Theory*, 2005, 26, 677–699.

- Chandrasekhar, Arun, Horacio Larreguy, and Juan Pablo Xandri**, “Testing Models of Social Learning on Networks: Evidence from a Lab Experiment in the Field,” Mimeo 2017.
- Choi, Syngjoo, Douglas Gale, and Shachar Kariv**, “Social learning in networks: a Quantal Response Equilibrium analysis of experimental data,” *Review of Economic Design*, 2012, 16, 135–157.
- , **Eduardo Gallo, and Shachar Kariv**, “Networks in the Laboratory,” Mimeo 2015.
- Crawford, Vincent and Nagore Iriberri**, “Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner’s Curse and Overbidding in Private-Value Auctions?,” *Econometrica*, 2007, 75 (6), 1721–1770.
- DeGroot, M.H.**, “Reaching a Consensus,” *Journal of the American Statistical Association*, 1974, 69 (345), 118–121.
- DeMarzo, Peter, Dimitri Vayanos, and Jeffrey Zwiebel**, “Persuasion Bias, Social Influence, and Unidimensional Opinions,” *Quarterly Journal of Economics*, 2003, 118 (3), 909–968.
- Enke, Benjamin**, “What you see is all there is,” Mimeo 2017.
- and **Florian Zimmermann**, “Correlation Neglect in Belief Formation,” *Mimeo*, 2015.
- Esponda, Ignacio and Emanuel Vespa**, “Hypothetical thinking and information extraction in the laboratory,” *American Economic Journal: Microeconomics*, 2014, 6 (4), 180–202.
- Eyster, Erik and Matthew Rabin**, “Cursed Equilibrium,” *Econometrica*, 2005, 73 (5), 1623–1672.

- **and** — , “Naïve Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2010, *2* (4), 221–243.
- **and** — , “Extensive Imitation Is Irrational and Harmful,” *Quarterly Journal of Economics*, 2014, *129* (4), 1861–1898.
- Fischbacher, Urs**, “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 2007, *10*, 171–178.
- Gale, Douglas and Shachar Kariv**, “Bayesian Learning in Social Networks,” *Games and Economic Behavior*, 2003, *45* (4), 329–346.
- Goeree, Jakob, Thomas Palfrey, Brian Rogers, and Richard McKelvey**, “Self-Correcting Information Cascades,” *Review of Economic Studies*, 2007, *74* (3), 733–762.
- Golub, Benjamin and Matthew O. Jackson**, “Naïve Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2010, *2* (1), 112–149.
- Kübler, Dorothea and Georg Weizsäcker**, “Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory,” *Review of Economic Studies*, 2004, *71* (2), 425–441.
- Lee, In-Ho**, “On the Convergence of Informational Cascades,” *Journal of Economic Theory*, 1993, *61*, 395–411.
- March, Christoph and Anthony Ziegelmeyer**, “Even Experts Follow Large Herds: Naïve Herding in the Laboratory,” *Mimeo*, 2015.
- McKelvey, Richard and Thomas Palfrey**, “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, 1995, *10*, 6–38.

Nagel, Rosemarie, “Unraveling in Guessing Games: An Experimental Study,” *American Economic Review*, 1995, 85 (5), 1313–1326.

Stahl, Dale O. II and Paul W. Wilson, “Experimental evidence on players’ models of other players,” *Journal of Economic Behaviour & Organization*, 1994, 25 (3), 309–327.

Weizsäcker, Georg, “Do we follow others when we should? A simple test of rational expectations,” *American Economic Review*, 2010, 100, 2340–2360.