An Experiment on Social Mislearning

Erik Eyster, Matthew Rabin and Georg Weizsäcker*

December 18, 2015

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 comments given by audiences at CEU, Cologne, DIW Berlin, Maastricht, Pittsburgh, and Tel Aviv University, for 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 (Starting Grant 263412).

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 whom they observe.

Accounting for redundancy proves challenging even in settings devoid of rational anti-imitation. Experimental evidence by Kübler and Weizsäcker (2004) and others demonstrates such failure. Even pre-dating the experimental evidence, doubts about whether people fully adjust for redundancy motivated researchers to develop models of *redundancy neglect*. DeMarzo, 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 in simple herding contexts 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 that are designed to be conducive to efficient learning when rationality is common knowledge, yet amenable to the detection of redundancy neglect when participants succumb to it. 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 the two treatments, each participant privately observes an integer (or "signal") as well as the public entries of all preceding participants. She then makes an entry herself and gets paid for entering a number as close as possible 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 mimic realistic challenges people might face when applying the logic of social learning, without demanding complicated math or inference of those who understand the logic.

The lessons learned about rationality and redundancy neglect are varied. We identify relatively mild effects of redundancy neglect in one treatment, and more redundancy neglect with much more severe consequences in the other treatment. In the first treatment, where—as in previous experiments—participants move single-file, 75% of participants employ their Bayesian-Nash-Equilibrium (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. This leads the 75% doing the "right" thing to over-imitate; they would be better off by down-weighting their immediate 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, overinfluence of initial movers and long-run extreme beliefs, appear neither strongly nor statistically significantly. In our second treatment, even mild redundancy neglect can be very socially harmful. Here, participants move four at a time, which creates a large set of informational redundancies (like many natural settings, but unlike all previous experimental settings that we know of), such that BNE predicts frequent anti-imitation. Deviations from these strategies are far more common. Early signals are heavily over-counted and push later entries towards costly extremities. On average, participants lose from social learning—they would do better by ignoring others' entries and announcing only their own signals. Nevertheless, we also uncover evidence of complex rational behavior: almost 16% of moves requiring anti-imitation conform to BNE. Although anti-imitation occurs infrequently, and although our experiments vindicate prior theoretical findings that very little over-imitation suffices to lead societies seriously astray, this is the first evidence we know of where people rationally anti-imitate.[1]

We specified few 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 analysis, 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 which we did hypothesize. BRTNI play requires that a player in period t add 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 that should appear implausible to participants. 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 weighting of all previous entries whose weights 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 (at best) weakly confirmed in

[1] We cannot fail to be impressed by the two participants who perfectly follow the BNE strategy in period 6, which calls for summing the entries of the four immediate predecessors, before subtracting off 3 times the entries of the four prior movers, adding 9 times the entries of the four players prior, subtracting off 27 times the entries of the four prior players, and finally adding 81 times the entries of the first four movers. (However, they may have arrived at the correct answer through the conceptually simpler procedure of iteratively determining signals: each period-1 signal equals entry; each period-2 signal equals entry minus the sum of period-1 signals; each period-3 signal equals entry minus the sum of period-1 and period-2 signals; etc. Unfortunately, given other participants' behavior, this strategy fares poorly.

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 added to her immediate predecessor's entry. When done correctly by all players, every action equals its target, and all players earn the maximum payoff. For ease of reference, we shall refer to players who play this "naive Bayesian" strategy as *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 BRNTI 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) and Nagel's (1995) 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 make errors—including, for example, sign errors when reporting intended entries or when adding their private signals. 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. Across all decisions in this treatment, participants would earn more by adding their signal to a number 31% closer to zero than number apparently used. A player whose predecessors are not Nebis does not maximize expected payoff by being Nebi; in fact, the 75% of participants who are Nebi would do better shading by 10%. 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 only deliver 71% of the maximum payment.[2]

Section 4 summarizes the results in the second, multi-file treatment. Here, BNE reasoning once again demands only simple arithmetic of the players (and once again predicts full efficiency), but it involves peculiar-looking behavior. In period 1, optimal behavior remains

[2]Yet Eyster 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 effects. That is, the theory that motivates this experiment and finds support in its data predicts severe long-run costs from even the mildest of over-counting.

trivial—a participant must just enter her signal. In period 2, she should just 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.[3] Surprisingly, we find 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 outnumbers Nebi 3:2. Behavior corresponds less well overall 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 veer off in the direction predicted by BRTNI. Participants make entries with much greater magnitude than those predicted by BNE, and on average they would earn more by strongly shading their interpretation of prior entries towards zero—optimally shading by 98%. Moreover, participants earn less than they would by relying purely on their private signals, even including period 2 in which participants clearly benefit from observing the first movers.

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, whereas redundancy neglect predicts

[3]The precise formula (see earlier footnote) is of course tied to the particular situation, but we once again emphasize that the BNE prediction of anti-imitation is not an artifact of clever or sinister experimental design, but rather inherent to almost all non-single-file settings.

that players in period t should implicitly weight signals of periods $t-2, t-3, \ldots$ with coefficients larger than one. Although the regression analysis in the single-file treatment paints a very mixed but more positive picture of BNE, the analysis in the multi-file case confirms the redundancy neglect portrayed by the descriptive analysis. Nine out of the ten point estimates of these coefficients exceed 1, with estimates ranging from 3.0 to 19.6. In the 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.

Putting our approach in perspective of the literature, we note that our experiments differ in at least three ways from the standard experimental set-up developed by Anderson and Holt (1998) 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 to generate identical responses. Third, and most important from an economics perspective, we move away from the traditionally studied but misleading narrow band of herding settings where the behavior predicted by BNE closely resembles a less-rational tendency to imitate, and where welfare costs of over-imitation are limited. By disentangling rational imitation from irrational over-imitation, our experiment better informs our understanding of many social-learning tasks than settings that have until now been the focus of experimental research.

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 will suggest that most initial departures from BNE predictions cannot 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 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.

We conclude the paper in Section 6. We first 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.) Rather, our aim is to shift focus from very special settings where BNE happens to be difficult to distinguish from intuitive imitative behavior towards the far more common settings in which BNE predicts behavior different from intuitive 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 6 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 $t = 1, ..., T$, 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 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

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

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 excludes signals of the other three participants moving concurrently.[4] The payoff function penalizes deviations from the target in a linear fashion up to the point where a participant's entry lies 96 away from the target, beyond which there is no punishment for further error.[5] Upon completion of each period, all participants receive an updated list of all previous entries.

Our 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. 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, our experimental setting approximates a situation with a binary state about which

[4]Including these three signals would not change the optimal strategy in the game.

[5]Despite facing flat incentives, participants whose guesses veer way off target lack any obvious alternative strategy to simply providing their best guess. 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.

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. This "wallet-game" signal structure has been tested in other experimental settings (as a 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.[6] 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.[7] In seven of the 14 sessions, participants play the single-file game, and the remaining 7 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. Because participants receive no feedback about the true value of the target until after all

[6]In 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.

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

decision-making, the experimental design does little to promote learning across the three games per player and precludes us from addressing learning across repetitions. Nevertheless, Appendix A includes additional results that separate the data between the three games per session.

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 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 actions 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 be 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. 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 whose magnitudes exceed 100 could not possibly reflect private signals alone. For neither case did we formulate hypotheses on how 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 1 and 2 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.[8]

Figure 1: Single-file mean absolute entries and targets

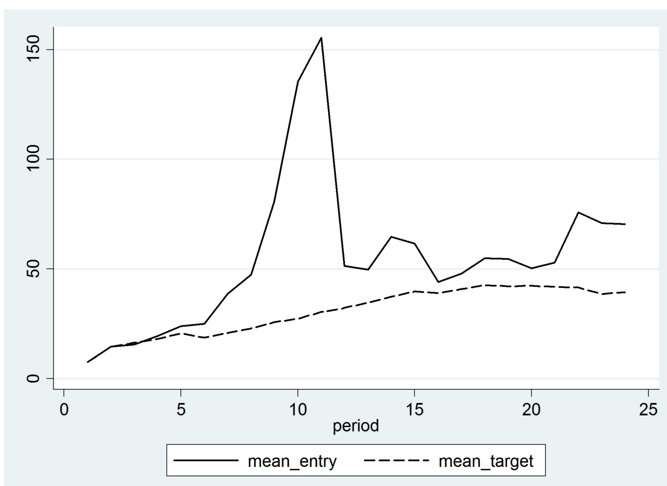

[8]By using absolute values of each entry, the figures treat positive and negative entries symmetrically. One may worry about a potential bias towards making positive-valued entries, but such a bias cannot be discerned in our data. In the single-file treatment, the random signals happen to be negative (49%) more often than 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 negative.

Figure 2: Single-file median absolute entries and targets

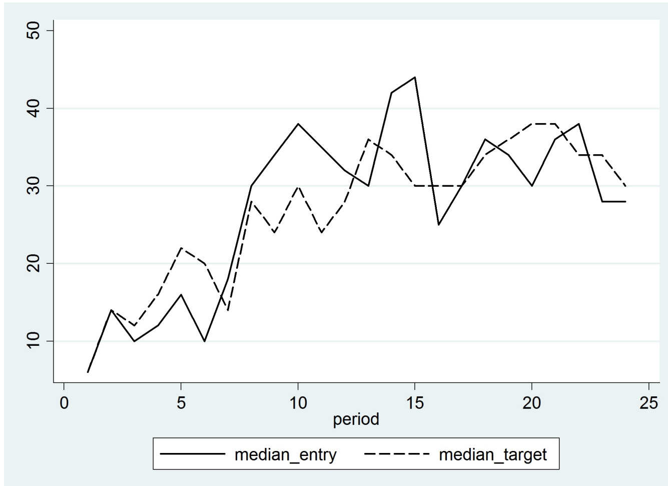

Figure 1 shows that after a few periods in which the mean entry nearly coincides with the mean target, a large disparity emerges before swiftly vanishing. The underlying fluctuation of entries derives solely from one game, Game 34. In it, several participants (those in periods 3, 4, 7, 9, 10) act like 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 2 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, namely 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.[9]

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. In our games,

[9] In 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).

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.[10]

[10]Because 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.

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	Other
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

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

In each of the first three periods, where 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 "Other" 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.[11] The single-file game is so simple that participants appear to understand the logic of BNE—they simply add their own signal to the previous period's entry—and can apply it even 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, for late movers the Nebi strategy is not empirically optimal. Indeed, in the last 21 rounds, while 75% follow their Nebi strategy, only 17% hit their target—because the play has previously departed from Nebi play. While playing Nebi may be arithmetically simple, best-responding to predecessors who play differently presents greater challenges.

Our best guess is that exactly zero participants think in terms of Level 3 in this game. Of the 504 choices, 54 match Level 3. But 53 of those occur in $t = 1$ and $t = 2$, where Level-3 play coincides with Level-2/BRTNI and Nebi play, or in $t = 3$, where it agrees with

[11]Nebi'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.

Nebi. Of the 441 choices beginning in $t = 4$, Nebi play appears 332 times, and Level-3 just once.[12]

The table 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 the point prediction of BRTNI often gets the qualitative difference from BNE right: of the 377 entries off target, 270 (72%) veer off in the direction of BRTNI.[13] Figure 3 shows the distribution of entries relative to their targets. Even apart from its outliers, the figure illustrates an asymmetry in the deviation from target, with a tendency towards BRTNI (for scaling purposes, the figure excludes six outliers in the right tail that deviate in the same direction as BRTNI).

Altogether, the evidence in the single-file game suggests that most participants employ the Nebi strategy. Thus, 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 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. 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 $s_{i,t}$ as $e_{i,t} = s_{i,t} + (e_{i,t} - s_{i,t})$;

[12] The 332 surely exaggerates how many people were thinking through the logic of BNE in choosing their actions. But even the single hit of Level 3 is exaggerated: it occurs in Round 6 of Game 40, where the Level-3 prediction coincidentally matches the player's signal. Because this same participant later on in Round 18 also plays her own signal, where it differs from the Level-3 action, she appears to be cursed (Level 1) instead of Level 3.

[13] As described in Footnote 10, 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.

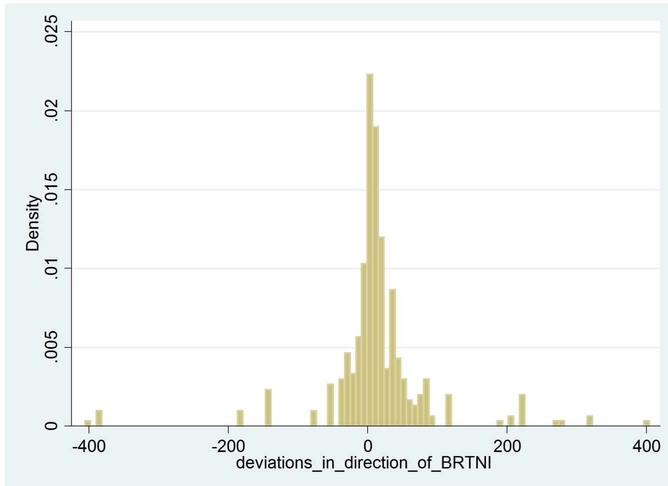

Figure 3: Single-file entry minus target. Positive values indicate deviations of the same sign as BRTNI's deviation. For scaling purposes, the 127 entries on target (zero deviation), and six outliers in the right tail, are not depicted.

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 $\gamma \geq 0$ is 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, namely 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 those who play the Nebi strategy, $\gamma = 0.90$ 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 1 and 2. They depict the mean and median entries relative to their targets across periods. The median in period t of Figure 5 corresponds to the median across the 21 games of the average of the four period-t entries, while the mean in period t of Figure 4 represents the mean across the 21 games of the average of the four period-t entries.

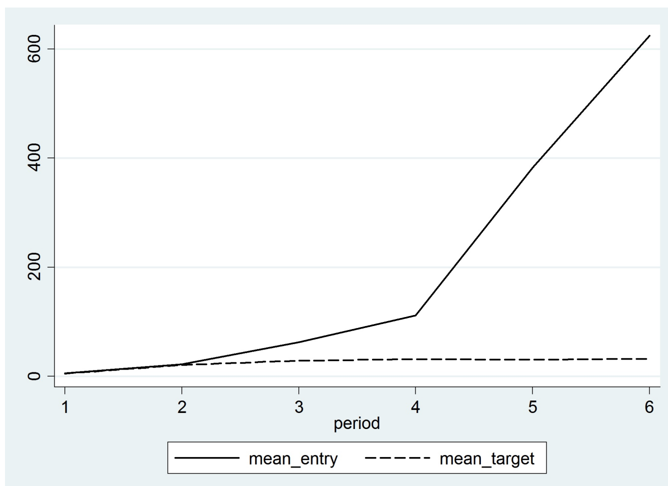

Figure 4: Multi-file mean absolute entries and targets

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 magnitude is 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 begin to accumulate 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 participants make choices that are too extreme on average, in several games

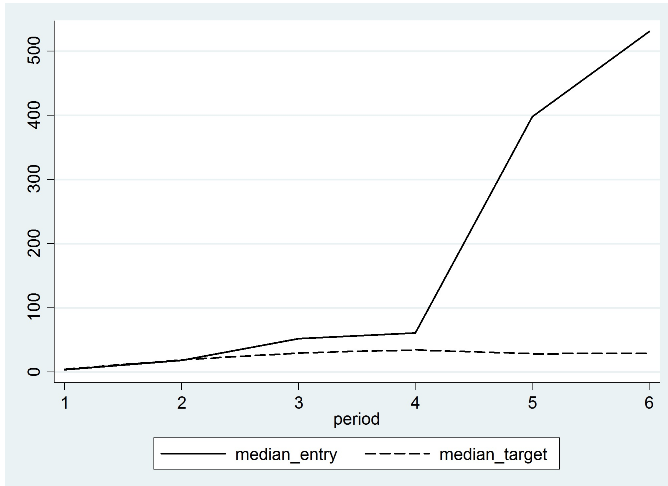

Figure 5: Multi-file median absolute entries and targets

a subset of players do appear to recognise 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.[14] The table 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.

	$\sum_{i=1}^{4}\sum_{t'=1}^{t-1} s_{i,t'}$	$s_{1,t}$	$e_{1,t}$	$s_{2,t}$	$e_{2,t}$	$s_{3,t}$	$e_{3,t}$	$s_{4,t}$	$e_{4,t}$
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

Table 2: Signals and entries in Game 17

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

Entries in the first two periods equal or approximate the 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 $t = 3$ 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 only one player chooses an extreme entry of -150, consistent with further 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 neglectors propagates extreme beliefs.

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.

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	Other
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: Multi-file entries consistent with different behavioral models,

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. Indeed, 63% of entries in the first three rounds hit their targets, and in the third game of each session, Nebi play, which includes anti-imitation, occurs more frequently than the simpler BRTNI play. In periods $4, 5, 6$, 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.[15]

Just as in the single-file treatment, BRTNI behavior diminishes with experience (see the tables in Appendix A). But regardless of whether participants follow BRTNI or another form of redundancy neglect, they make far too many extreme entries, which Appendices A and B

[15] Of the 220 unexplained observations, 11 can be explained by enriching one of the proposed models by allowing participants to flip signs of their private signals or flip the signs of what they infer from their predecessors, or both.

document for nearly all games.

Overall, only 6% of entries from the final three periods hit their target. Figure 6 depicts these deviations from target. 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.

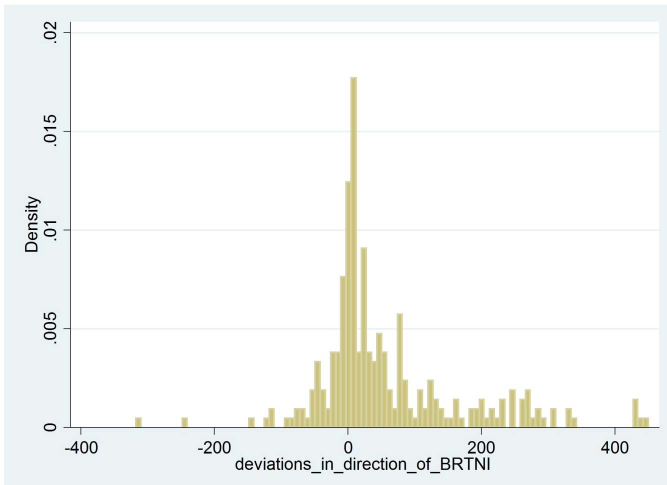

Figure 6: Multi-file entries off target. Positive values indicate deviations of the same sign as BRTNI's deviation. For scaling purposes, the 172 entries on target (zero deviation), 56 outliers in the right tail, and 4 outliers in the left tail are not depicted.

Altogether, the multi-file treatment induces strong herding that leads participants severely astray. 44% of the data are consistent with BRTNI, and much more are consistent with a general propensity to neglect redundancy, which leads to ever increasingly extreme and off-target predictions. Participants in later periods make entries which reveal that they 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 earned 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 such an effect.

Finally, Figure 7 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 first four signals sum to something positive versus the average entries and targets in multi-file games whose first four signals sum to something negative. Because the signals are iid, the sign of the first four signals does not predict later signals, and, hence, the targets (dotted line) remain stable on average after $t = 1$. The 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 20 in Appendix A gives an analogous picture for single file that shows no significant momentum.

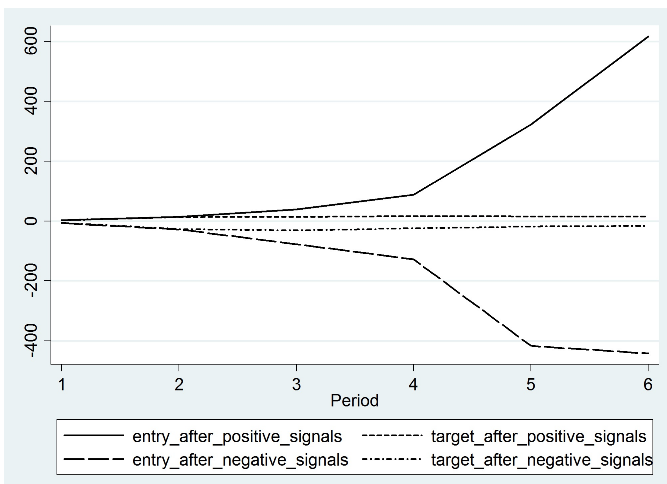

Figure 7: Average entry and target in multi file, separate by period and by the sign of the sum of $t = 1$ signals in the same game.

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 and of redundancy neglect and anti-imitation. We focus on the multi-file treatment and only briefly report results on the single-file treatment, where three-quarters of participants adhere to Nebi play.

Although our regressions organize the data in a simple manner, because they impose a rather rigid structure on how players react to one another, they fail to address the substantial heterogeneity in behavior identified in Section 4. For example, our linear-in-entries regression framework does not accommodate the possibility that players in period t may follow a moderate action by one previous player in period t' while at the same time ignoring an extreme action by another previous player in period t''. Moreover, it does not account for the co-existence of different sets of players prone to redundancy neglect to different degrees. The linear regressions merely summarize the central tendencies in behavior: even if participants in the multi-file treatment rarely match Nebi's prediction precisely, they may still show a propensity to anti-imitate where appropriate. Likewise, even if BRTNI precisely fits only a minority of data, the regressions may uncover that players have a general tendency to over-imitate.

Our first set of results investigates over-imitation in the multi-file treatment by describing the connection between early signals and later entries. This parallels Figure 7'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', $\overline{s}_{t'} = \sum_{i=1}^{4} s_{i,t'}$. Nebi predicts that all coefficients should equal one, since all signals are correctly accounted for in equilibrium.[16] BRTNI makes the same predictions for $t = 1$ and $t = 2$. For periods

[16]Nebi/BNE also predicts that a constant regressor has a zero coefficient, and we therefore omit the

$t \geq 3$, BRNTI makes two distinct qualitative predictions.[17] First, it predicts that for each i and $t' < t-1$, the estimated effect of $\overline{s}_{t'}$ on $e_{i,t}$ should exceed one. Second, it predicts that for each i and $t' < t$, the effect of $\overline{s}_{t'}$ on $e_{i,t+1}$ should exceed the effect of $\overline{s}_{t'}$ 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.

	e_1	e_2	e_3	e_4	e_5	e_6
$\overline{s}_1$	-	1.025 (0.056)	2.983 (0.426)	3.073 (0.832)	13.513 (3.113)	19.609 (5.447)
$\overline{s}_2$	-	-	0.999 (0.231)	3.951 (1.307)	9.112 (2.459)	10.101 (10.426)
$\overline{s}_3$	-	-	-	1.066 (0.914)	6.711 (3.912)	10.362 (7.240)
$\overline{s}_4$	-	-	-	-	- 4.786 (3.750)	-1.553 (9.511)
$\overline{s}_5$	-	-	-	-	-	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

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

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 estimated coefficients far larger than one in periods $t' + 2, t' + 3, \ldots$ (the smallest point

constant. Empirically, the inclusion of a constant regressor leaves the results essentially unchanged.

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

estimate being 3.0 and the largest 19.6). In addition, along each of the first rows in the table, the estimated coefficients increase monotonically, as predicted by BRTNI and many of its relatives in the family of 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, participants' own signals do not have significant effects on their actions.

Table 5 presents analogous regressions for the single-file treatment. To readily compare coefficients with the multi-file treatment, we group periods $t = 1, \ldots, 4$ into "super-period" $\tilde{t} = 1$, periods $t = 5, \ldots, 8$ 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 participants' super-period-$\bar{t}$ entries, $e_{i,\bar{t}}$, on the period-t signals they receive, $s_{i,t}$, as well as on the sum of all signals in every prior super-period $\bar{t}'$, $\bar{s}_{\bar{t}'} = \sum_{i=1}^{4} s_{i,\bar{t}'}$. Table 5 presents the regression results, where, as in Table 4, player indexes are omitted from the dependent variables $e_{\bar{t}}$ for conciseness.

	$e_{\bar{1}}$	$e_{\bar{2}}$	$e_{\bar{3}}$	$e_{\bar{4}}$	$e_{\bar{5}}$	$e_{\bar{6}}$
$\bar{s}_{\tilde{1}}$	-	1.302 (0.697)	5.751 (4.849)	1.672 (1.027)	1.664 (.701)	0.680 (1.897)
$\bar{s}_{\tilde{2}}$	-	-	1.036 (0.875)	1.211 (1.166)	1.376 (.956)	-1.942 (.911)
$\bar{s}_{\tilde{3}}$	-	-	-	1.502 (.558)	1.584 (.289)	0.017 (1.112)
$\bar{s}_{\tilde{4}}$	-	-	-	-	0.496 (.477)	2.553 (.589)
$\bar{s}_{\tilde{5}}$	-	-	-	-	-	-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

Table 5: Single-file regressions of super-period-$\bar{t}$ entries on current and past signals. Standard errors in parentheses are clustered by session.

BNE/Nebi play predicts that all coefficients should equal one, since all signals are correctly accounted for in later periods. Despite most point estimates across the first five super-periods exceeding one, the differences from one are statistically insignificant. (In the sixth super-period, two of the estimated coefficients of the effect of past signal on actions take the wrong sign, and all have wide confidence intervals.) Moreover, the regressions provide no indication that early signals exert increasingly strong influence on later and later actions.

We now turn to a direct test 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. Our next regression looks for evidence of anti-imitation in the multi-file treatment; since Nebi, BRTNI, and most other models of interest do not predict anti-imitation in the single-file treatment, we exclude it from this part of the analysis.[18] We regress participants' period-t entries $e_{i,t}$ on their period-t signals $s_{i,t}$ and on lagged entries, $\overline{e}_{t'}$ with $t' < t$. BNE and Nebi predict a coefficient on $s_{i,t}$ equal to one and coefficients on $\overline{e}_{t'}$ 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 should get the same coefficient of 1. Other rules embedding different forms of redundancy neglect predict that all of the coefficients should be non-negative, and that their sum should exceed one. Table 6 shows the regression results, where once again player indexes are omitted from the dependent variables e_t for conciseness.

[18]Note, however, that Level-3 does predict anti-imitation in the single-file treatment.

	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

Table 6: Multi-file regressions of period-t entries on current signals and past entries. Standard errors in parentheses are clustered by session.

The coefficients clearly 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. 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. Overall, the regressions of Table 6 confirm that the central tendency of the data patterns does not include anti-imitation.

6 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.

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 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 underling 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 22 than Level 1 or Level 2, although of course we think that Level 22 is not the right conceptualization. The data patterns also suggest that models which incorporate decision noise such as Quantal Response Equilibrium (McKelvey and Palfrey (1996)) or noise-enhanced versions of cognitive-hierarchy models like Camerer, Ho and Chong (2004) would fare no better: in the games' rich action spaces they would fail to match the data as frequently as does BNE, at least in the single-file treatment.[19] Not only would these models miss systematic deviations made out or 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—may turn out backwards here because our many Nebi participants very often fail to take into account their predecessors' over-counting (or other mistakes). Once again, we did not aim to test any of these models, and refrain from doing so ex post. Yet 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 games in which the enhanced solution concepts offer worse predictions than traditional solution concepts.

[19] In 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 should be compared to Figure 1.

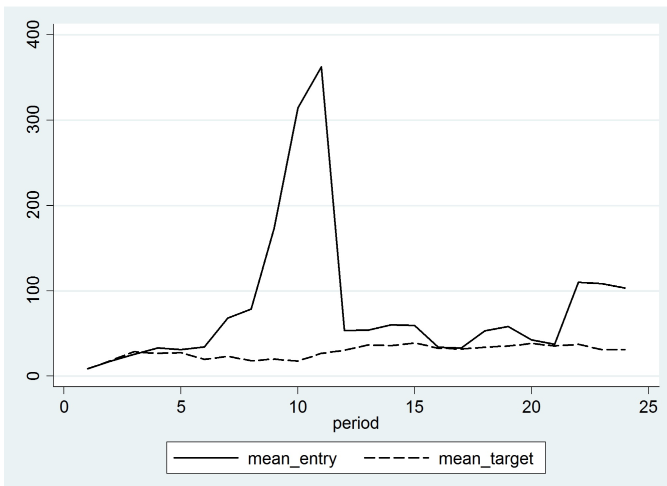

Figure 8: Single-file mean absolute entries and targets for first games

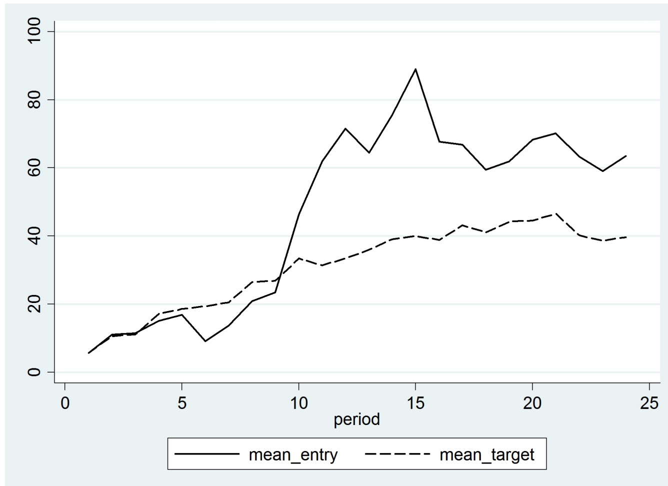

Figure 9: Single-file mean absolute entries and targets for second games

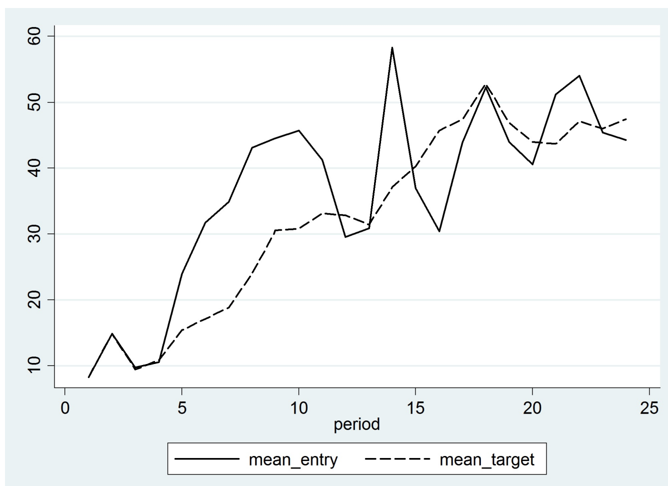

Figure 10: Single-file mean absolute entries and targets for third games

.

The next three figures should be compared to Figure 2.

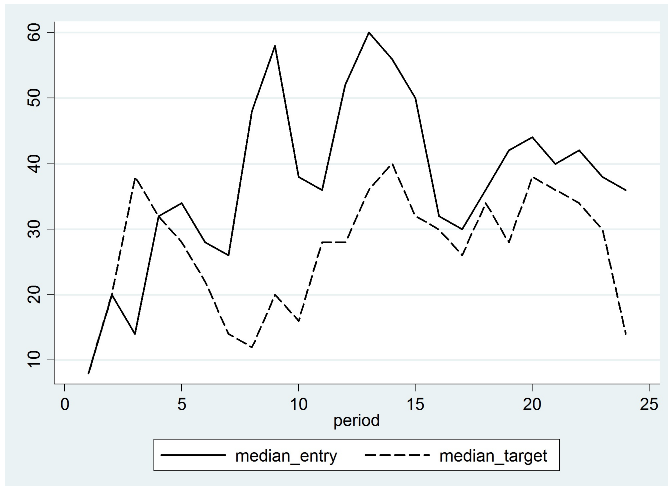

Figure 11: Single-file median absolute entries and targets for first games

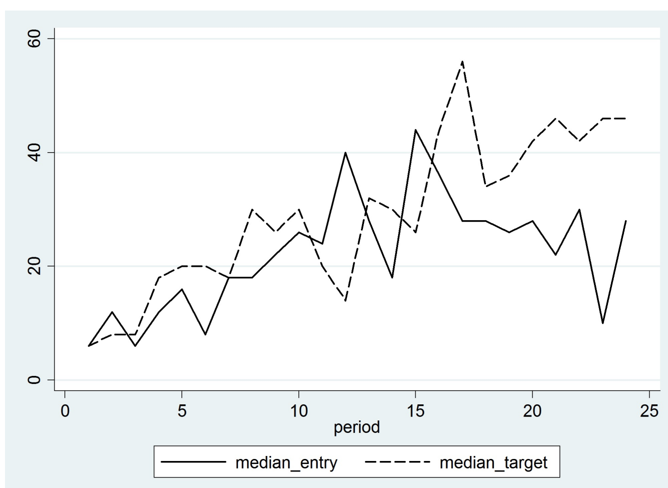

Figure 12: Single-file median absolute entries and targets for second games

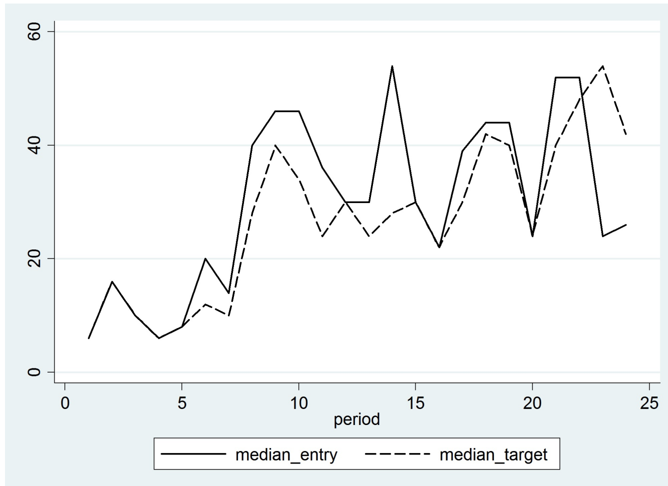

Figure 13: Single-file median absolute entries and targets for third games

.

The next three figures should be compared to Figure 4.

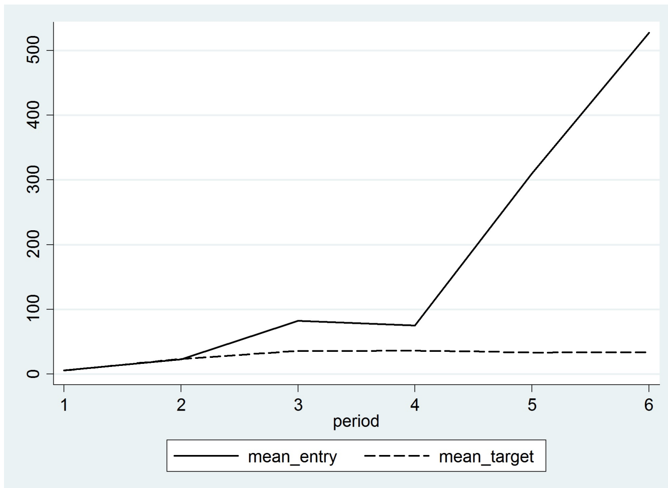

Figure 14: Multi-file mean absolute entries and targets for first games

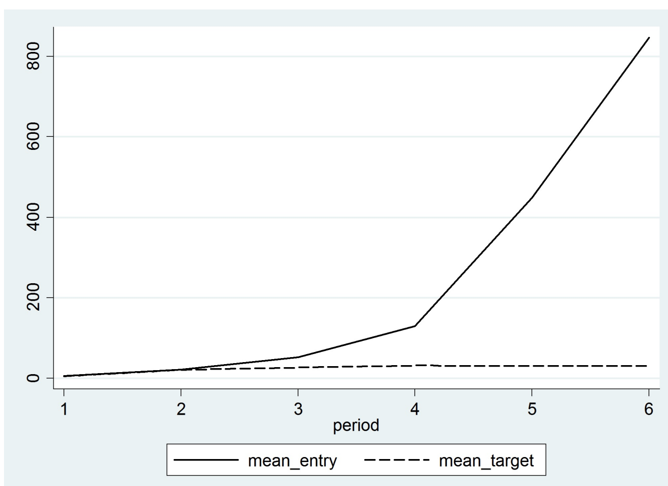

Figure 15: Multi-file mean absolute entries and targets for second games

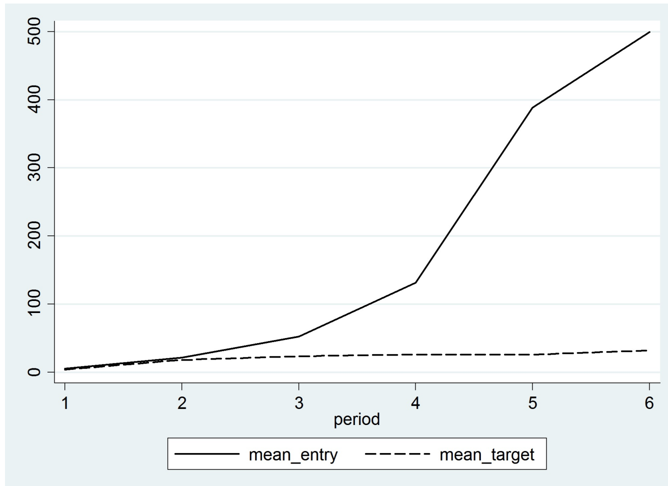

Figure 16: Multi-file mean absolute entries and targets for third games

The next three figures should be compared to Figure 5.

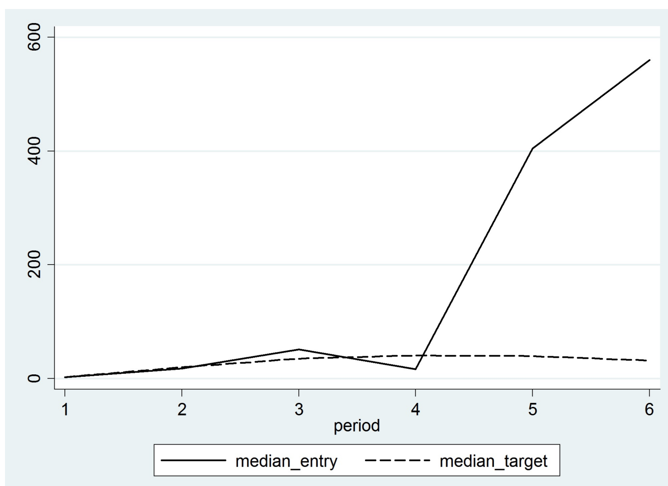

Figure 17: Multi-file median absolute entries and targets for first games

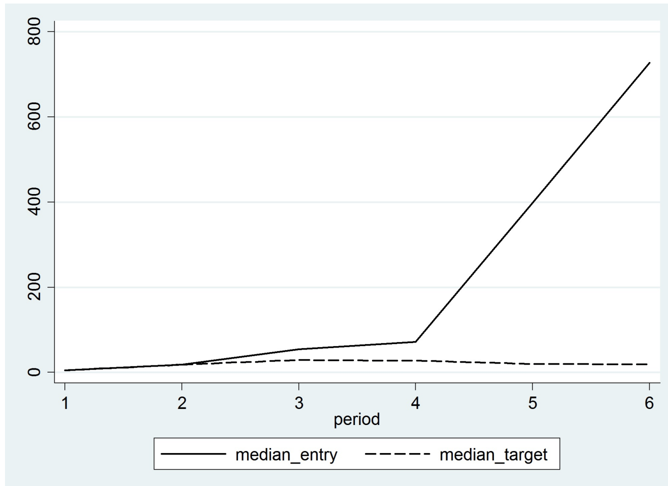

Figure 18: Multi-file median absolute entries and targets for second games

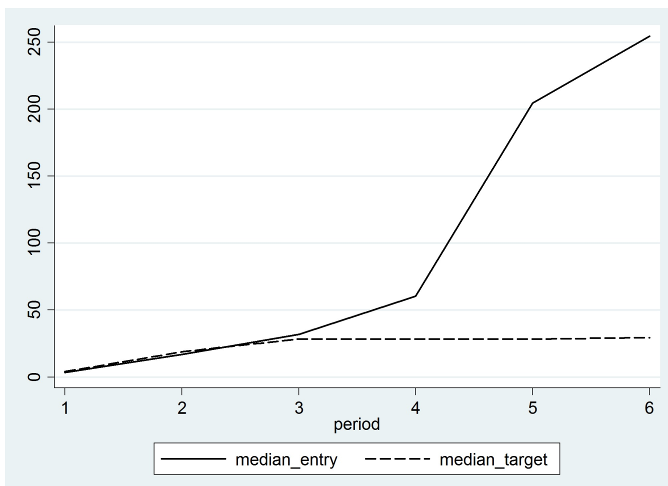

Figure 19: Multi-file median absolute entries and targets for third games

.

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

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	Other
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 7: Single-file entries consistent with different behavioral models for first games

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	Other
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 8: Single-file entries consistent with different behavioral models for second games

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	Other
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 9: Single-file entries consistent with different behavioral models for third games

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

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	Other
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 10: Multi-file entries consistent with different behavioral models for first games

Period	Cursed ($k = 1$)	BRTNI ($k = 2$)	Level 3 ($k = 3$)	Nebi ($k > t - 1$)	Other
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 11: Multi-file entries consistent with different behavioral models for second games

Period	Cursed ($k=1$)	BRTNI ($k=2$)	Level 3 ($k=3$)	Nebi ($k>t-1$)	Other
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 12: Multi-file entries consistent with different behavioral models for third games

The following Figure 20 is analogous to Figure 7 but uses the single-file data. It shows the between-game average entries and targets in single file and by period, separated according to a positive versus negative sum of the first four signals in the same game. Apart from the episode in Game 34, entries do no discernibly deviate from target. This also implies that having a positive versus negative sum of early signals does not induce a systematic deviation from target.

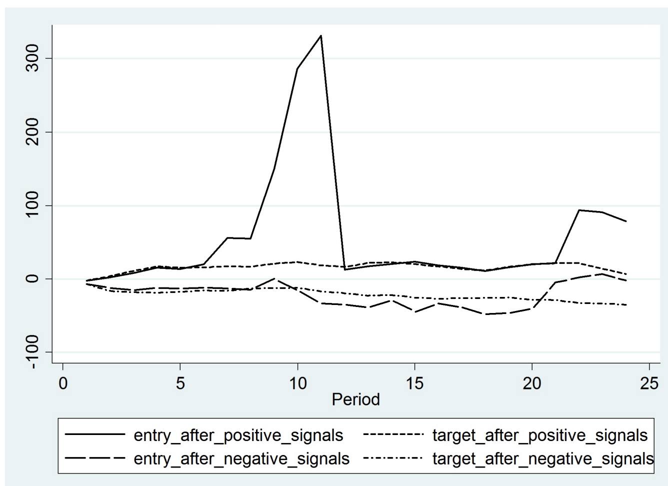

Figure 20: Average entry and target in single file, separate by period and by the sign of the sum of $t = 1, ..., 4$ signals in the same game.

References

Anderson, Lisa R. and Charles A. Holt, "Information Cascades in the Laboratory," *American Economic Review*, 1998, *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.

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.

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 and Florian Zimmermann, "Correlation Neglect in Belief Formation," *Mimeo*, 2015.

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.

Goeree, Jakob, Thomas Palfrey, Brian Rogers, and Richard McKelvey, "Self-Correcting Information Cascades," *Review of Economic Studies*, 2007, *74* (3), 733–762.

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.

March, Christoph and Anthony Ziegelmeyer, "Even Experts Follow Large Herds: Naïve Herding in the Laboratory," *Mimeo*, 2015.

McKelvey, Richard and Thomas Palfrey, "A Statistical Theory of Equilibrium in Games," *Japanese Economic Review*, 1996, *47* (2), 186–209.

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.