Entry and exit decisions under public and private information: An experiment

Aleksei Chernulich* John Horowitz† Jean Paul Rabanal ‡
Olga A. Rud § Manizha Sharifova ¶

September 29, 2021

Abstract

We design an experiment to study how reversible entry decisions are affected by public and private payoff disclosure policies. In our environment, subjects choose between a risky payoff, which evolves according to an autoregressive process, and a constant outside option payoff. The treatments vary the information disclosure rule on the risky payoff, such that in the public information treatment the risky payoff is always observable, while in the private information treatment, the risky payoff is observable only to the participants who enter the market. We find that under private information, market entry is higher, which suggests that subjects engage in exploration and place value on information.

Keywords: experiment, entry and exit decisions, bandit problems, information provision, forecasting

JEL codes: D81, D83, D84, G11, G170, C91

*Center for Behavioral Institutional Design, New York University Abu Dhabi
†Department of Economics, Ball State University
‡Department of Economics and Finance, University of Stavanger
§Department of Economics and Finance, University of Stavanger
¶Department of Economics, University of the Pacific

1 Introduction

A large number of retail transactions occur on centralized online platforms which allow sellers and buyers to seamlessly exchange goods and services. These transactions are easily recorded, often in real time, and offer sellers rich data on consumer preferences. In turn, this data can be used by sellers to forecast future demand with greater precision. Typically, the production decision involves the question of how much to produce, or the intensive margin, and whether to produce, or the extensive margin. The latter is also known as a market entry decision. In this paper, we report on a laboratory experiment which studies how information affects market entry decisions (IN) when the outside option is a constant payoff (OUT).

In our environment, the decision to enter the market is fully reversible– that is, market participants can switch between IN and OUT decisions at any time. The payoff to IN follows a stationary AR(1) process, which reflects market conditions. The participant does not have any market power, and therefore takes prices as given. The payoff to OUT is less than the expected payoff to IN. For simplicity, we assume zero transaction costs for switching between IN and OUT and provide subjects with information about the risky payoff generation process. Our main research question is how the information provided to a subject affects their entry decision. While market transparency in relation to buyers is an actively developing field,[1] our experiment allows us to study the impact of transparency on the supply side of the market. In private information treatment, which is a typical bandit problem, participants select between a safe arm (OUT), and a risky arm (IN), where the platform provides information in almost real-time. If a participant selects OUT, then they will not observe the returns to the risky payoff. In order to observe the return to the risky arm, this option must be selected. In the public information treatment, we provide information on the risky payoff regardless of whether the participant is IN or OUT.

To draw predictions for IN and OUT decisions across both treatments, we assume that agents forecast future revenues under two different rules: (i) rational expectations, and (ii) behavioral expectations that allow for well-known forecasting biases such as extrapolation and stickiness (Landier et al., 2019). While both rules suggest that IN will be observed more often under public information, we find that subjects select IN more often when information is private. This result suggests that there is a demand for information, and that participants do not exhibit a high degree of risk aversion

[1] See Kaya and Liu (2015), Fuchs et al. (2016), and Bergemann and Hörner (2018).

since they are willing to explore the risky option. The last result is supported by other bandit experiments in the lab (Hoelzemann and Klein, 2021).

Our design is motivated by market entry decisions when information flows rapidly. We modify the standard bandit problem found in economics and finance (Bergemann and Välimäki, 2008) to study reversible entry decisions under different disclosure rules. In a related experiment, Grosskopf et al. (2006) find that providing counterfactual information in a bandit problem can increase risky behavior. However, this result disappears with experience, suggesting that subjects become less sensitive to additional information over time. Yechiam and Busemeyer (2006) show that counterfactual information can increase risky option selection when the negative outcome is rare and large. Biele et al. (2009) employ a Markov process with two states for the risky option, H and L, which are unknown to subjects. They find that players do not learn to become risk averse. In our experiment, we enrich the set of possible outcomes, and providing more opportunities for the subjects to familiarize themselves with the payoff realization process due to the nearly continuous environment.[2]

To formulate our predictions we draw from past literature on expectation formation where the predicted variable is exogenous.[3] According to Assenza et al. (2014) and Mokhtarzadeh and Petersen (2020), access to historical data in forecasting experiments can encourages more adaptive, and trend-chasing expectations. Landier et al. (2019) asked subjects to forecast 40 realizations of a risky asset and found evidence of both sticky (Coibion and Gorodnichenko, 2015; Bouchaud et al., 2019), and extrapolative (Bordalo et al., 2018) expectations.[4] We simplify the forecasting tasks by focusing on the binary decision of IN or OUT, which indirectly measures market expectations. Our findings complement existing experimental literature on switching behavior (Anufriev et al., 2016; Anufriev et al., 2018 and Anufriev et al., 2019). In these experiments, participants were offered an opportunity to switch between investment alternatives. In contrast to our paper, participants in these studies (i) were not informed about the payoff generating process, and (ii) were evaluated only in a public information environment.

[2]The payoff realization process occurs twice per second, and helps capture the instantaneous feedback of modern retail markets.

[3]See Schmalensee (1976), Dwyer et al. (1993), Hey (1994), Kelley and Friedman (2002), Glaser et al. (2007), and Beshears et al. (2013).

[4]For empirical evidence using expectations survey data, please refer to Amromin and Sharpe (2013), Gennaioli et al. (2016) and Wen (2018). For an overview of the interactions of individual forecasting rules, and the aggregate macro behavior they co-create, refer to Assenza et al. (2014).

2 Environment

For each $t = \{1 \ldots, T\}$, a player seeks to maximize profit π_t, by choosing between two actions; IN ($a = 1$) and OUT ($a = 0$). Choosing $a = 1$ yields a stochastic payoff $\pi_t = x_t + 100$, in which x_t is driven by an autoregressive process of order one–AR(1)– as specified in equation (2), and ϵ follows a standard normal distribution. OUT, the outside option, yields a constant payoff $r + 100$.

$$\pi_t = \begin{cases} x_t + 100 & \text{if} \quad a = 1, \\ r + 100 & \text{otherwise}, \end{cases} \tag{1}$$

where

$$x_{t+1} = \rho x_t + \sigma \cdot \epsilon_t. \tag{2}$$

The payoff structure we employ has two important features worth noting. First, the payoff for IN x_t incorporates a risk premium, and therefore $r < E(x_t)$. Second, to avoid negative payoffs we add a constant component equal to 100 to both alternatives. We study two environments with different information disclosure rules: public information and private information. In the **private information** treatment, the subject does not observe the payoff to IN if they choose OUT. In the **public information** treatment, we provide counterfactual information on the foregone payoff when the subject selects OUT. In both environments, the subject's decision between IN and OUT at every t is driven by the subject's beliefs regarding the one period ahead value x_{t+1}, denoted as $F_t x_{t+1}$. Consequently,

$$a = \begin{cases} 1 & \text{if} \quad F_t x_{t+1} > r, \\ 0 & \text{otherwise}. \end{cases} \tag{3}$$

We allow beliefs $F_t x_{t+1}$ to accommodate a large family of expectations, from Rational Expectations (RE) with $F_t x_{t+1} = E_t x_{t+1} = \rho x_t$, to Sticky and Extrapolative Expectations (SEE) which capture the belief formation for AR(1) processes (Landier et al., 2019). Under SEE, the belief $F_t x_{t+1}$ combines extrapolative expectations, in which subjects overreact to unexpected innovations (Bordalo et al., 2018), and sticky expectations (Coibion and Gorodnichenko, 2015, Bouchaud et al., 2019), in which subjects demonstrate inertia in updating expectations. We use a nonrecursive specification for the one period ahead forecast, which depends on current and past RE, denoted as $\sum_{k=0}^{n} E_{t-k} x_{t+1}$, as well as the history of x,

$$F_t x_{t+1} = (1-\lambda) \sum_{k=0}^{n} \lambda^k E_{t-k} x_{t+1} + \gamma \sum_{k=0}^{n} \lambda^k (x_{t-k} - E_{t-k-1} x_{t-k}) \cdot \mathbb{1}_{t-k}, \tag{4}$$

where $\lambda \in [0,1]$ is the degree of stickiness, $\gamma > 0$ captures the importance of extrapolative beliefs, and $\mathbb{1}_{t-k}$ takes the value of one when either (i) the environment is public information, and therefore x_{t-k} is always observed, or (ii) the environment is private, and the subject selected IN for $t-k$. If there is no stickiness to belief updating ($\lambda = 0$), and no trend extrapolation ($\gamma = 0$), then the SEE belief is exactly RE, $F_t x_{t+1} = E_t x_{t+1}$.

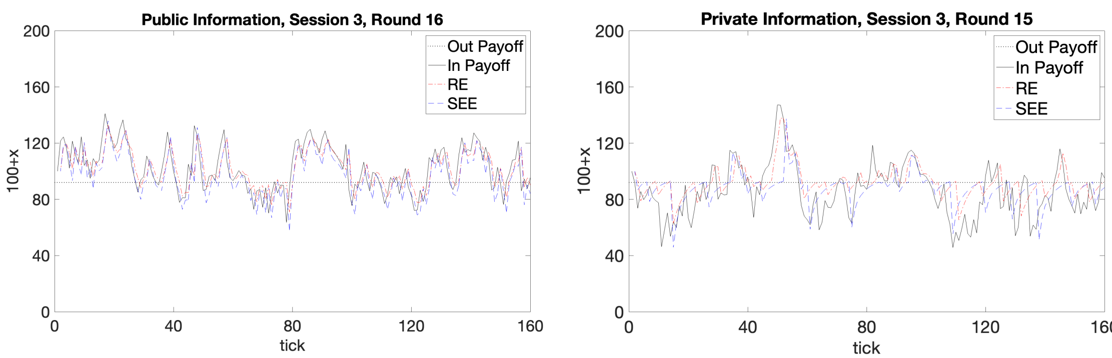

Figure 1: Payoff from choosing IN and the evolution of belief under rational expectations (RE) and sticky and extrapolative expectations (SEE). The left panel depicts an example of public information treatment. The right panel depicts an example of private information treatment. The horizontal dashed line represents the value of OUT option and the solid line represents the value of IN option.

Figure 1 shows examples of RE and SEE belief formations using a series of actual x realizations observed by participants in public and private information treatments with $\{\rho, \sigma, T, r, \lambda, \gamma\} = \{.85, 12, 160, -8, .21, .41\}$. We employ such parameter values for the following reasons. The value of ρ has to be high enough to create an important half-life for innovations ϵ. However, it should not be too close to one in order to avoid a random walk. The value of σ follows the empirical work of Landier et al. (2019), though we assume a smaller value to mitigate fatigue from facing very volatile series. The value for OUT, r, is obtained from simulations, which create meaningful OUT spells. A smaller value of r incentivizes subjects to stay IN, and therefore limits the ticks in which we can observe players selecting OUT. The last two parameters (λ and γ) are estimates of behavioral expectations parameters reported by Landier et al. (2019).[5] The solid black line in Figure 1 shows the value of IN, the RED line represents an agent's belief under RE, and the blue line the belief under SEE. We assume that the simulated player makes a decision according to equations (3) and (4), and does not have information on the risky payoff when selecting OUT. The beliefs of the player evolve upwards after choosing OUT, because the value of IN follows a mean reversion process centered around zero plus the constant of 100.

Hypothesis: *Under both behavioral rules, RE and SEE, the frequency of* OUT *will be lower in public information treatment than in the private information treatment.*

We perform numerical simulations using the observed values of x_t (a total of 160-ticks per round) from our experimental sessions to formulate our Hypothesis. We simulate RE- and SEE-type behavior for all treatments and report average values for statistics of interest in Table 1, which include frequency of staying OUT, payoff as a fraction of ex-post optimal payoffs, and switching values of x_t.[6] The predictions show that OUT frequency is smaller in the public information treatment. This is due to the mean reversion process of x_t and the lower value expected payoff to OUT. This prediction is consistent across forecasting rules.

We analyze the drivers of behavioral differences across treatments according to the value of x with respect to the payoff to OUT. Conditional on $x < r$, the simulated player behaves similarly under both rules on average, choosing the outside option approximately 70 percent of the time under RE, and opting OUT a bit more frequently under SEE. The main difference arises when $x \geq r$, where the frequency of IN is smaller in the private information treatment (.77 under SEE and .82 under RE) compared to the public information treatment (.85 in SEE and .88 under RE).

In the second section of Table 1, we calculate how the simulated choices perform in terms of the ex-post optimal (maximum) payoff. The predicted behavior is close to optimal under both treatments, and under both forecasting rules. The next section of Table 1 shows the values of x when simulated players switch to OUT and IN. We find that players switch OUT, for both first and all switching decisions, when the value of OUT is below 92, which is in line with our expectations. Regarding the value of x that

[5]Landier et al. (2019) results are robust to a family of values of $\rho \in \{0, .2, .4, .8, 1\}$, and a value of $\sigma = 20$. Given our tick size, the equivalent standard deviation is $14.142 = 20/\sqrt{2}$ for our design.

[6]Information on OUT spells for simulated behavior is available in the Appendix D.

Table 1: Predictions under sticky and extrapolative expectations (SEE) and rational expectations (RE)

	SEE[c]		*RE*	
	public	*private*	*public*	*private*
Frequency OUT[a]	.36	.39	.31	.33
Freq. OUT $\|x < r$[b]	.75	.67	.69	.61
Freq. IN $\|x \geq r$	.85	.77	.88	.82
Payoff as a fraction of ex-post optimal	.97	.97	.97	.97
Payoff $\|x < r$	.97	.96	.96	.96
Payoff $\|x \geq r$	.98	.98	.98	.98
$x + 100$ *when switching OUT (all)*				
Mean	89.45	86.52	86.34	84.11
SD	13.04	14.11	12.44	13.84
$x + 100$ *when switching OUT (first)*				
Mean	91.74	89.19	86.31	87.00
SD	15.62	11.95	13.86	12.88
$x + 100$ *when switching IN*				
Mean	103.07	97.87	101.08	96.13
SD	13.30	20.39	13.30	17.64

Notes: a. Frequency of OUT is computed using a tick count of when OUT is selected.
b. Frequency of OUT conditional on the realized value of x being below the payoff for OUT.
c. SEE (RE) represents the simulation of a player with sticky and extrapolative expectations (rational expectations) using the data generated in the experiment of x for each treatment; 80 series of 160 ticks.

triggers the switch to IN, we observe it to be slightly higher with public information than with private information. However, the difference is not statistically significant. The predicted comparative statics in Hypothesis are robust to increasing risk aversion.[7]

Alternatively, we may also find that subjects opt IN more often in the private information treatment due to informational demand. Thus, to obtain information on the payoff to x, subjects may (i) delay their decision to switch from IN to OUT, and/or (ii) shorten the OUT spells by re-entering prematurely to gain information about the risky payoff. This alternative hypothesis implies that the value of x that triggers the decision to switch from IN to OUT may be higher in the public information treatment, where players do not gain anything from delaying the decision to switch from IN to OUT. To accurately identify the value of x that triggers switching and separate the effect of exploration, we present summary statistics for the first switching decision,

[7]Risk neutrality does a good job in explaining behavior in dynamic environments. For example, see Magnani and Munro (2020) who also state their predictions under risk neutrality in a dynamic experiment. For robustness, we allow for risk aversion and assume a power utility, u^{α} where $\alpha \in \{.3, .5\}$.

when the information sets are comparable across treatments.

3 Laboratory procedures

The experiments were conducted at the MonLee laboratory in Monash University using oTree software (Chen et al., 2016). Subjects were recruited online via SONA software and included undergraduate students across all fields. We assigned all participants to one of the two possible treatments: (i) public information, where the information on the foregone payoff is always available, or (ii) private information, where the risky payoff (IN) is unknown when the player opts OUT, and earns a constant payoff.[8] We elicited risk attitudes in all sessions, following the protocol of Crosetto and Filippin (2013) given that risk preferences may explain the heterogeneity of choices within a treatment.[9] In the instructions, we present the subjects with the underlying AR(1) process and the parameters used. After reading the instructions, subjects answered four control multiple-choice questions.[10] If a subject answered a question incorrectly, then the experimenter privately discussed with the participant the relevant section of the instructions.

Each session included two practice rounds, followed by 20 paid rounds. The risky asset payoff was updated every half a second, for 80 seconds (160 ticks) per round. The value of the risky asset was the same within a session, but all realizations of AR(1) were different between sessions. In total, 83 subjects participated in the experiment. Forty-one subjects participated in the public information treatment, and 42 participated in the private information treatment. Table 2 presents an overview of all laboratory sessions, and an Online Appendix provides dynamic graphs for all experimental sessions.

[8]The instructions for the public information treatment are in the Online Appendix. The key differences in wording across instructions are: *When you switch to OUT, you will stop observing the variable payoff of IN* for private information treatment, and *When you switch to OUT, you will still observe the variable payoff of IN* for public information treatment.

[9]To measure risk attitudes, Crosetto and Filippin (2013) ask subjects to decide how many of 100 boxes they want to collect. Earnings increase linearly with the number of boxes collected. However, if subjects select a box with a bomb inside, then their earnings are zero. In our experiment, risk-neutral expected earnings (50 boxes collected) are \$2. Appendix A presents the individual choices in the risk elicitation task.

[10]We asked subjects to answer the following: (i) what is the average payoff of IN?, (ii) if you select OUT, then you accumulate points according to (100, x, 92, 0), (iii) if you switch from OUT to IN, can you switch again and go OUT? (iv) Does the current value of x affect the value of x in the next period?

Table 2: Overview of sessions

	Public	*Private*
Profit (points per tick)		
Mean	101.16	100.11
SD	17.25	17.85
Profit without show-up fee (\$)		
Mean	10.12	10.01
Show-up fee (\$)	10	10
Number of subjects	41	42
Number of sessions	4	4

Notes: Subjects were paid in Australian dollars.

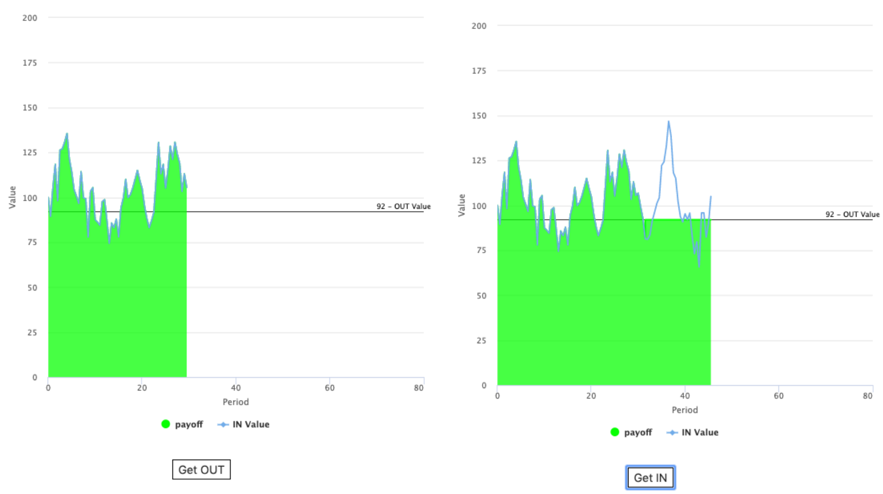

Figure 2: User interface in the public information treatment: (i) left panel shows the payoff (in green) for staying IN up to tick 30, and (ii) right panel shows the payoff (in green) if a subjected switched OUT at tick 30.

Figure 2 shows the user interface (UI) in the public information treatment for two different decisions. The left panel of Figure 2 shows the UI as seen by the subjects when they select IN, and the right panel shows the UI when they select OUT. The UI for the private information treatment is similar to the public information treatment, except that when the player selects OUT, they no longer observe the payoff to IN. The value of IN, $x + 100$, is depicted as a blue line while the value of OUT appears as a horizontal line at the ordinate value of $r + 100 = 92$. For each subject, the default initial state is IN. If the current strategy is IN, subjects can switch OUT by clicking the

"GET OUT" button located at the bottom of the interface, and if the current strategy is OUT, subjects can switch IN by clicking the "GET IN" button also located at the bottom of the interface. Players can switch between IN and OUT each tick, which lasts half a second, for $T-1$ ticks. The green shaded area represents the accumulated payoffs. When a subject selects OUT, the payoffs accumulate at the constant rate of 92 points per tick (see right panel of Figure 2).

After each round ends, we show subjects the points accrued in that round, and the cumulative points earned from all non-practice rounds. The experimental sessions lasted about 50 minutes. At the end of the session, the points earned across all rounds were added and converted to cash at the exchange rate of \$.003125 per 100 points. Excluding the show-up fee of \$10, subjects received on average \$10.12 in the public information treatment and \$10.01 in the private information treatment (see Table 2). The two treatments have similar payoffs, which is consistent with our predictions, due to the mean reversion process that governs the evolution of the value of IN. Despite the similarities in average values, we observe important differences in behavior when the value of x goes below the outside option payoff.

4 Results

We begin our discussion of results with Figure 3, which shows an example of a round from each treatment.[11]

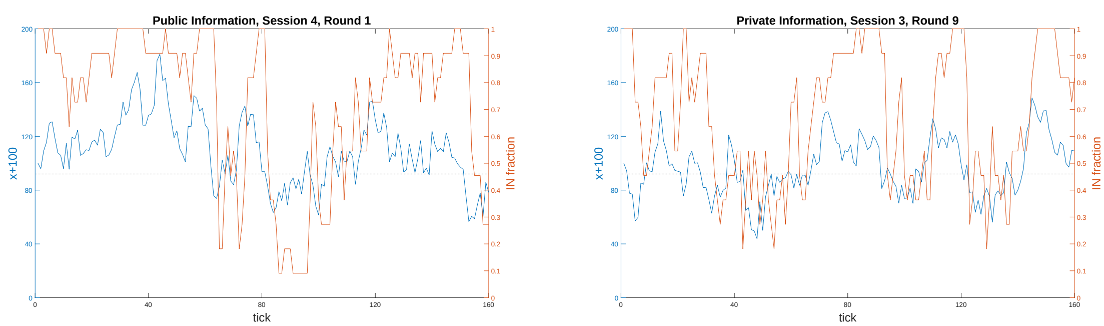

Figure 3: Evolution of x (blue) and fraction of players (red) in a session selecting IN for the public (left) and private (right) treatments. The black dotted line is the outside option payoff.

[11] Appendix E has the complete time-series observations for all experimental sessions.

The left panel of Figure 3 presents the results for the public information treatment, while the right panel presents the results for the private information treatment. The blue line depicts the value of $x + 100$, measured against the left y-axis, while the red line shows the fraction of subjects who choose IN, measured against the right y-axis, at time t. The black line is the outside option payoff. We observe that subjects actively move IN and OUT of the market throughout the round. When the value of IN is high, more players select the risky option (such that the fraction of subject IN approaches 1), and when the value of IN is low, more players choose the outside option.

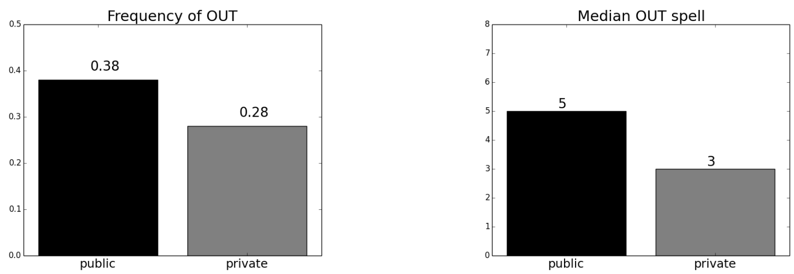

Figure 4: Summary of results (pooled data): (i) frequency of OUT (left panel), and (ii) median spell OUT (right panel). Under RE, the frequency of OUT is 0.31 and 0.33 for public and private treatments, respectively, while the median OUT spell is 2 and 4 for public and private treatments, respectively.

Figure 4 provides a summary of the observed frequency of OUT spells and the duration of the median OUT spell for both treatments. The black bar shows the results for the public information treatment, while the grey bar shows the results for the private information treatment. According to the left panel, subjects stay OUT more frequently in the public information treatment. Similarly, the median OUT spell duration is also longer in the public information treatment.

Next, we study the cumulative distribution function (CDF) in Figure 5 which shows the fraction of ticks when subjects select OUT. The CDF for the public information treatment first order stochastically dominates the distribution for the private information treatment, which is consistent with the larger mean frequency of OUT in Figure 4.

To study subject behavior over time, we present summary statistics in Table 3 using data from (i) all rounds, and (ii) rounds 11 through 20. We find that experience

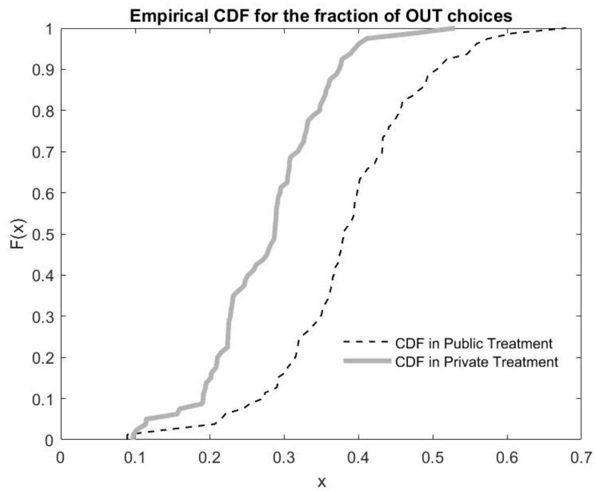

Figure 5: CDF of fraction of OUT choices in private v. public information treatments (pooled data).

does not affect subject behavior, with both data samples showing similar outcomes for each each treatment. To better understand the observed behavior, we compute the frequency of IN conditional on the value of x being equal to or greater than the outside option payoff (IN — $x \geq r$), and the frequency of OUT conditional on the value of x being below the outside option payoff (OUT — $x < r$). When IN is more profitable, subjects select IN 70 percent of the time in the public information treatment and 78 percent of the time in the private information treatment. When OUT is more profitable, subjects play OUT 58 percent of the time in the public information treatment and 39 percent of the time in the private information treatment. The second section of Table 3 calculates the relative payoff as a fraction of the ex-post optimal payoff. Consistent with the higher frequency of IN when IN is more profitable, we observe that payoff is higher in the private information relative to the public information. Overall, players perform slightly worse than the predicted payoff of 0.97 under both forecasting rules.

In the third section of Table 3, we compute the mean value of x that triggers players to switch from IN to OUT. For the public information treatment, the value is close to the outside option payoff of 92. For the private information treatment, players wait until $x + 100$ drops to 88 to switch from IN to OUT. If we look at the first switching decision, players switch OUT at about 88 in the public information treatment and 85 in the private information treatment.[12] The observed frequency is sightly above the

[12] Appendix D presents data on the duration of OUT spells for the 10th, 25th, 50th, 75th, and 90th percentiles, where the top percentiles indicate shortest duration. While there is no difference in

Table 3: Summary statistics

	All rounds		*Rounds 11-20*	
	public	*private*	*public*	*private*
Frequency OUT[a]	.39	.28	.42	.29
Obs.	128,000	131,200	64,000	64,000
Freq. OUT $\|x < r$[b]	.58	.39	.61	.40
Freq. IN $\|x \geq r$	.70	.78	.68	.78
Payoff as a fraction of ex-post optimal	.95	.96	.93	.96
Payoff $\|x < r$	.94	.93	.93	.93
Payoff $\|x \geq r$	.95	.97	.93	.97
$x + 100$ *when switching OUT (all)*				
Mean	91.71	88.19	91.58	87.20
SD	17.77	17.64	17.21	17.47
$x + 100$ *when switching OUT (first)*				
Mean	88.77	85.27	87.50	85.08
SD	17.81	15.41	18.08	15.64
$x + 100$ *when switching IN*				
Mean	103.96	95.75	103.52	94.51
SD	17.41	20.14	16.96	19.51

Notes: a. Frequency of OUT is computed using a tick count of when the OUT strategy is observed.
b. Frequency of OUT conditional on the realized value of x being below the risky-free payoff.

RE prediction of 86 in the public information treatment, and below the prediction of 87 in the private information treatment.

The last section of Table 3 presents the value of $x + 100$ that triggers the decision to select IN. In the public information treatment, subjects require approximately 104 points to select IN again. In the private information treatment, where subjects do not observe the value of IN while OUT, subject require a lower payoff of 96 to re-enter the market. The lower trigger value in the private information treatment is consistent with the shorter OUT duration in the private information treatment, where subjects switch sooner due to lack of information.

duration of an OUT spell at the 10th and 25th percentiles between the two treatments, the median duration of an OUT spell is two ticks greater in the public information treatment than the private information treatment.

Result 1. *Subjects select OUT more often in the public information treatment.*

According to the linear probability model presented in Table 4, subjects are more likely to select OUT in the public information treatment, contrary to the hypothesis. The dependent variable takes the value of one when the subject selects OUT, and zero when the subject selects IN for all specifications except (IV), where the definition is reversed. Further, in specifications (IV) and (V), the probability is conditional on the value of x. The dummy variable *Private* is the treatment effect, which equals one if the subject is in the private information treatment and zero otherwise. The dummy variable *Round* is the trend effect, which controls for learning. We also include the risk elicitation task as a control in specifications (II)-(V). We find that, on average, the frequency of OUT in the private information treatment is 10 percentage points lower than in public information treatment. This coefficient is robust to risk preferences (II-III) and learning (III).

Table 4: Linear probability model

	(I) OUT	(II) OUT	(III) OUT	(IV) IN $\|x \geq r$	(V) OUT $\|x < r$
Intercept	$.38^{***}$	$.41^{***}$	$.38^{***}$	$.72^{***}$	$.62^{***}$
	(.02)	(.02)	(.02)	(.01)	(.02)
Private	$-.10^{***}$	$-.10^{***}$	$-.10^{***}$	$.06^{*}$	$-.17^{***}$
	(.01)	(.02)	(.02)	(.02)	(.02)
Round	—	—	$-.00^{***}$	—	—
			(.02)		
Controls (Risk)	No	Yes	Yes	Yes	Yes
R^2	.01	.01	.01	.01	.03
N	259 200	259 200	259 200	170 359	88 841

Notes: The Intercept captures the public information treatment. Standard errors are in parenthesis, clustered at the subject level and are computed via bootstrapping. $^{***}p \leq .01$, $^{**}p \leq .05$, $^{*}p \leq .1$

Specification (IV), which restricts the sample to when the risky option outperforms the outside option, suggests that players in the private information treatment stay IN more often than players in the public information treatment. The largest difference in behavior is observed in specification (V), which restricts the sample to observations where the risky option under-performs the outside option. In this subsample, the frequency of OUT in the private information treatment is 17 percentage points lower. If information had no value, then we would not observe OUT more often in the private information treatment (see Table 1). Thus, specification (V) suggests that subjects in the private information treatment value information because they are willing to stay

IN. To confirm that these results are robust, Table 7 in Appendix B presents the linear regression results for rounds in which the value of IN is below 80 for at least 40 ticks. We conclude that the treatment differences are robust to when players are in markets with a low rate of return.

Result 2. *Duration of an OUT spell, or uninterrupted time spent OUT without switching, is longer in the public information treatment.*

To analyze the duration of an OUT spell, we use a Weibull survival function,

$$S(t; p, z_j\beta) = \exp(-\lambda_j t_j^p),$$

where t is the number of ticks that a player chooses OUT, p is the shape parameter and $\lambda_j = \exp(z_j\beta)$, which includes the regressors (z_j) and the coefficient (β). The hazard rate is computed as

$$h(t) = f(t)/S(t) = -\frac{d\ln S}{dt}.$$

The estimated parameters of the hazard function are presented in Table 5, and the survival function $S(t)$ is shown in Figure 6. The standard errors in the parametric estimation are clustered at the subject level. We find that $p < 1$, which indicates that $h(t)$ is a decreasing function. Note that in each round, we observe multiple OUT spells.

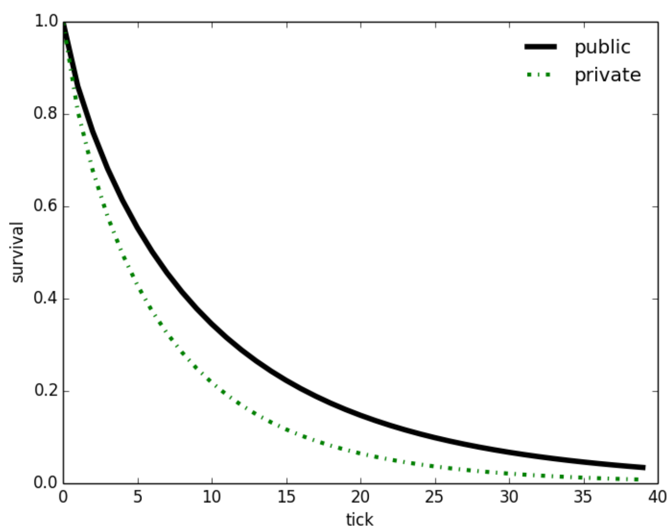

Figure 6: Weibull survival function: $OUT \rightarrow IN$

Table 5: Hazard function

	(I)	(II)
Intercept	-1.89^{***}	-1.74^{***}
	(.06)	(.08)
Private	$.36^{***}$	$.38^{***}$
	(.14)	(.14)
$\log(p)$	$-.16^{***}$	$-.16^{***}$
	(.03)	(.03)
Control (Risk)	No	Yes
ψ	6.96	8.55
N	86 169	86 169

Notes: To compute the hazard ratios, we use an exp function on relevant coefficients.
$^{***}p \leq .01$, $^{**}p \leq .05$, $^{*}p \leq .1$

The survival function confirms that subjects in the public information treatment stay out longer (solid black line) than in the private information treatment (dashed green line). We estimate the survival function using the parameters from Table 5 and find that the coefficient for *Private* is .36, and the hazard rate is 1.43 (= exp(.36)) in the private information treatment relative to public information treatment. The hazard rate does not significantly change when we control for risk preferences in specification (II). Shorter OUT spells in the private information treatment confirm that information has value– subjects choose IN to determine and evaluate the payoff of x relative to the outside option OUT. In the public information treatment, subjects know the value of x and can evaluate the relative payoff without switching prematurely. Hence, subjects are willing to opt OUT of the risky option more often in the public information treatment because they do not lose access to information.

Result 3. *Subjects select OUT faster in the public information treatment.*

To determine when subjects switch from IN to OUT, we use a Tobit regression. Since we are interested in a point estimate rather than the duration of an event, a Tobit regression can provide a more precise estimate than a Weibull survival analysis. The decision to switch OUT is dependent on observing a sufficiently low value of x. Therefore, we address possible censoring issues by using the Tobit regression.

Table 6 summarizes the results of the Tobit regressions for the value of x when subjects switch from $IN \rightarrow OUT$ in specifications (I) and (II). Specification (I) analyzes all IN/OUT decisions, while specification (II) focus only on the first IN/OUT switch. In the public information treatment, subjects switch when $x + 100$ is around 90. In

Table 6: Switching value of $x + 100$

	(I) $x+100$: $IN \rightarrow OUT$ All	(II) $x+100$: $IN \rightarrow OUT$ First	(III) $x+100$: $OUT \rightarrow IN$ All
Intercept	90.25^{***}	89.20^{***}	99.71^{***}
	(.99)	(1.36)	(.11)
Private	-3.76^{**}	-3.42^{**}	$-.66^{***}$
	(.91)	(1.65)	(.10)
Controls (Risk)	Yes	Yes	Yes
Pseudo R^2	.001	0.04	.000
N	14 343	1 606	166 791

Notes: All specifications are estimated using a Tobit regression. The intercept indicates the value of $x + 100$ in the public information treatment. Standard errors are in parenthesis, clustered at the subject level and are computed via bootstrapping.
$^{***}p \leq .01$, $^{**}p \leq .05$, $^{*}p \leq .1$

the private information treatment, subjects switch when x is about 3.75 points lower. In other words, subjects in public information treatment do not wait as long to exit (select OUT). This difference can be explained by the fact that the value of x is always available in the public information treatment, and therefore, the payoff to each strategy is clear. On the other hand, in the private information treatment selecting OUT reduces the information available. Thus, waiting to select IN suggests that subjects demand information on the relative payoff.

Result 4. *Subjects wait longer to select IN in the public information treatment.*

Specification (III) in Table 6 shows the value of x when subjects switch from OUT$\rightarrow$ IN. We find that in the public information treatment subjects wait longer to re-enter, and that when subjects switch, the value of x is close to its mean, as specified in equation (2). In the private information treatment, the subject is uninformed about x and therefore its particular value is not as meaningful. In this environment, subjects switch IN to learn the payoff to x, while in the public information environment subjects react to the value of x. The shorter duration of the OUT spell is consistent with a lower value of x observed in the OUT$\rightarrow$ IN decision in the private information treatment.

5 Discussion

In this paper, we study market entry decisions, where the payoff to entry is governed by a stationary AR(1) process, while market exit guarantees a constant payoff, under alternative information disclosure policies. In the public information treatment, we provide information on forgone payoffs to market entry, while in the private information treatment the subjects learn about market payoffs only if they enter the market. We introduce a nearly continuous environment where the risky payoff is updated every 0.5 seconds,[13] and omit point forecast elicitation.

Our results show that market entry is higher by about 10 percentage points when we omit information about foregone payoffs. While we observe strong variation in individual behaviour across treatments, the payoffs under both public and private information treatments are quite similar. The small difference may be due to the mean reversion process which governs the evolution of the risky payoff.[14] It would be interesting to study whether a process other than mean reversion leads to a different conclusion. However, we hypothesize that the behavior will be consistent due to demand for information in the private information treatment, but with larger payoff differences.

The analysis presented in this paper is motivated by modern online retail platforms, though we believe that our results are applicable to other settings. For example, managers could possibly increase participation rates in a venture by selectively disclosing payoff information to clients. Furthermore, one can extend our design to analyze how the decisions of others affect individual entry decisions. For example, in related bandit experiments, providing information about the decisions of others can help maximize profit (Hanaki et al., 2018) and, in an exponential bandit problem, increase free-riding on the information produced by partners (Hoelzemann and Klein, 2021). Alternatively, one can also study how almost continuous-time affects the evolution of prices in learning-to-forecast experiments (Hommes et al., 2005, and recently Arifovic and Petersen, 2017 and Kopányi et al., 2019). We leave these ideas for future research.

6 Acknowledgements

We are grateful for the insightful discussions and comments received from the Editors Ragan Petrie and Roberto Weber, two anonymous referees, Dan Friedman, John Duffy, Tibor Neugebauer, Luba Petersen, Nick Feltovich, Sébastien Pouget, Peter Bossaerts, Philip Drummond, Diego Aycinena, Pedro Romero, Robert Durand as well as from participants in numerous seminars and workshops. This project was approved by

[13]For other nearly continuous environments, refer to Cason et al. (2014), Friedman et al. (2015), and Bosch-Rosa (2018).

[14]In our sessions, the mean payoff to OUT is 0.67 standard deviations below the mean payoff to IN.

the human subjects committee at Monash University. Aleksei Chernulich is grateful for financial support from Tamkeen under the NYUAD Research Institute award for Project CG005.

References

Amromin, Gene and Steven A Sharpe, "From the horse's mouth: Economic conditions and investor expectations of risk and return," *Management Science*, 2013, *60* (4), 845–866.

Anufriev, Mikhail, Aleksei Chernulich, and Jan Tuinstra, "A laboratory experiment on the heuristic switching model," *Journal of Economic Dynamics and Control*, 2018, *91*, 21–42.

_ , Te Bao, and Jan Tuinstra, "Microfoundations for switching behavior in heterogeneous agent models: An experiment," *Journal of Economic Behavior & Organization*, 2016, *129*, 74–99.

_ , _ , Angela Sutan, and Jan Tuinstra, "Fee structure and mutual fund choice: An experiment," *Journal of Economic Behavior & Organization*, 2019, *158*, 449–474.

Arifovic, Jasmina and Luba Petersen, "Stabilizing expectations at the zero lower bound: Experimental evidence," *Journal of Economic Dynamics and Control*, 2017, *82*, 21–43.

Assenza, Tiziana, Te Bao, Cars Hommes, Domenico Massaro et al., "Experiments on expectations in macroeconomics and finance," *Experiments in Macroeconomics*, 2014, *17* (2014), 11–70.

Bergemann, Dirk and Johannes Hörner, "Should first-price auctions be transparent?," *American Economic Journal: Microeconomics*, 2018, *10* (3), 177–218.

_ and Juuso Välimäki, "Bandit problems," *The New Palgrave Dictionary of Economics: Volume 1–8*, 2008, pp. 336–340.

Beshears, John, James J Choi, Andreas Fuster, David Laibson, and Brigitte C Madrian, "What goes up must come down? Experimental evidence on intuitive forecasting," *American Economic Review*, 2013, *103* (3), 570–74.

Biele, Guido, Ido Erev, and Eyal Ert, "Learning, risk attitude and hot stoves in restless bandit problems," *Journal of Mathematical Psychology*, 2009, *53* (3), 155–167.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Diagnostic expectations and credit cycles," *The Journal of Finance*, 2018, *73* (1), 199–227.

Bosch-Rosa, Ciril, "That's how we roll: An experiment on rollover risk," *Journal of Economic Behavior & Organization*, 2018, *145*, 495–510.

Bouchaud, Jean-Philippe, Phillip Kürger, Augustin Landier, and David Thesmar, "Sticky Expectations and the Profitability Anomaly," *The Journal of Finance*, 2019, *74* (2), 639–674.

Cason, Timothy N, Daniel Friedman, and Ed Hopkins, "Cycles and instability in a rock–paper–scissors population game: A continuous time experiment," *Review of Economic Studies*, 2014, *81* (1), 112–136.

Chen, Daniel L, Martin Schonger, and Chris Wickens, "oTreeAn open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Coibion, Olivier and Yuriy Gorodnichenko, "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 2015, *105* (8), 2644–78.

Crosetto, Paolo and Antonio Filippin, "The bomb risk elicitation task," *Journal of Risk and Uncertainty*, 2013, *47* (1), 31–65.

Dwyer, Gerald P, Arlington W Williams, Raymond C Battalio, and Timothy I Mason, "Tests of rational expectations in a stark setting," *The Economic Journal*, 1993, *103* (418), 586–601.

Friedman, Daniel, Steffen Huck, Ryan Oprea, and Simon Weidenholzer, "From imitation to collusion: Long-run learning in a low-information environment," *Journal of Economic Theory*, 2015, *155*, 185–205.

Fuchs, William, Aniko Öry, and Andrzej Skrzypacz, "Transparency and distressed sales under asymmetric information," *Theoretical Economics*, 2016, *11* (3), 1103–1144.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer, "Expectations and investment," *NBER Macroeconomics Annual*, 2016, *30* (1), 379–431.

Glaser, Markus, Thomas Langer, Jens Reynders, and Martin Weber, "Framing effects in stock market forecasts: The difference between asking for prices and asking for returns," *Review of Finance*, 2007, *11* (2), 325–357.

Grosskopf, Brit, Ido Erev, and Eldad Yechiam, "Foregone with the wind: Indirect payoff information and its implications for choice," *International Journal of Game Theory*, 2006, *34* (2), 285–302.

Hanaki, Nobuyuki, Alan Kirman, and Paul Pezanis-Christou, "Observational and reinforcement pattern-learning: An exploratory study," *European Economic Review*, 2018, *104*, 1–21.

Hey, John D, "Expectations formation: Rational or adaptive or ...?," *Journal of Economic Behavior & Organization*, 1994, *25* (3), 329–349.

Hoelzemann, Johannes and Nicolas Klein, "Bandits in the Lab," *Quantitative Economics*, 2021, *12* (3), 1021–1051.

Holzmeister, Felix and Armin Pfurtscheller, "oTree: The bomb risk elicitation task," *Journal of Behavioral and Experimental Finance*, 2016, *10*, 105–108.

Hommes, Cars, Joep Sonnemans, Jan Tuinstra, and Henk Van de Velden, "Coordination of expectations in asset pricing experiments," *The Review of Financial Studies*, 2005, *18* (3), 955–980.

Kaya, Ayca and Qingmin Liu, "Transparency and price formation," *Theoretical Economics*, 2015, *10* (2), 341–383.

Kelley, Hugh and Daniel Friedman, "Learning to forecast price," *Economic Inquiry*, 2002, *40* (4), 556–573.

Kopányi, Dávid, Jean Paul Rabanal, Olga A Rud, and Jan Tuinstra, "Can competition between forecasters stabilize asset prices in learning to forecast experiments?," *Journal of Economic Dynamics and Control*, 2019, *109*, 103770.

Landier, Augustin, Yueran Ma, and David Thesmar, "Biases in Expectations: Experimental Evidence," 2019.

Magnani, Jacopo and David Munro, "Dynamic runs and circuit breakers: an experiment," *Experimental Economics*, 2020, *23* (1), 127–153.

Mokhtarzadeh, Fatemeh and Luba Petersen, "Coordinating expectations through central bank projections," *Experimental Economics*, 2020, pp. 1–36.

Schmalensee, Richard, "An experimental study of expectation formation," *Econometrica*, 1976, *44* (1), 17–41.

Wen, Quan, "Asset growth and stock market returns: A time-series analysis," *Review of Finance*, 2018, *23* (3), 599–628.

Yechiam, Eldad and Jerome R Busemeyer, "The effect of foregone payoffs on underweighting small probability events," *Journal of Behavioral Decision Making*, 2006, *19* (1), 1–16.

Appendix

A. Risk-elicitation results

We elicit risk attitudes following the protocol of Crosetto and Filippin (2013) implemented in oTree (Holzmeister and Pfurtscheller, 2016). The median boxes collected in the public (private) information treatment is 35 (40). Using a Wilcox test, we cannot reject that the distribution of boxes is equal across treatments (p-value of .14). Figure 7 shows the frequency of boxes collected in the two treatments.

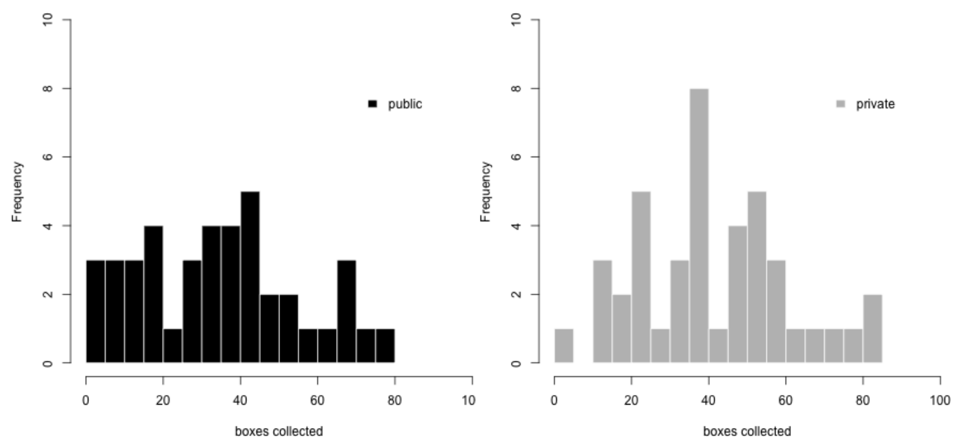

Figure 7: Histogram of boxes collected in the public information (left) and private information treatment (right)

B. Additional robustness checks

In this section, we limit our sample to sessions that experienced low realizations of the risky market. Specifically, we work with rounds in which the value of IN went below 80 for at least 40 ticks (of a total of 160). In this case, we obtain eight (nine) rounds for the public (private) information of a total of 78 (78). We replicate Table 4 using this restricted sample. The results are depicted in Table 7.

The treatment differences are consistent with what we observe in the regressions with the full sample (Table 4). The coefficient for *Private* is $-.12$, which confirms

Table 7: Linear probability model (restricted sample)

	(I) OUT	(II) OUT	(III) OUT	(IV) IN $\|x \geq r_f$	(V) OUT $\|x < r_f$
Intercept	$.51^{***}$	$.50^{***}$	$.48^{***}$	$.65^{***}$	$.65^{***}$
	(.02)	(.02)	(.03)	(.02)	(.03)
Private	$-.12^{***}$	$-.12^{***}$	$-.13^{***}$	$.08^{**}$	$-.20^{***}$
	(.02)	(.02)	(.02)	(.03)	(.03)
Round	—	—	$-.00^{***}$	—	—
			(.02)		
Controls (Risk)	No	Yes	Yes	Yes	Yes
R^2	.02	.02	.01	.00	.04
N	28 480	28 480	28 480	14 058	14 422

Notes: The Intercept captures the public information treatment. Standard errors are in parenthesis, clustered at the subject level and are computed via bootstrapping. $^{***}p \leq .01$, $^{**}p \leq .05$, $^{*}p \leq .1$

that players in the private information treatment chose IN more often compared to the public information treatment. In the public information treatment, as one would expect, players opt for OUT more often in the restricted sample (.51) compared to the full sample (.38).

Furthermore, the performance slightly improves when $x < r_f$ compared to the full sample in the public information treatment. Now, players opt for OUT with a frequency of .65 (compared to .62 in the full sample).

In sum, the main results presented with the full sample are in line with restricting our analysis to rounds in which players experienced a significant number of low realizations for the IN payoff.

C. Additional simulations

In this section, we replicate the first section of Table 1 for three different sub-samples of data. We divide all sessions in three groups according to the frequency with which risky investment (IN) outperforms the outside option (OUT): low, medium and high. To identify thresholds for each group we plot a histogram of all frequencies and set two threshold–0.6 and 0.7, that divide the sample into three groups of approximately the same size.

Table 8 contains descriptive statistics for three groups separately. Frequencies of choosing OUT in private treatment under Landier et al. (2019) expectations and RE differ approximately by 20% in all groups.

Table 8: Predictions for three groups of frequencies ($x > r_f$): low, medium and high.

	Frequency $(x > r_f) \leq 0.6$				$0.6 <$Frequency $(x > r_f) \leq 0.7$				$0.7 <$Frequency $(x > r_f)$			
	Landier et. al.		RE		Landier et. al.		RE		Landier et. al.		RE	
	public	private	public	private	public	private	public	private	public	private	public	private
Frequency OUT[a]	0.43	0.70	0.43	0.51	0.33	0.59	0.32	0.41	0.23	0.46	0.21	0.29
Freq. OUT$\vert x < r_f$[b]	0.76	0.84	0.75	0.74	0.73	0.81	0.70	0.68	0.66	0.78	0.63	0.62
Freq. OUT$\vert x \geq r_f$	0.84	0.44	0.85	0.70	0.88	0.52	0.89	0.73	0.91	0.63	0.92	0.81

Notes:
a. Frequency of OUT is computed using a tick count of when OUT strategy is observed.
b. Frequency of OUT conditional on the realized value of x being below the risky-free payoff.

D. Tables with OUT spells analysis

Table 9: Simulated predictions under sticky and extrapolative expectations (SEE) and rational expectations (RE) forecasting rules.

	SEE		*RE*	
	public	*private*	*public*	*private*
OUT spell (ticks)[a]				
10th	1	1	1	1
25th	1	2	1	2
Median	2	4	2	4
75th	4	6	4	6
90th	7	8	8	8

a. An OUT spell is counted as the number of ticks in which the player stays OUT without switching.

Table 10: Observed behavior in all rounds and in late rounds (11-20).

	All rounds		*Rounds 11-20*	
	public	*private*	*public*	*private*
OUT spell (ticks)[a]				
10th	1	1	1	1
25th	2	2	2	2
Median	5	3	5	3
75th	11	6	11	6
90th	20	12	22	12

a. An OUT spell is counted as the number of ticks in which the player stays OUT without switching.

E. Plots of sessions. (For Online publication)

Here, we present more examples and collect dynamics from all sessions of both treatments. We provide graphs of the fraction of players choosing IN in the experiment and evolution of IN option payoffs.

Treatment *public* information

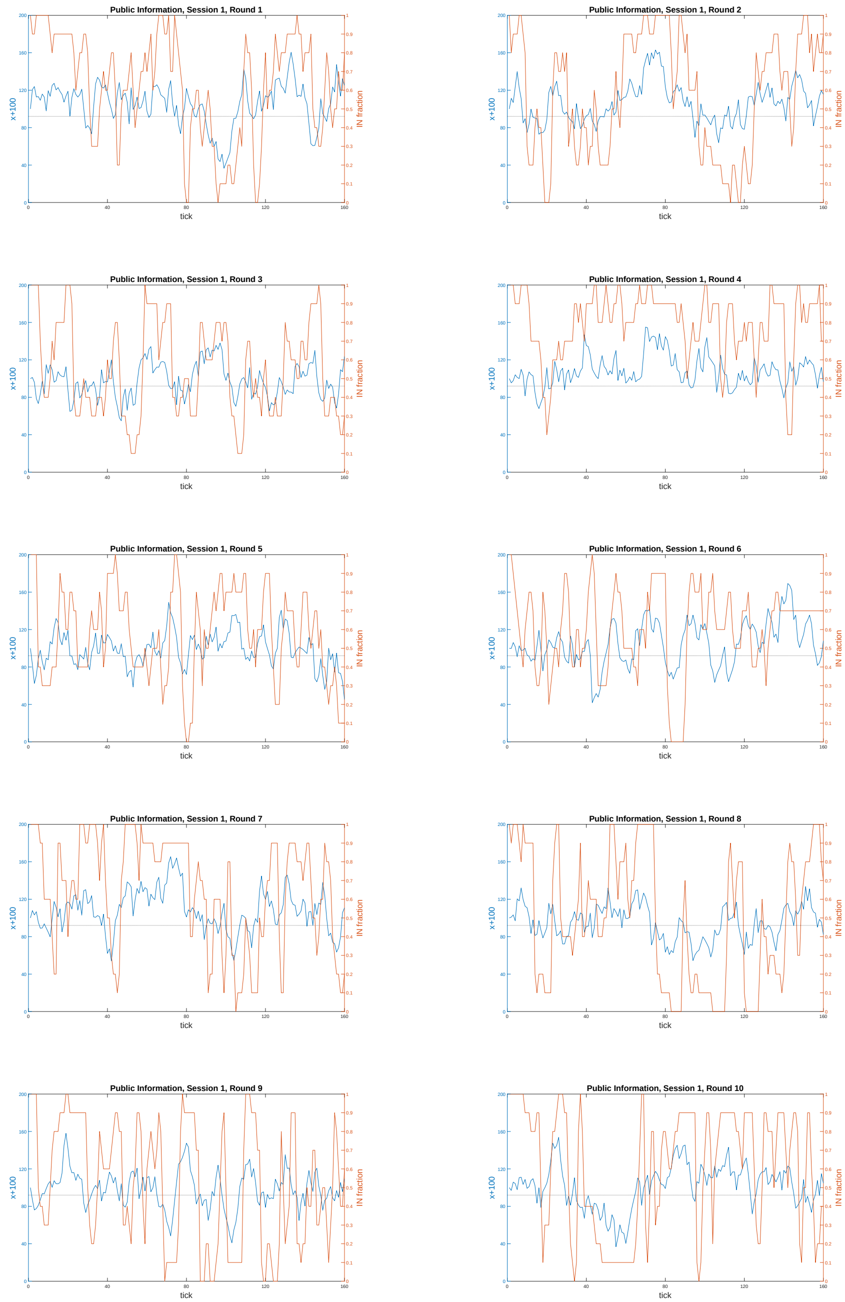

Figure 8: Time series of choices and x_t (public) - Session 1 Round 1-10

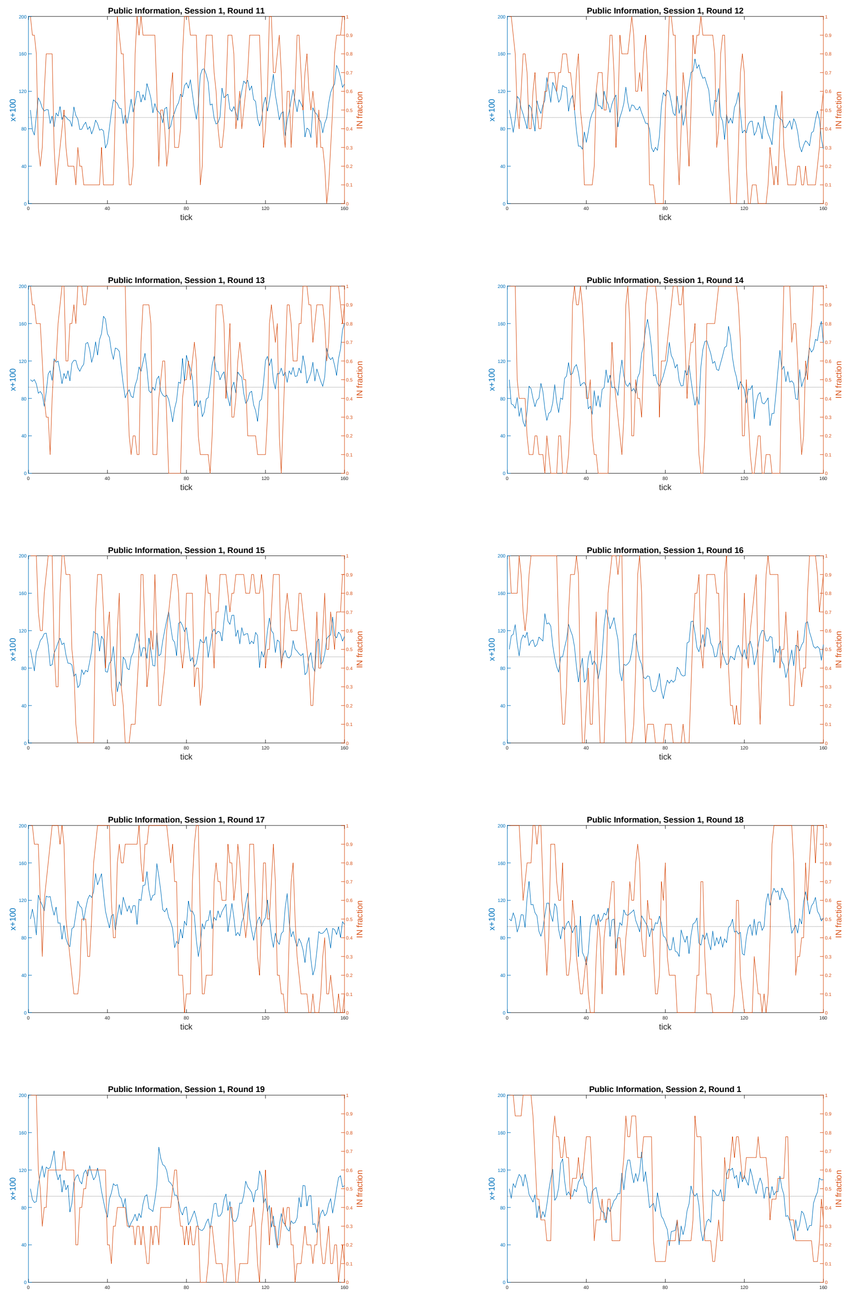

Figure 9: Time series of choices and x_t (public) - Session 1 Round 11-20, Session 2 Round 1

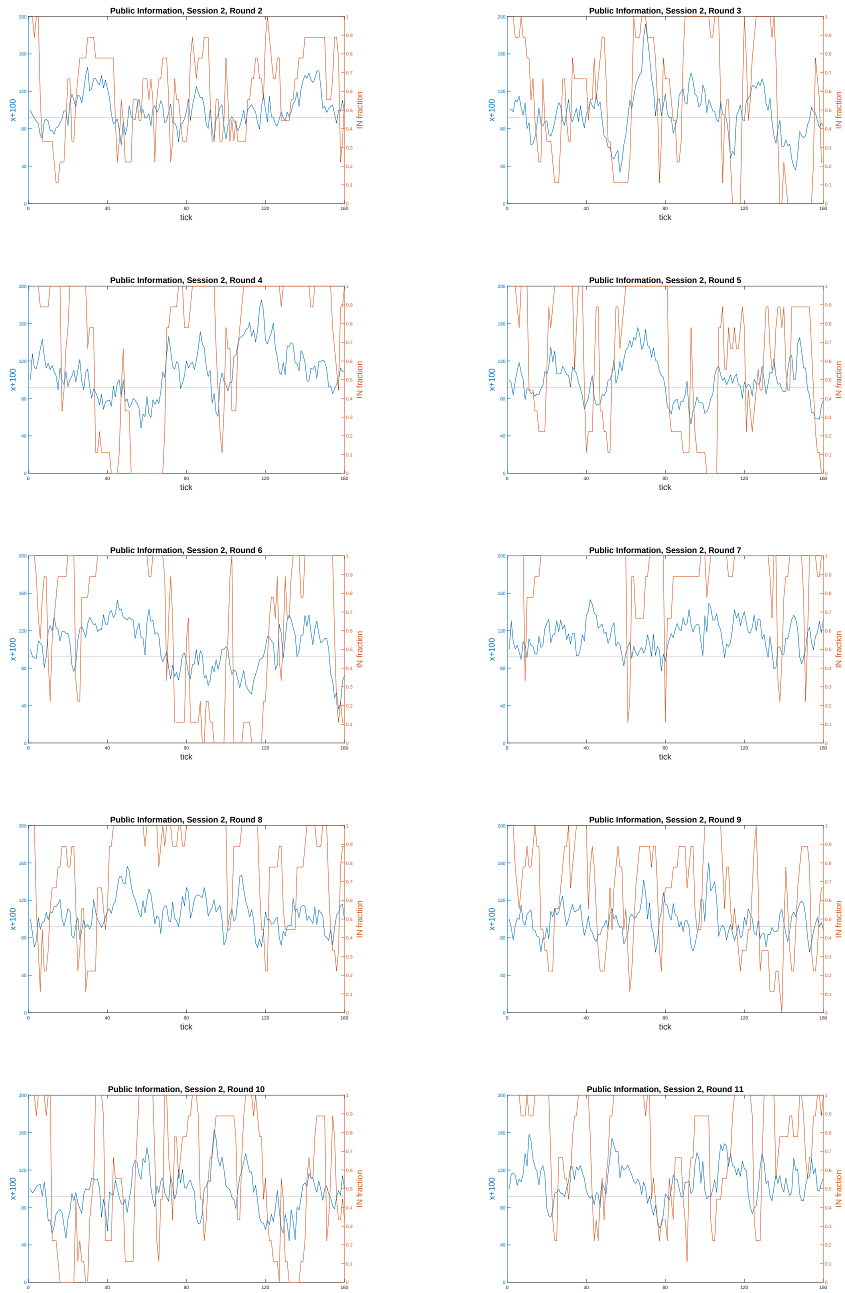

Figure 10: Time series of choices and x_t (public) - Session 2 Round 2-11

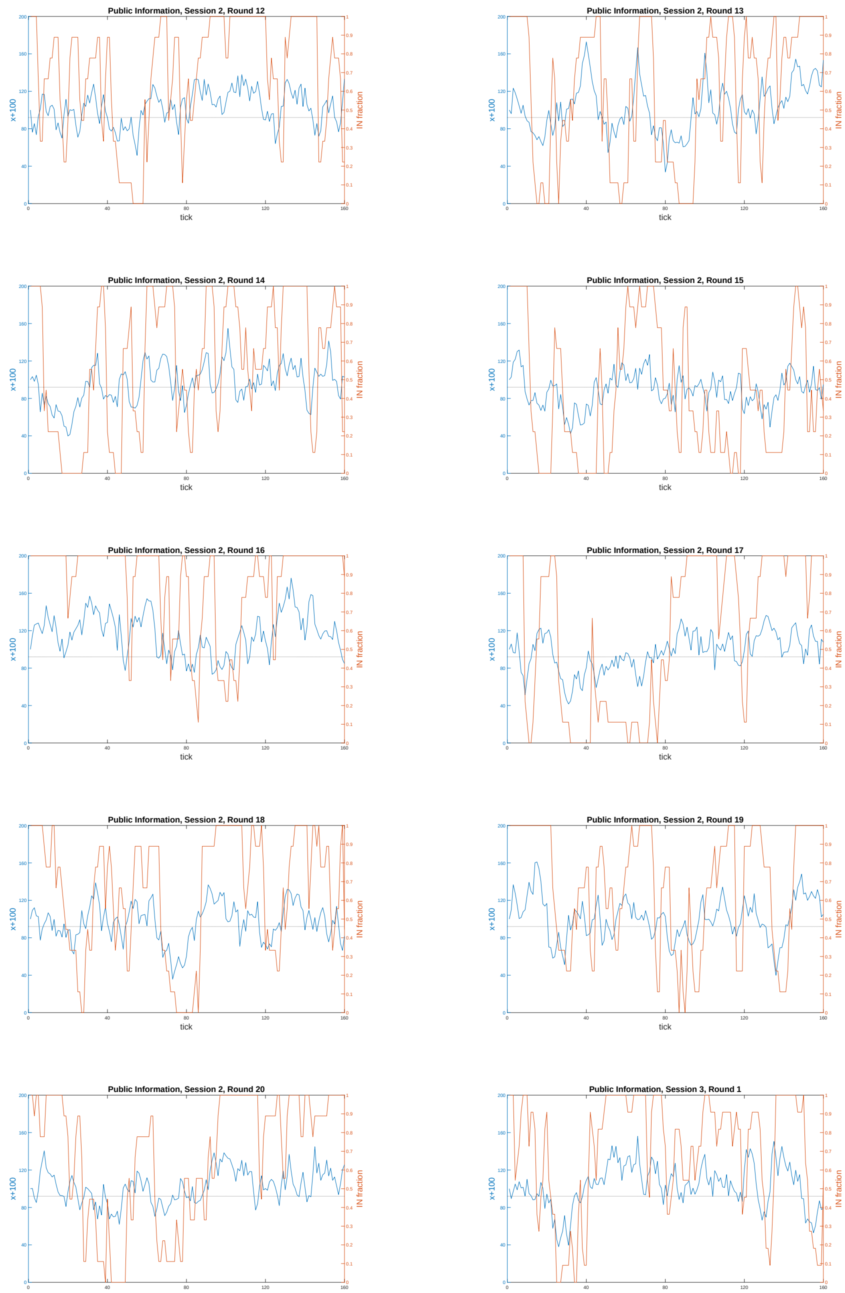

Figure 11: Time series of choices and x_t (public) - Session 2 Round 12-20 and Session 3 Round 1

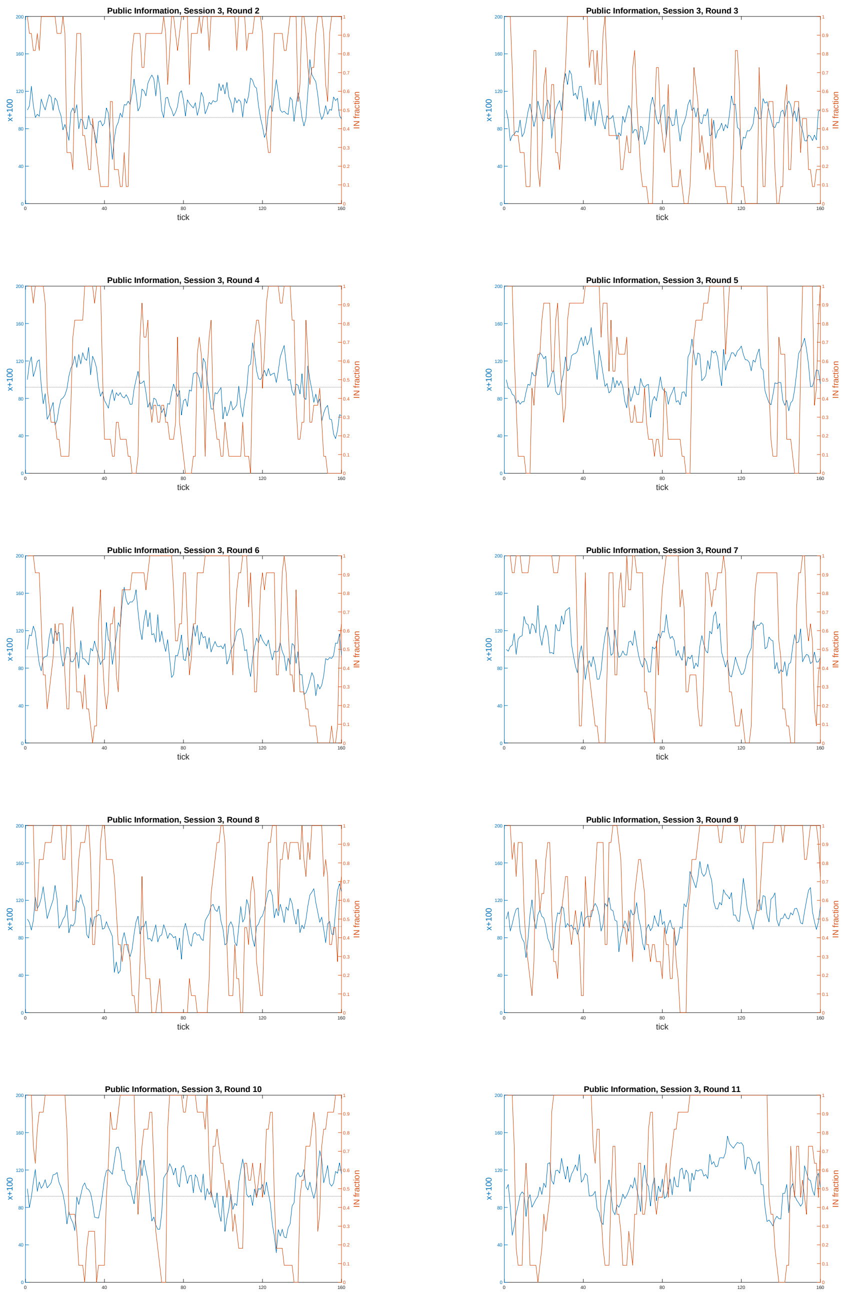

Figure 12: Time series of choices and x_t (public) - Session 3 Round 2-11

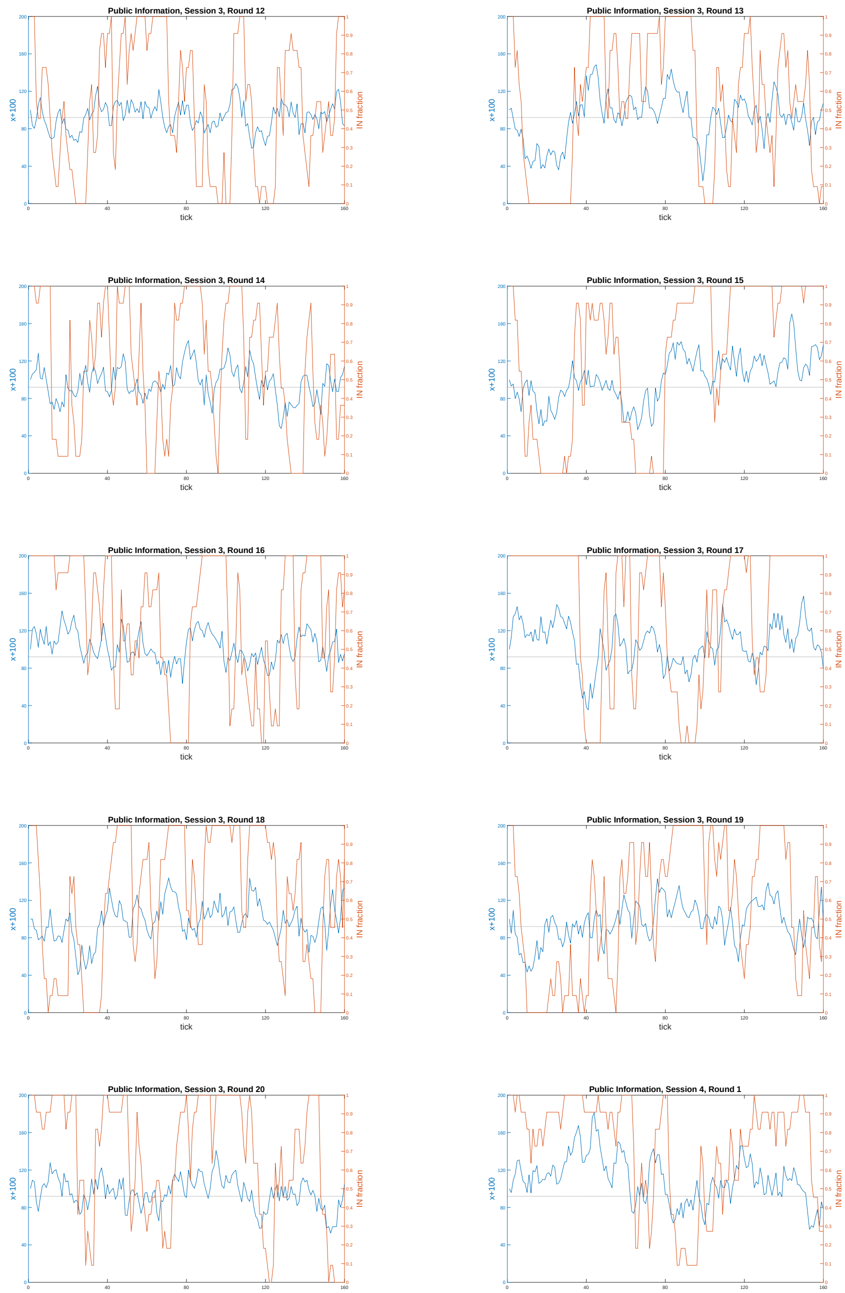

Figure 13: Time series of choices and x_t (public) - Session 3 Round 12-20, Session 4 Round 1

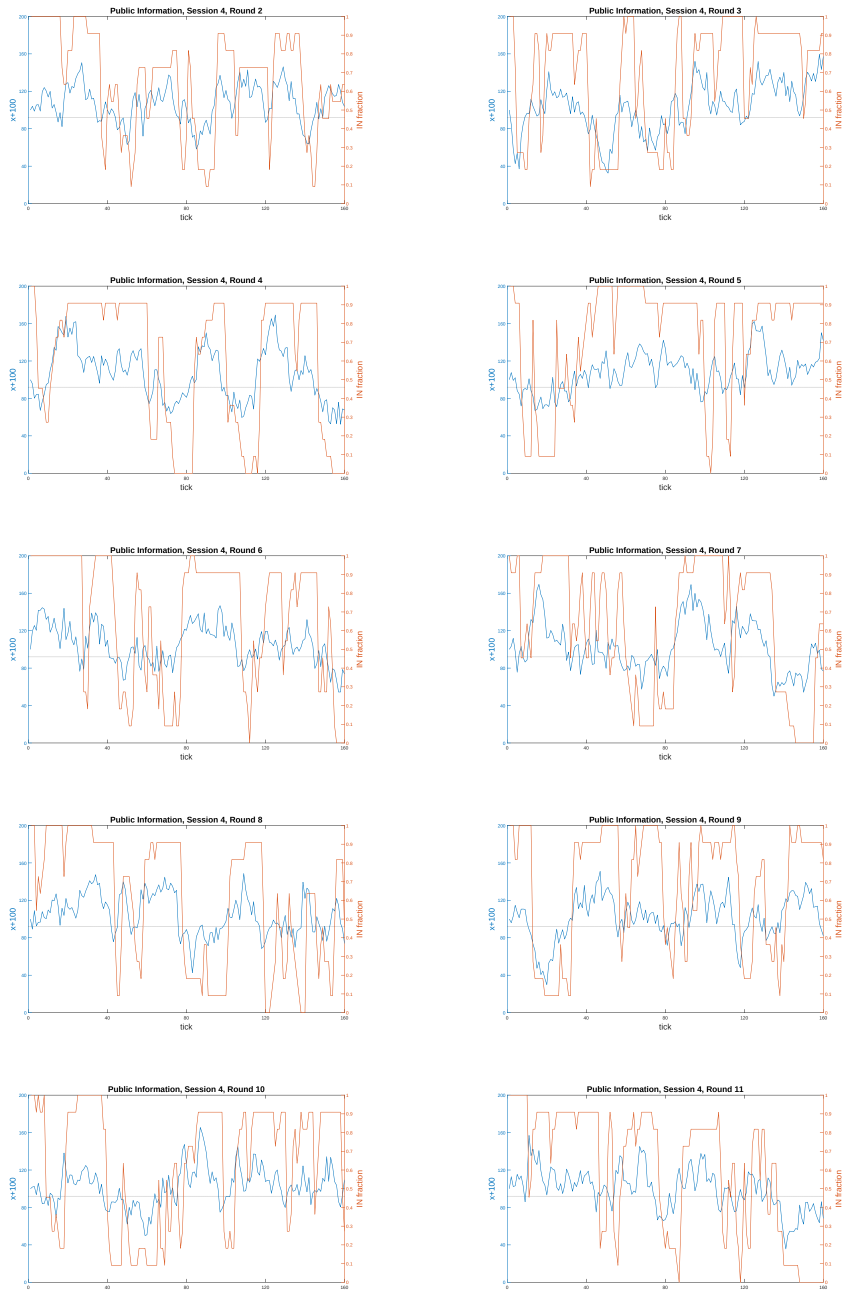

Figure 14: Time series of choices and x_t (public) - Session 4 Round 2-11

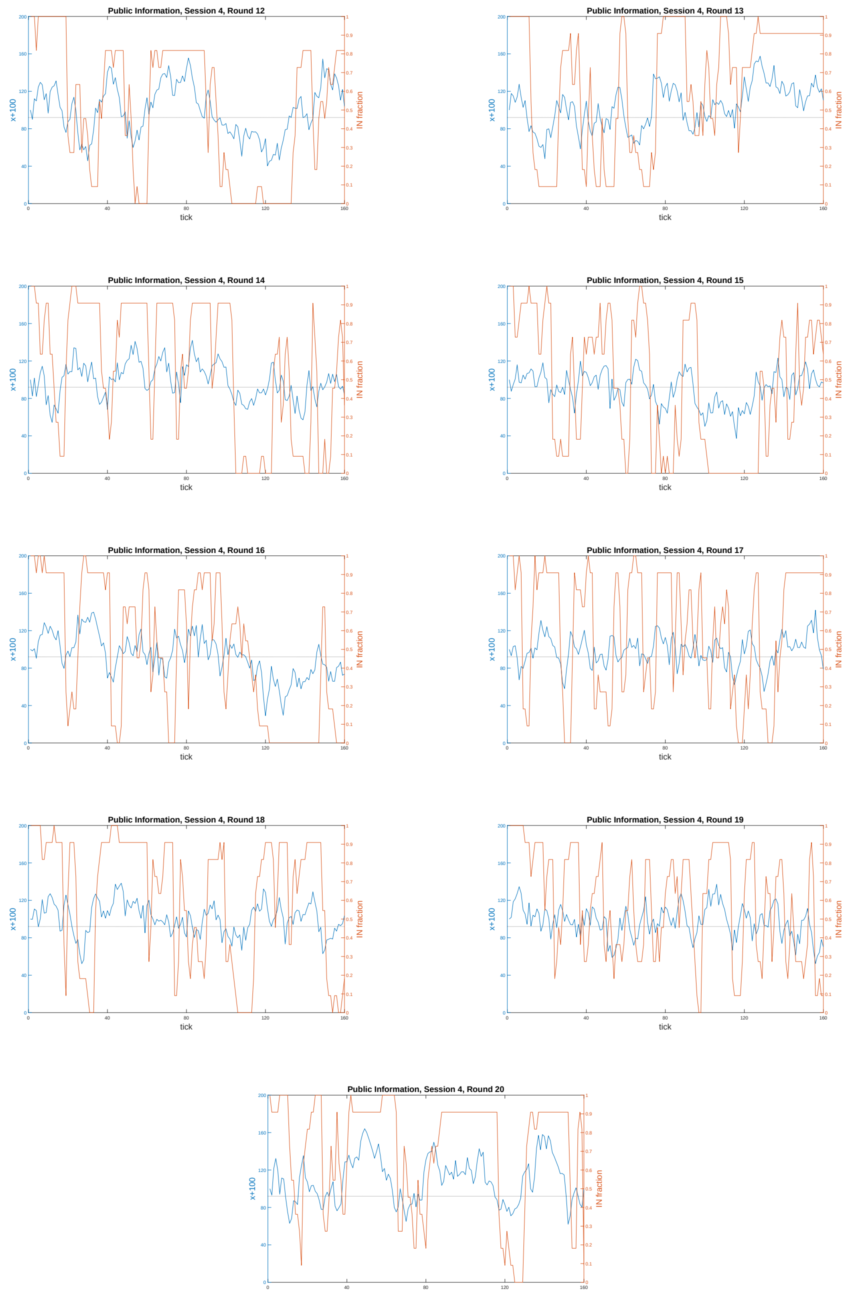

Figure 15: Time series of choices and x_t (public) - Session 4 Round 12-20

Treatment *private* information

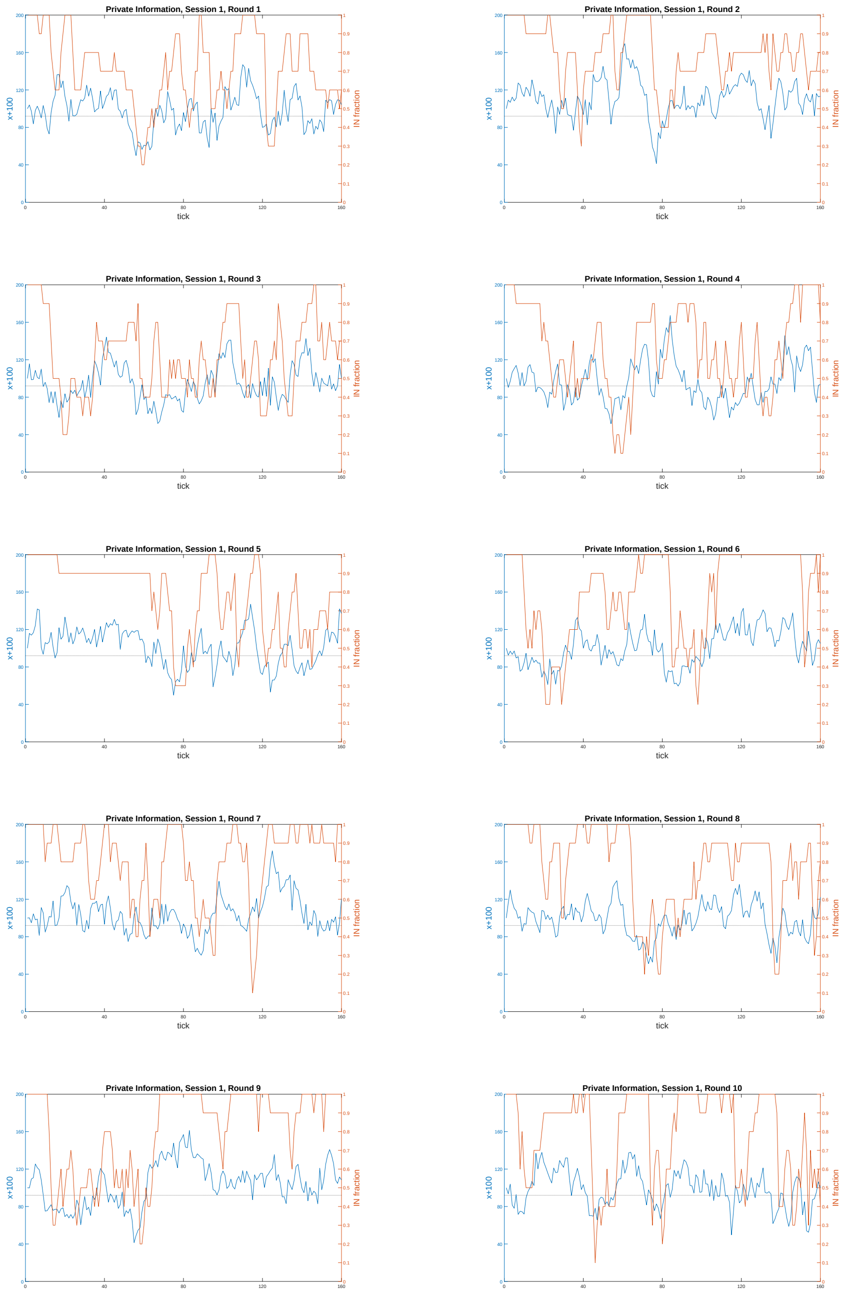

Figure 16: Time series of choices and x_t (private) - Session 1 Round 1-10

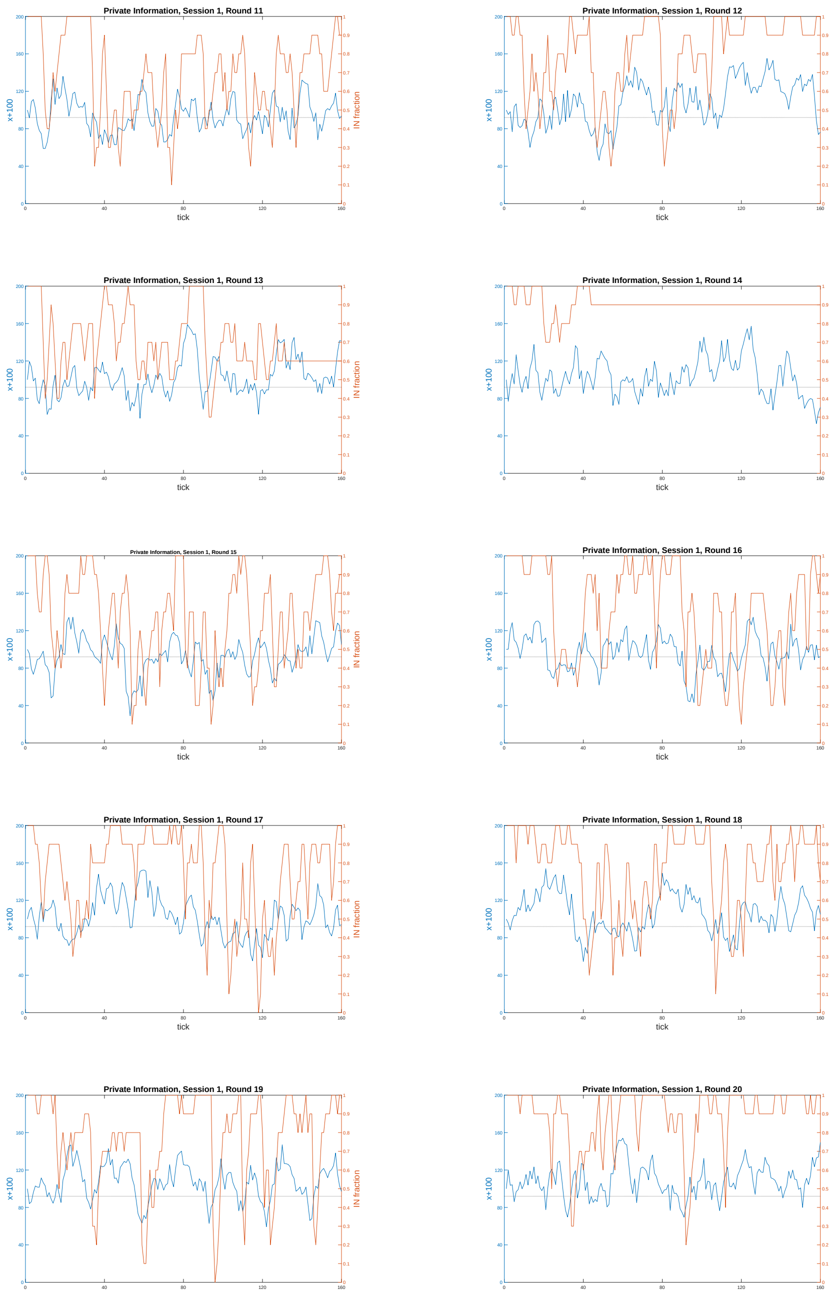

Figure 17: Time series of choices and x_t (private) - Session 1 Round 11-20

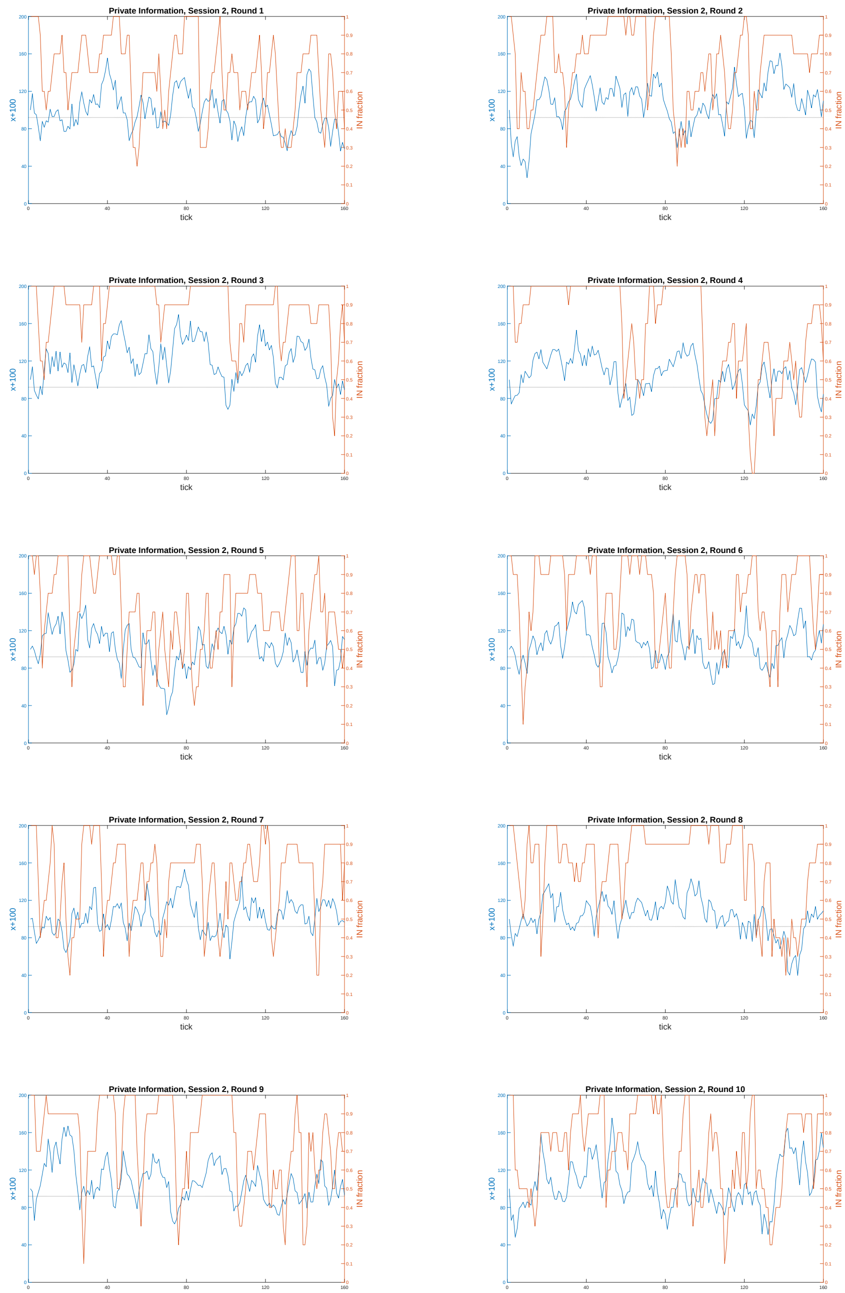

Figure 18: Time series of choices and x_t (private) - Session 1 Round 11-20

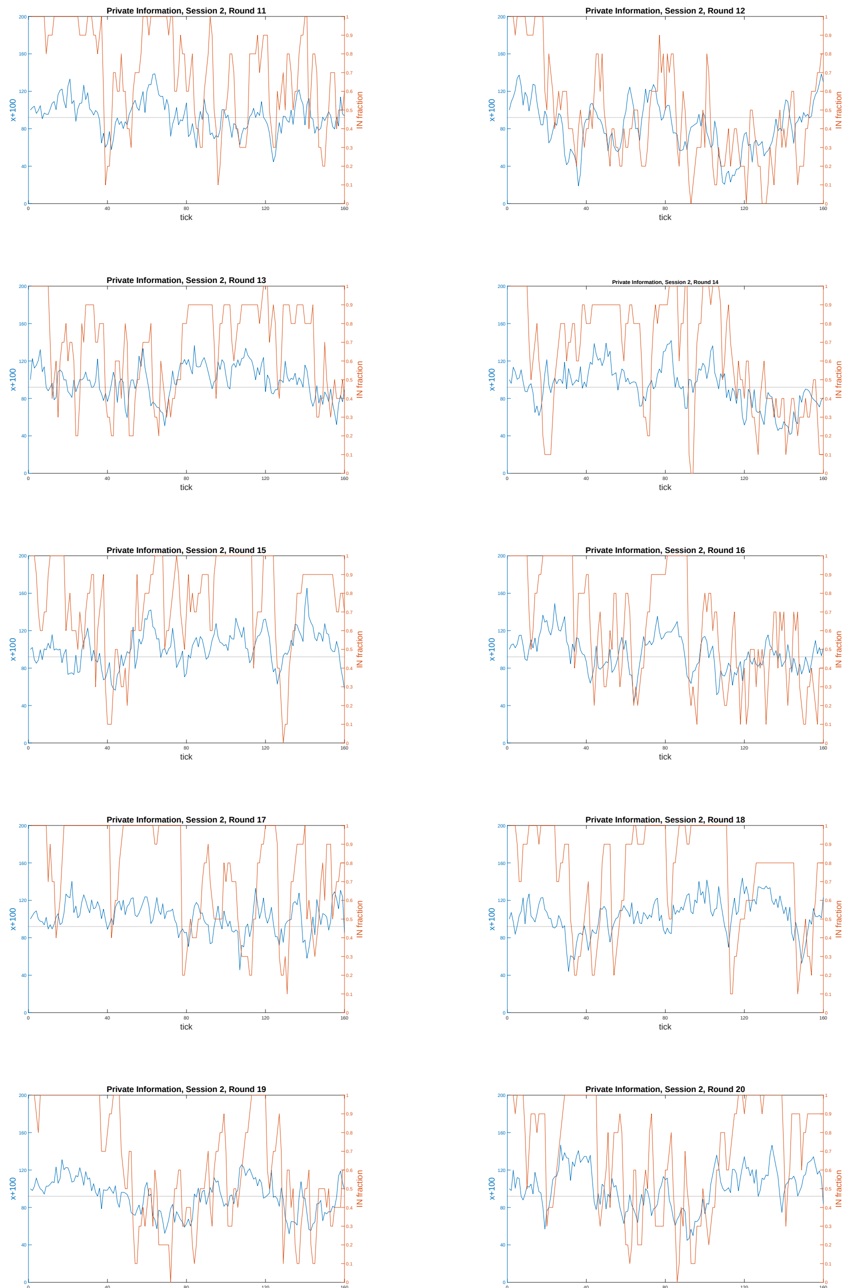

Figure 19: Time series of choices and x_t (private) - Session 2 Round 11-20

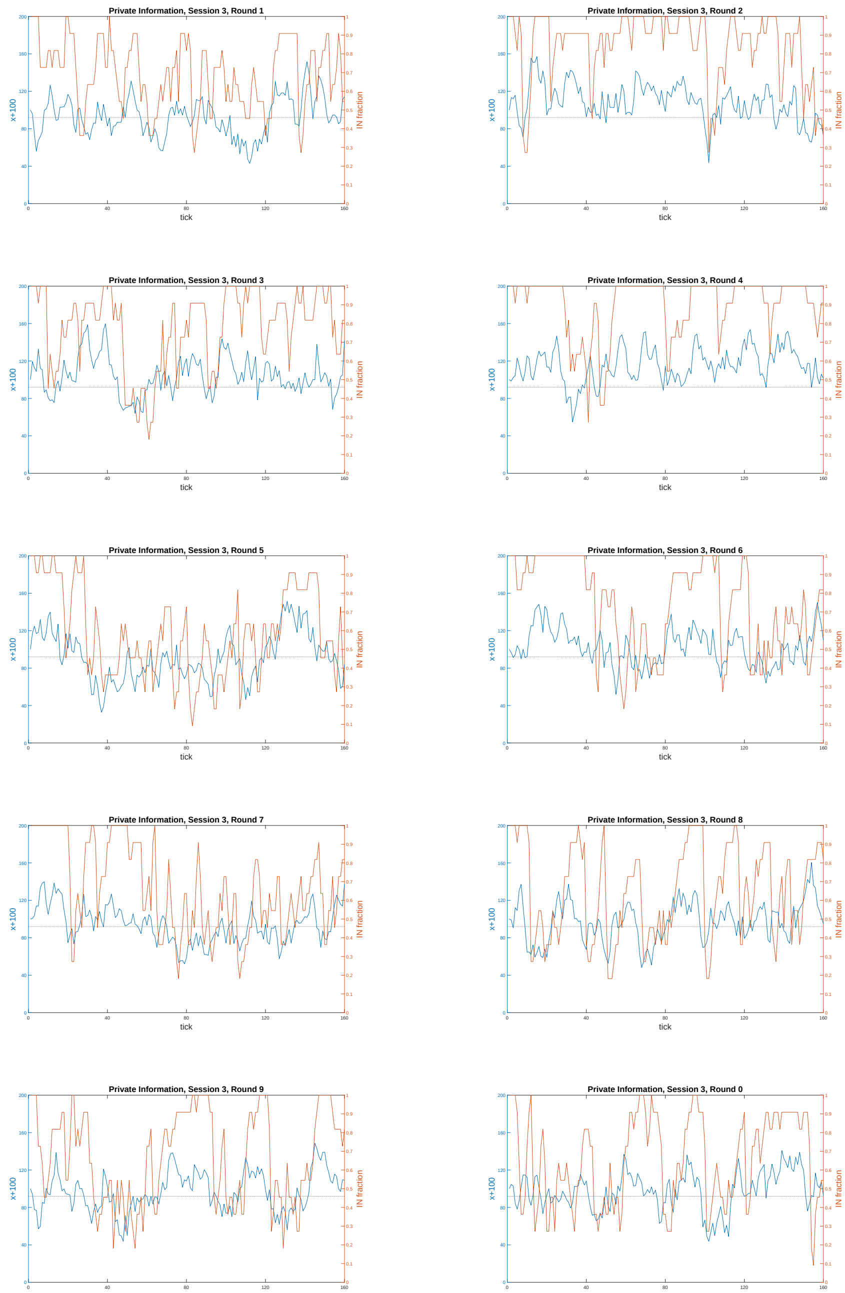

Figure 20: Time series of choices and x_t (private) - Session 3 Round 1-10

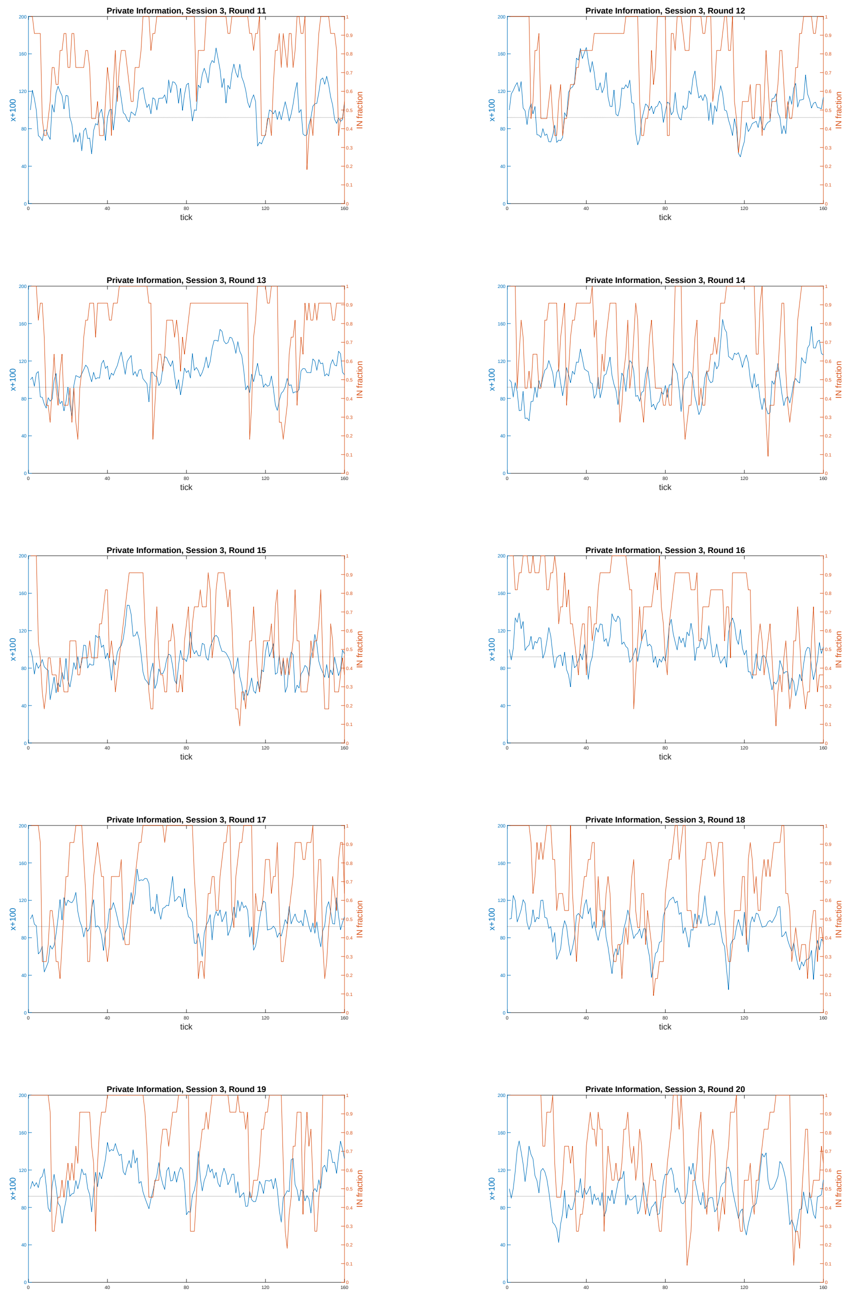

Figure 21: Time series of choices and x_t (private) - Session 3 Round 11-20

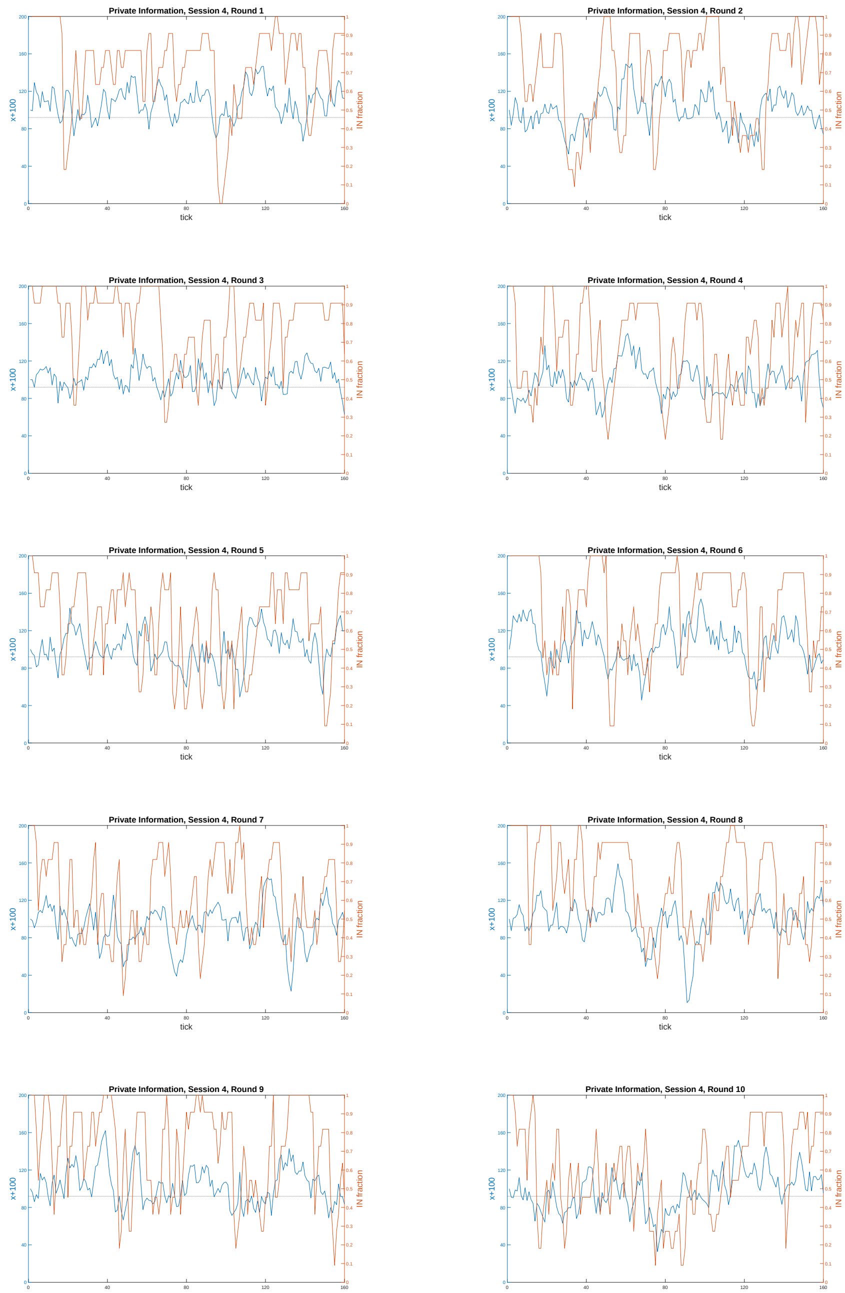

Figure 22: Time series of choices and x_t (private) - Session 4 Round 1-10

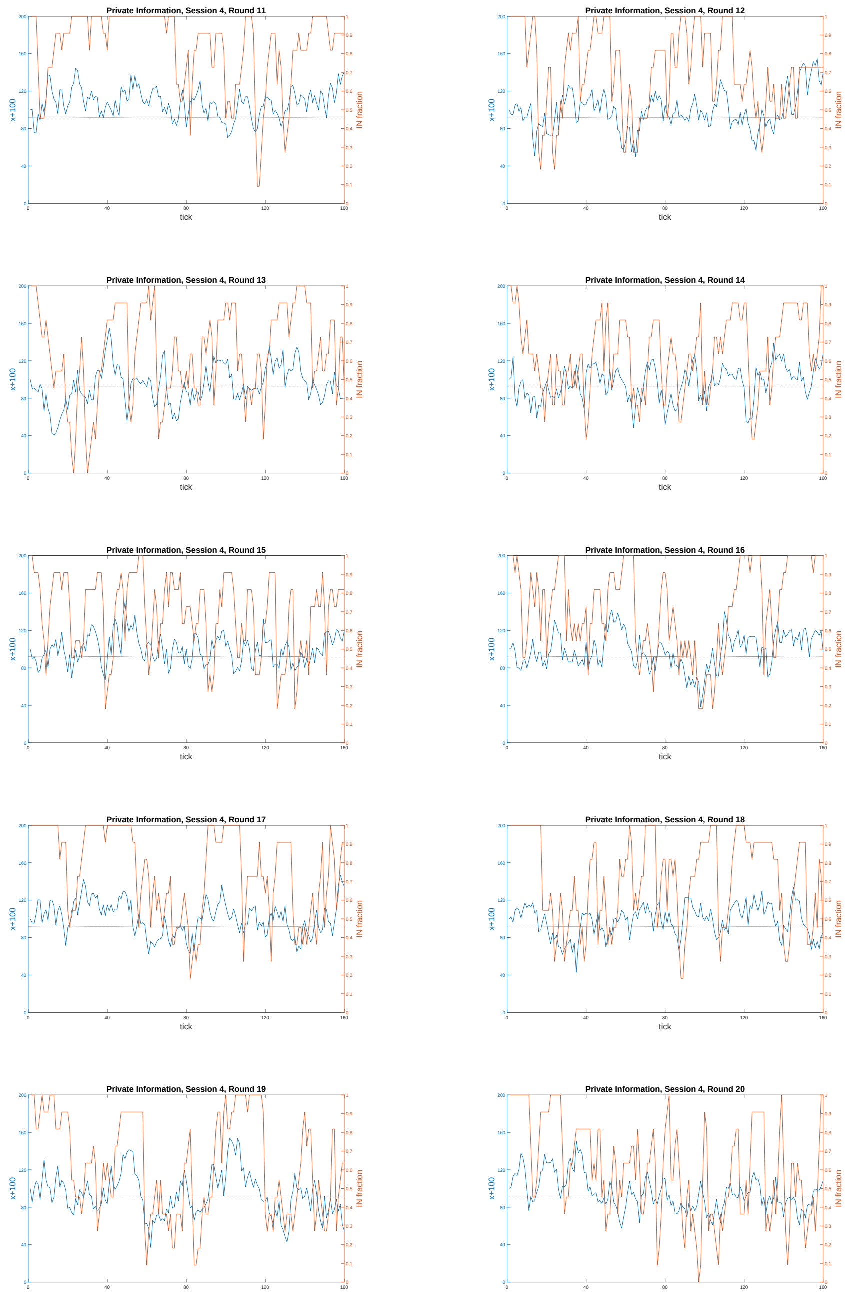

Figure 23: Time series of choices and x_t (private) - Session 4 Round 11-20