

The Times They Are a-Changing: Experimenting with Dynamic Adverse Selection

By FELIPE A. ARAUJO, STEPHANIE W. WANG AND ALISTAIR J. WILSON*

We examine a common-value dynamic matching environment where adverse selection accrues slowly over time. Theoretical best responses are therefore time varying, and where the prior experimental literature suggests that sequential environments might lead to greater understanding of adverse selection in this dynamic setting. However, while a sophisticated minority in our experiment do condition on time and are close to a best response, the majority use a stationary response, even after extended experience. In an environment with persistent uncertainty, our results indicate that sequentiality is insufficient for the large majority of participants to recognize the effects of adverse selection.

JEL: C72, C92, D90

Keywords: Dynamic adverse selection, sequentiality, experiments

The passage of time carries with it important strategic import across a myriad of economic settings. In labor markets, longer periods of unemployment can serve as a negative signal to prospective employers. For durable goods like houses, extended time on the market can reduce sellers' bargaining positions. In insurance markets, protracted spells without policy coverage can raise red flags with underwriters on new policy applications. Dynamic selection forces are present in day-to-day consumer interactions, where the expected produce quality at a farmer's market will fall through the day as early-bird shoppers pick through the best offerings. And it can be a force in more esoteric markets, such as for academic paper publication, where observables like the length of time a working paper has been in circulation can act as a negative signal.

While the above examples motivate how situations with dynamically accruing selection are commonplace, evidence for how decision makers respond to adverse selection is predominantly derived from behavior in static situations such as sealed-bid common-value auctions and market-for-lemons environments. While this literature certainly demonstrates initial failures of the Bayes-Nash equilibrium predictions, there is some evidence that experience pushes participants *toward* the equilibrium. For example, though many fall for the winner's curse early on in experiments, learning pushes them to bid less as they gain experience. Moreover,

* Araujo: Lehigh University, 266 Rauch Business Center, Bethlehem, PA 18015, fad220@lehigh.edu. Wang: University of Pittsburgh, Wesley W Posvar Hall, Pittsburgh, PA 15260, swwang@pitt.edu. Wilson: University of Pittsburgh, Wesley W Posvar Hall, Pittsburgh, PA 15260, alistair@pitt.edu. We would like to thank the following for their very helpful comments and input: Ignacio Esponda, David Huffman, Alessandro Lizzeri, Dan Levin, Muriel Niederle, Ariel Rubinstein, Emanuel Vespa, Lise Vesterlund and Georg Weizsäcker

the literature has identified sequentiality of the decision as an important predictor for when participants will adapt to selection. Though strategically more complex, there is evidence that dynamic environments such as ascending-price auctions remove winner's curse behavior. As such, while much of the behavioral literature is pessimistic about decision-maker's ability to comprehend and adapt to adverse selection, there is some room for optimism for the standard theory in dynamic, sequential settings.

Our experiments create a test-tube for examining these ideas. Specifically, in our dynamic matching environment, participants make sequential decisions where the best response should condition on the passing of time. In our main adverse-selection treatments participants are formed into groups with each member initially assigned an object with an independently assigned common value. While the assigned object's value is initially uncertain, participants receive information on its value at an exogenously determined point in time. At this point they are given a chance to exchange it for a rematching option. However, the rematching pool in our game is whatever object is not held by the other participants, and given-up objects become the rematching pool for subsequent movers. As such, adverse selection increases over time.

Given the accruing selection, *equilibrium* behavior is highly sensitive to the passing of time, and can be characterized by a cutoff strategy with greater willingness to keep low-valued objects at later points in time. While precise computation of the equilibrium best-response is undoubtedly complicated, a feature of our environment is that so long as all participants use a cutoff at each point in time—giving up objects below this level, keeping those above—then the precise best-response is insensitive to the level of other's time-varying cutoffs. In particular, even a simple belief that others use a fixed, stationary cutoff leads to essentially the same predicted behavior as equilibrium. However, while equilibrium-like play can be reached with minimal strategic sophistication in beliefs, it is also possible that equilibrium behavior is reached in the long run through experience. A simple requirement for convergence to equilibrium in our setting (as all information sets are visited with positive probability) is that participants allow for the possibility that the best response is time-varying. Experienced bad (good) outcomes that occur later (earlier) in the game must be considered distinctly when forming an expectation over the value of rematching.

Our experimental results find that aggregate behavior is *qualitatively* in line with the equilibrium predictions: participants are less likely to give up lower-valued objects as time passes. However, the point estimates are significantly different from the quantitative predictions. We show that this mismatch is driven by individual-level heterogeneity. In particular, a large proportion of participants fail to respond to the passage of time at all, exhibiting a stationary response even after extensive experience. The qualitative decrease overall is instead driven by a minority of participants (approximately a third), whose behavior is much closer to the equilibrium predictions, both qualitatively and quantitatively. Similar effects

and participant heterogeneity are found across four robustness treatments: three of which vary the feedback participants are given, and one that converts the dynamic adverse selection into a decision problem.

Taken together, our baseline and robustness results point to approximately two-thirds of the participants maintaining a stationary response that omits time as a factor. While these stationary participants show no movement towards learning the qualitative conditioning variable (time) as the session proceeds, they do show a significant unconditional learning effect, adjusting the level of their stationary response in reaction to experienced adverse selection, though without recognizing the pattern that their worst experienced outcomes tend to come at later points in time, and their better outcomes at earlier points.

While the behavior of the majority in our experiments sound a note of caution for the equilibrium prediction, there is also a glass-one-third-full interpretation of our results: A substantial minority do condition on time. This subgroup's late-session behavior is well-explained by the Bayes-Nash equilibrium, where their time-conditioning emerges early on in the experimental sessions. Moreover, we find direct evidence that their behavior comes about through a more direct understanding of the equilibrium mechanics: written statements in a peer-advice treatment indicate their ability to explain why the dynamic selection is occurring to others. Considering the marketplace selection effects likely to be present for professionals in finance, human resources, and actuarial disciplines, it is therefore plausible that our sophisticated minority is selected at much higher rates into the professional markets where understanding such forces is paramount. As such, behavior in such markets may be far closer to the equilibrium predictions.

After a brief literature review in the next sub-section, the remainder of the paper is structured as follows: Section two contains the experimental design and procedures, while section three presents the model and hypotheses. The main results are presented in section four, while section five discusses heterogeneity in the response, and how it shifts with experience. Finally, section six concludes.

A. Literature Review

Our study contributes to the growing theoretical (Eyster and Rabin, 2005; Jehiel, 2005; Jehiel and Samet, 2007; Jehiel and Koessler, 2008; Esponda, 2008) and experimental (Esponda and Vespa, 2014, 2018) literature on failures to account for others' private information in strategic settings. Experimental and empirical studies have primarily focused on three settings: auctions (see Kagel and Levin, 2002 for a survey), voting, and informed sellers. One well-documented case is the *winner's curse*, the systematic overbidding found in common-value auctions. A leading theoretical explanation for this effect is that bidders fail to infer decision-relevant information on the value for the item they are bidding on, conditional on a relevant hypothetical: their bid winning the auction. Modeling this, Eyster and Rabin (2005) allow for participants to best respond to others' *expected* action, failing to incorporate (or imperfectly incorporating, if partially cursed) how

others' actions are correlated with their private information.

A number of experimental studies have focused on determining the extent to which the winner's curse can be explained by this conjecture. For example, Charness and Levin (2009) have participants engage as buyers in an individual version of the informed-seller problem. Ivanov, Levin and Niederle (2010) have players bid in a common-value second-price auction where the value of the object is the highest signal in the group (the maximal game), thereby controlling for beliefs about their opponents' private information. Both studies continue to find deviations from the standard equilibrium prediction, suggesting that incorrect beliefs about other players' information are one source, but not the only one, of the failure to best respond.

For cursed behavior in voting, Esponda and Vespa (2014) find that most participants in a simple voting decision problem with minimal computational demands are unable to think hypothetically. That is, they do not condition their votes on the event that their vote is pivotal (and the subsequent information on the common state). Moreover, a smaller fraction of participants is also unable to infer the other (computerized) voters' information from their actual votes. Similarly, Esponda and Vespa (2018) found that most participants were not able to correctly account for sample-selection driven by other players' private information.¹ Our experimental setup expands this literature by offering a novel setting that can be used to explore various models of learning in connection with people's misspecified perceptions of the strategic setting.

Thus far, the experimental literature has focused on the importance of sequential rather than simultaneous play in reaching equilibrium-like behavior. For example, a significant share of participants who received explicit feedback about the computerized players' choices in the sequential treatment of Esponda and Vespa (2014) were able to correctly extract information from those observed choices. Similarly, players are more likely to adjust their thresholds to account for the selection problem if they were actually pivotal in the previous round (Esponda and Vespa, 2018). A number of experiments on sealed-bid vs. clock auctions have found closer to equilibrium bidding behavior when bidders are able to observe the decisions of other bidders (Levin, Kagel and Richard, 1996; Kagel, 1995). Carrillo and Palfrey (2009) find that second movers in a sequential version of a two-sided adverse selection setup behave more in line with equilibrium predictions than first movers (and players in the simultaneous version). Ngangoué and Weizsäcker (forthcoming) provides another example, where traders neglect the information contained in the hypothetical value of the price in a simultaneous market, but react to realized prices in line with standard theory in a sequential market. However, while the literature has identified sequentiality as the key

¹See also Enke (2020), who examines belief updating in selected samples, and Jin, Luca and Martin (forthcoming) who examine the response to empty messages in a disclosure environment. In both studies, sample selection creates a tension between the naive expectation and the correct one. In our setting we instead examine the within-participant response to an observed conditioning variable, namely *time*, and focus more on how participants learn about the selection forces.

to participants understanding the equilibrium thinking, our paper suggests that sequentiality alone might not be enough in more complex settings.

Our study also speaks to the substantial theoretical literature interested in dynamic adverse selection environments (Hendel, Lizzeri and Siniscalchi, 2005; Daley and Green, 2012; Gershkov and Perry, 2012; Chang, 2018; Guerrieri and Shimer, 2014; and Fuchs and Skrzypacz, 2015). One focus has been on asset markets where sellers have private information about the quality of the asset (Chang, 2018; Guerrieri and Shimer, 2014). Similarly, the current and past owners of an object in our setup could know the value of the object, while those who have never held the object do not. Although our players only make a binary choice on whether to keep the object or trade it for another in the early rounds of the experiment, they state a cutoff value for trading the object in later rounds, much like the price setting done by sellers and buyers in asset markets. Our experimental results suggest that these models should take seriously the questions of how equilibrium might plausibly be reached in the long-run if agents fail to identify relevant conditioning variables.

I. Design

We conducted 20 experimental sessions with 336 undergraduate participants. The experiments were all computer-based and were run at the Pittsburgh Experimental Economics Laboratory (PEEL). Sessions lasted approximately 90 minutes and payment averaged \$25.60, including a \$6 participation fee. In total we examine six distinct treatments: our main treatment and control; three variants of the main treatment that check robustness; and one variant of our main treatment that convert it into a decision problem. In the next two sections we focus on describing the design and results for the core treatment/control pair: a *Selection* (S) treatment, with dynamic adverse selection; and a *No Selection* (NS) control with no adverse selection and a stationary best response.

Our *S* and *NS* sessions both consist of 21 repetitions of the main supergame, broken up into: part (i) (supergames 1–5), which introduces participants to the environment; part (ii) (supergames 6–20) and part (iii) (supergame 21), which add strategy methods; and part (iv), which elicits information on risk preferences and strategic thinking. Before each part, instructions were read aloud to the participants, alongside handouts and an overhead presentation.² The environment in both treatments has a similar sequential structure with one key difference: in *S* supergames three randomly chosen participants are matched together into a group to play a game with a shared rematching pool; while in the *NS* supergames an individual participant makes choices in an isolated decision problem where the rematching pool is unaffected by others. We next describe the *S*-treatment environment in more detail before coming back to describe the *NS*-treatment.

²Detailed instructions, presentation slides, and screenshots of the experimental interface are available in online Appendices B and C.

A. Selection (S) Treatment

The primary uncertainty in each of our supergames is generated by drawing four numbered objects, labeled as *balls A–D*. Each ball is assigned a random value θ through an independent draw over the integers 1–100 (with proportionate common monetary values from \$0.10 to \$10.00) according to a distribution F with an expected value of 50.5.³ A group of three players are randomly assigned a mover position, which we refer to as *first*, *second* and *third* mover. Each group member is initially assigned one of the four balls (without overlap). As the three players each hold an initial ball, a single ball remains unheld; this unheld ball is the initial rematching object in our game.

An example matching is illustrated in Figure 1, where the first line shows an initial matching. In the illustrated example the first mover is matched to *Ball B*, the second mover to *Ball A* and the third mover to *Ball D*, or $\langle 1B, 2A, 3D \rangle$ for short; hence the leftover rematching ball is *Ball C*.

Though players know which of the four balls they have been initially assigned, they do not start out knowing its value, nor the balls (or values) held by other group members. Information is exogenously provided through the following process: In each round, the players each flip a fair coin, where if it lands heads they learn their held ball's value; if tails they do not and must wait for the next round to flip again. Rounds are broken into sub-periods where the players move sequentially in order of their mover role: first, second, then third. Finally, if a player has not seen their held ball's value in rounds one or two (flipping tails in both) the held ball's value is revealed to them in round three with certainty. Each participant makes only one payoff-relevant decision per supergame. This occurs in the round that they see their ball's value, where they either:

Keep hold-on to the currently held known-value ball as the final supergame outcome.

Switch Take the unknown rematching ball (whichever ball is currently unheld by another player) as the final supergame outcome, and give up their currently held ball (which becomes the rematching ball for subsequent movers).

To make clear the process and intuition of the game, consider the example illustrated in Figure 1(A). The figure takes the point of view of the first mover, where elements in black represent information that is known to the first mover at each point in time, while elements in gray represent unknowns. In the example, though the first mover knows she is holding *Ball B* in the first round, its value remains unknown to her as she fails the coin flip. The first mover does *not* know the balls the other two players are initially holding, nor their coin flip outcomes,

³The distribution used in our experiments is a discrete uniform with additional point masses at the two extreme points. The probability mass function puts a $\frac{51}{200}$ mass on the two values 1 and 100 and a $\frac{1}{200}$ weight on each of the integers 2–99. This distribution was chosen to make the selection problem more salient, and to generate sharper equilibrium predictions

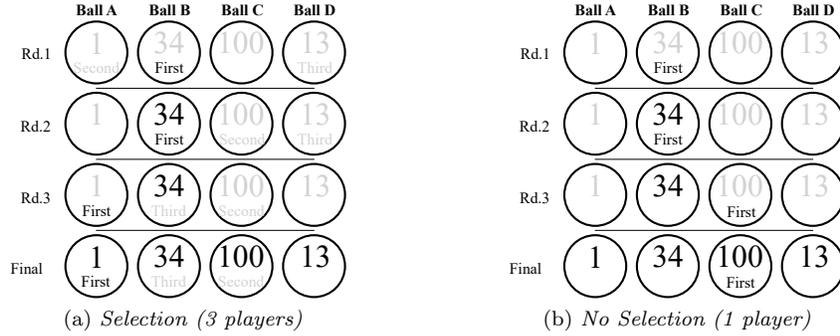


Figure 1. : Example Supergames

nor their decisions. She only knows that they are present and that their decisions are potentially affecting the rematching ball.

In the illustrated example, the initial matching is $\langle 1B, 2A, 3D \rangle$. Unknown to the first mover, the second mover in round one flips a head and sees his held-ball's value is one, and decides to switch, while the third mover flips tails and does not learn her value. The interim matching at the beginning of round two is therefore $\langle 1B, 2C, 3D \rangle$ where the rematching ball is the released *Ball A*. In the second round, the first mover flips a head, and sees that her held ball's value is 34. She decides to switch, and rematches to the currently unheld *Ball A* (the matching becomes $\langle 1A, 2C, 3D \rangle$). After her round-two decision (and again, unknown to the first mover) the second mover does not act as he has already made a decision, while the third mover flips a head and decides to give up her 13-ball, rematching to *Ball B* that was given up by the first mover (moving the match to $\langle 1A, 2C, 3B \rangle$). By round three, all three participants have made a decision, and so the final matching is $\langle 1A, 2C, 3B \rangle$. At the end of the supergame all four balls' values are made common knowledge—though which balls other players are assigned to is not—and the first mover learns that the ball she rematched to has a value of one.

Supergames one to five exactly mirror the procedure above: Participants make a binary decision to *keep* or *switch* only in the round where their assigned ball's value is revealed. The second part adds a partial strategy method. Specifically, in supergames 6–20 participants are asked to provide a cutoff in each round, indicating the lowest value for which they would keep their held ball, contingent on seeing its value that round. If they receive information, the decision to keep or switch is then resolved according to the stated cutoff; if they do not get information, they must wait until the next round, when they provide another cutoff. As in supergames 1–5, participants only implemented decision is in the round where the ball's value is revealed. Finally, in part (iii) we use a complete strategy method in which participants are not told whether or not information was received in each round, and we collect their minimum-acceptable cutoff values in

all three rounds of supergame 21 with certainty.^{4,5}

Strategic feedback on the other participants is purposefully limited in our baseline *S* game, where we examine the effects of alternative feedback in our robustness treatments. At the end of each *S*-supergames, each participant sees the values of the four drawn balls, the particular ball/value they are holding at the end of the supergame and (if they switched) the ball/value they were initially assigned to. Participants *do not see* strategic feedback. That is, they observe neither the identity of the balls held by the other two group members at the end of the supergame, nor the balls others were initially holding, nor their choices.

Participants' final payment for the session is the sum of: a \$6 show-up fee; \$0.10 times the value of their final held ball (\$0.10 to \$10.00) from two randomly selected supergames from 1–20; and \$0.10 times the value of their final held ball in supergame 21. Excluding the part (iv) payments the experiment therefore has a minimum possible payment of \$6.30 and a maximum of \$36.00.

B. No Selection (NS) Control

Our *No Selection* (NS) control is designed to have the same structure as the *S*-treatment game, except that we turn off the dynamic adverse selection. This is achieved by making a single change to the environment: each group has just one member. As such, each supergame is a decision problem with a single participant in the role of first mover. As there are four balls and only one of them is held, the rematching pool is the three unheld balls. Whenever the first-mover sees their initially assigned ball's value, if they decide to switch their ball, they are randomly rematched to one of the three unheld balls. Our *NS* sessions therefore replicate the incentives and timing from the *S* sessions, but without the possibility of other group members decisions contaminating the rematching pool. We illustrate a parallel example supergame for the NS environment in Figure 1(B).

II. Model and Hypotheses

The environments described above are dynamic assignment problems over a finite set of common-value objects (the balls). The objects (the long-side) are initially assigned independently of value to the short-side of the market (the participants). Private information on the held object's value arrives randomly over time, according to an exogenous process (in the experiment, the coin flips).

⁴In expectation one-quarter of participant data in supergames 6–20 will have data from all three round cutoffs, one quarter with cutoffs from rounds one and two only, and one half of the data only has an elicited first-round cutoff.

⁵In part (iv) at the end of each session we collect survey information, and incentivize the following elicitations: (a) risk preferences (using the dynamically optimized sequential experimentation method from Chapman et al. 2018); (b) a three-question Cognitive Reflection Test (Frederick, 2005); and (c) a continuous version of the Monty Hall problem. One participant per session was selected for payment in the part (iv) elicitation.

With a single decision maker, the rematching pool is never affected by other participants' decisions. As such, the risk-neutral prediction in our *NS* treatment is that participants are stationary and use a minimal acceptable cutoff of $\mu_{NS}^* = 51$ for retaining a ball. The cutoff rule gives up balls valued 50 or below (beneath the expected value of 50.5) and keeps balls valued 51 or higher.

In order to make clear the difference between control and treatment, from this point forward we normalize all values by subtracting 51 (for both outcomes and cutoffs), measuring all responses relative to the stationary prediction in the control. As such, the risk-neutral prediction for our control is normalized to a zero cutoff in each round.⁶

HYPOTHESIS 1 (NS-Treatment): *Participants use a stationary cutoffs in the NS treatment (where the risk-neutral prediction is for a zero cutoff in all rounds)*

In contrast to the control, when there are multiple players the arrival of private information over time leads to accruing adverse selection on the rematching pool. Whenever other players give up objects with observed low values, and keep objects with high values, the rematching pool will become selected. As private information arrives slowly over time, adverse selection accrues dynamically. In early periods, it is less likely that others have received information, so the rematching pool is less likely to be selected. In later periods, it is more likely that others have received information, increasing the likelihood that the rematching pool was selected.

Because the environment is sequential, the equilibrium predictions can be solved for inductively, starting from the first mover in round one, and where the best response at each point is entirely backward looking. This is in contrast to many other situations that examine “cursed” behavior over hypothetical future events. For example, in common-value auctions, optimal decision-making requires the bidder to act as if concentrating solely on the hypothetical event that their bid wins the auction, inferring information contained in winning on others' signals and hence the object's value. Similarly, in common-value voting, the voter has to focus on the hypothetical event that her vote is pivotal. In our environment, the optimal response is instead conditioned on an experienced event—the passing of time—where hypothetical thinking relates to how other participants have acted in *previous* periods.

From the point of view of player i making a decision at time t there are two distinct random variables: the initially assigned object value θ_i^0 , an iid draw from the CDF F_θ ; and the rematching object, θ_i^R , with a distribution that varies over time due to other players' rematching choices. Once the player's held-object value becomes known, the optimal risk-neutral response is to give up held objects if their

⁶Risk-aversion (risk-lovingness) will lead to negative (positive) cutoffs under the normalization. For the control treatment all expected-utility decision makers are predicted to be stationary across rounds, irrespective of risk preferences, as the passing of time conveys no information on the rematching value.

value is lower than the expected value of rematching, and to keep higher values.⁷ The rematching random variable θ_t^R and the policy cutoff μ_t^* are calculated inductively from a first-mover seeing her object's value in the first round ($t = 1$).⁸ For the base case, the rematching option for the first mover in round one (θ_1^R) is just an *iid* draw from the initial-value distribution F_θ , as no other participant has had a chance to exchange their object yet. The policy for a player moving at $t = 1$ (first-round, first-mover) is identical to the cutoff to the *NS* rule: a zero cutoff, if risk neutral.

For the inductive step we define the event that the player moving at time t sees their value as \mathcal{I}_t , and the joint event that they both see their value *and* choose to switch as \mathcal{S}_t (with complement \mathcal{S}_t^C). Given the rematching random variable θ_t^R and the policy cutoff μ_t^* from period t , the rematching random variable for a player observing their held-value in period $t + 1$ is defined by:⁹

$$\theta_{t+1}^R | \mathcal{I}_{t+1} = \Pr \{ \mathcal{S}_t; \mu_t^* | \mathcal{I}_{t+1} \} \cdot (\theta | \theta < \mu_t^*) + \Pr \{ \mathcal{S}_t^C; \mu_t^* | \mathcal{I}_{t+1} \} \cdot (\theta_t^R | \mathcal{I}_{t+1}, \mathcal{S}_t^C).$$

The optimal policy cutoff for the player making a decision at $t + 1$ is then given by the expected value of rematching, $\mu_{t+1}^* = \mathbb{E}(\theta_{t+1}^R | \mathcal{I}_{t+1})$.

Figure 2 illustrates the unique risk-neutral PBE prediction by round and role. The predictions for the *S*-treatment decrease from a predicted cutoff of 0 for the first mover in the first round (equivalence to the *NS* control), to a cutoff of -28 for the third mover in the third round.¹⁰ To put the extent of the adverse selection in context, if the other two agents were fully informed on the other three values and perfectly sorted so the remaining unheld ball was the worst of the three, its expected value would be $\mu_{(3)} = -34.6$. As such, by the end of the last round, over 75 percent of the adverse selection possible with perfect sorting has occurred.

Within each role the PBE predictions indicate strictly decreasing cutoffs, reflecting the increased adverse selection as the game unfolds.¹¹ This decreasing pattern across rounds holds in equilibrium for both risk-loving and risk-averse preferences.

Moreover, decreasing cutoffs across rounds is predicted even without equilibrium beliefs on others' behavior. For example, a simple belief that other par-

⁷In our experiments the action set is discrete as the normalized values are in $\Theta = \{-50, -49, \dots, +49\}$, and so the cutoff can be summarized instead by $\min\{x \in \Theta : x \geq \mu_t^*\}$, the minimal acceptable ball value.

⁸For the theory, instead of indexing time by the round number, we do it by round-mover. So the first-mover in round one is $t = 1$; the second mover in round one is $t = 2$; the third mover in round one is $t = 3$; the first mover in round two is $t = 4$; *etc.*

⁹For example, given the base case the next step in the induction has the second mover see her value and infer that $\Pr \{ \mathcal{S}_1; \mu_1^* | \mathcal{I}_2 \} = \Pr \{ \mathcal{S}_1 \} = \Pr \{ \mathcal{I}_1 \} \cdot F_\theta(\mu_1^*) = \frac{1}{4}$, given a one-half probability the first mover observes their value, and a one-half probability that their observed ball's value is lower than the first-round cutoff. The effective CDF for the rematching pool in period two is therefore $\frac{1}{4} \cdot F_\theta(x | \theta < 0) + \frac{3}{4} \cdot F_\theta(x)$, with expected value μ_2^* .

¹⁰The risk-neutral PBE cutoffs for rounds one/two/three, respectively, are: (0, -16, -23) for first movers; (-9, -20, -26) for second movers; and (-16, -23, -28) for third movers.

¹¹Cutoffs are flat going from the third mover to first mover in each round due to the conditioning in equation (1), where a first mover who sees their value in round two (the event \mathcal{I}_4 in the induction) knows that they did not switch in round one.

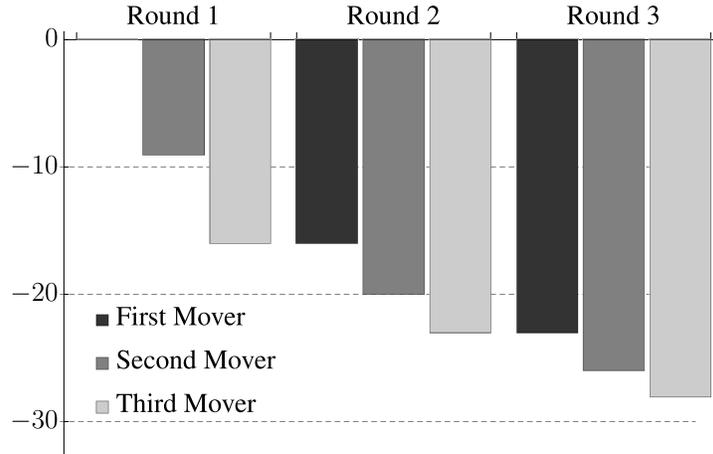


Figure 2. : Predicted adverse-selection accruing over supergame

ticipants use a stationary (non-boundary) cutoff rule that gives up low-valued objects and keeps high ones yields best-response cutoffs with quantitatively very similar predictions to the equilibrium.¹² The reason for this is that using a cutoff μ' that differs from the equilibrium cutoff has two largely offsetting effects in the first term on the right-hand-side of (II). On the one hand, increasing the cutoff increases the selection probability, $\Pr\{\mathcal{S}_t; \mu'\}$, as there are more values for which the initial object is exchanged. On the other, it decreases how bad that selection is, making the rematching distribution given a prior switch $\theta | \theta < \mu_t^*$ more favorable to succeeding agents. For our experimental parameterization, the robustness to others' cutoffs is a useful feature of the environment: even with substantial deviations from equilibrium play, the empirical best response remains essentially identical to the PBE. As such, accurate beliefs on the cutoff behavior of others is not essential for having a decreasing cutoff. Instead, the strategic sophistication required for a decreasing cutoff is more qualitative, understanding that information arrives slowly over time for all players, and that other players will give up low-valued objects and keep high-valued ones.

Our hypotheses over cutoffs in the S-treatment are therefore:

HYPOTHESIS 2 (S-Treatment): *Participants use strictly decreasing cutoffs over the rounds of the S-Treatment supergame.*

In addition to the qualitative direction of cutoffs within treatments, we can also predict no difference between treatment and control for the first-mover in the first round of the S-treatment, as at this point there is no selection.

¹²In Figure A1.1 in the online Appendix we indicate that the first-mover's cutoff decrease between rounds one and two is essentially the same across a large array of beliefs for the second- and third-movers round-one cutoff.

HYPOTHESIS 3 (First decision equivalence): *The distribution of first-round first-mover cutoffs in the S-treatment is identical to that in NS.*

Below we outline the aggregate experimental results from the treatment and control, where our focus will be on the late-session play after participants have acquired extensive experience within each environment. After outlining the main results (and checking their robustness) we turn to subject heterogeneity and how the response evolves over the session as the participants gain experience in section IV.A.

III. Aggregate Results

We now describe the main experimental results comparing the behavior in the environments with and without adverse selection, and examining the three hypotheses above.¹³ Aggregate results for the *S*- and *NS*-environments are illustrated in Figure 3. The figure presents normalized cutoff data from participants in the first-mover role in all supergames where a cutoff is elicited (6–21). The focus on first movers provides the cleanest comparison across treatments because (i) the PBE prediction is identical for the first-movers in the first round, and (ii) the changes in the optimal cutoffs across rounds are largest for first-movers.¹⁴ While the equilibrium theory (the circles) and empirical best response (the triangles) predict no adverse selection in *NS*, the prediction in the *S* treatment is for selection to accrue from no effect in round one, to a substantial negative cutoff in round three.

Three patterns emerge from inspecting Figure 3: (i) aggregate behavior does respond to the passage of time in *S* supergames, however, the adjustment to the adverse selection falls short of the equilibrium predictions (and the empirical best response); (ii) behavior is qualitatively different across treatment and control; and (iii) while aggregate behavior in the *NS* control is statistically indistinguishable from the risk-neutral prediction (a zero cutoff in every round), behavior in the *S* treatment is significantly different.

Table 1 provides random-effects regression results to complement the figure. The table reports normalized cutoff estimates for first movers across rounds one to three, recovered by regressing cutoffs on a set of mutually exclusive treatment-round dummies. We separately estimate first-mover behavior in supergames 11 to 20 (Panel A) and in the full-strategy method supergame 21 (Panel B).¹⁵ The estimated aggregate cutoff $\hat{\mu}_t^j$ for session type $j \in \{NS, S\}$ and supergame round

¹³Experimental data was obtained by the authors using an undergraduate student population. The data and replication code for the empirical parts of this paper can be accessed via the data repository. See the data citation Araujo, Wang and Wilson (2021) for full URL.

¹⁴Results and conclusions are statistically and numerically similar with a focus on all rounds and mover roles (see online Appendix A for details).

¹⁵Qualitatively similar results for supergames 6 to 20 are in online Appendix Table A3.1; while results for the second- and third-mover roles are in Table A3.2.

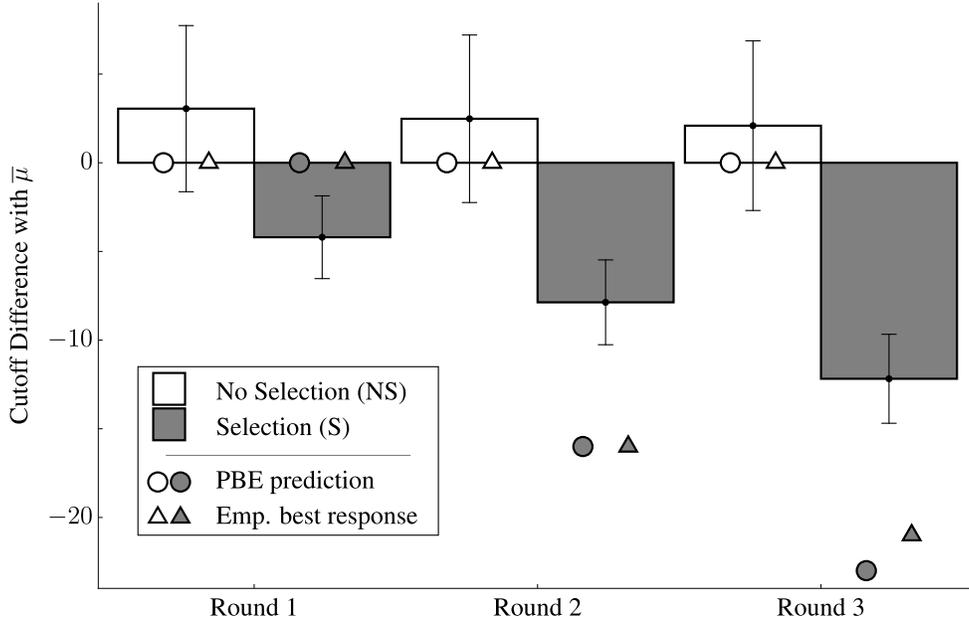


Figure 3. : First-mover cutoffs (supergames 6–21)

Note: Bars depict 95 percent confidence intervals from a random-effects estimation across all cutoffs in supergames 6–21. Empirical best responses calculated using cutoff distributions.

$t \in \{1, 2, 3\}$ then allows us to make statistical inference over the equilibrium hypotheses.

Hypothesis 1 predicts stationary cutoffs across the supergame in the *NS*-control. Inspecting the average cutoffs for the control in Table 1, we find that the average first-round cutoffs is +3.8, slightly higher than the risk neutral prediction of 0. While this decreases slightly over the course of each supergame, we cannot reject stationarity. Examining each *NS*-coefficient in turn we test whether the cutoffs used in each treatment-round are equal to the coefficients used in round one, reporting the p -values in the $H_0 : \hat{\mu}_t^j = \hat{\mu}_1^{NS}$ column. Individually, neither the second- nor third-round’s *NS* coefficients are significantly different from the first round. Examining Hypothesis 1 directly with a Wald test for the same cutoff for all three rounds of the control ($H_0 : \hat{\mu}_1^{NS} = \hat{\mu}_2^{NS} = \hat{\mu}_3^{NS}$) we fail to reject a stationary response for both supergames 11–20 ($p = 0.355$) and for supergame 21 ($p = 0.875$).

Beyond stationarity, we also fail to reject the stronger hypothesis that aggregate behavior in *NS* is both stationary and equal to the risk-neutral prediction. Examining each *NS*-coefficient separately, we fail to reject the risk-neutral predictions of a zero cutoff in all three rounds separately (p -values in the $H_0 : \hat{\mu}_t^j = \mu_t^*$ column) and jointly ($p = 0.238$ from a Wald test).

RESULT 1 (Control Stationarity): *In line with Hypothesis 1, we cannot reject*

Table 1—: Average first-mover cutoff by round in *No Selection* and *Selection**Panel A: Supergames 11 to 20*

Treatment	Cutoff	Theory	N	Estimate	p-values	
					$\hat{\mu} = \hat{\mu}_1^{NS}$	$\hat{\mu} = \mu_t^*$
NS	Rd. 1, $\hat{\mu}_1^{NS}$	[0]	330	+3.8 (2.4)	—	0.119
	Rd. 2, $\hat{\mu}_2^{NS}$	[0]	170	+3.1 (2.5)	0.213	0.216
	Rd. 3, $\hat{\mu}_3^{NS}$	[0]	83	+2.9 (2.5)	0.268	0.244
	Joint tests:			0.335[†]		0.238[§]
S	Rd. 1, $\hat{\mu}_1^S$	[0]	460	−4.4 (1.2)	0.003	0.000
	Rd. 2, $\hat{\mu}_2^S$	[-16]	212	−8.0 (1.3)	0.000	0.000
	R. 3, $\hat{\mu}_3^S$	[-23]	113	−11.9 (1.4)	0.000	0.000
	Joint Tests:			0.000[†]		0.000[§]

Panel B: Supergame 21

Treatment	Cutoff	Theory	N	Estimate	p-values	
					$\hat{\mu} = \hat{\mu}_1^{NS}$	$\hat{\mu} = \mu_t^*$
NS	Rd. 1, $\hat{\mu}_1^{NS}$	[0]	33	+1.8 (2.5)	—	0.479
	Rd. 2, $\hat{\mu}_2^{NS}$	[0]	33	+1.9 (2.5)	0.907	0.449
	Rd. 3, $\hat{\mu}_3^{NS}$	[0]	33	+2.3 (2.5)	0.620	0.360
	Joint tests:			0.875[†]		0.817[§]
S	Rd. 1, $\hat{\mu}_1^S$	[0]	33	−7.5 (3.0)	0.018	0.013
	Rd. 2, $\hat{\mu}_2^S$	[-16]	33	−11.1 (3.0)	0.001	0.107
	R. 3, $\hat{\mu}_3^S$	[-23]	33	−14.7 (3.0)	0.000	0.006
	Joint Tests:			0.000[†]		0.000[§]

Note: Figures derived from a single random-effects least-squares regression for the relative cutoff (choice-51) against treatment-round dummies. Standard errors in parentheses, risk-neutral predicted cutoffs in square brackets. There are 170/137/33 Total/*Selection*/*No Selection* first-mover participants across supergames 11–20, and 55/22/33 in supergame 21. *Selection* treatment includes data from *S* and *S-Peer* for supergames 11–20 as both treatments are identical up to supergame 21 (see section III for details); exclude participants in the second- and third-mover roles (these figures given in the online Appendix). †—Univariate significance tests columns examine differences from either the first-round coefficient from the control ($H_0 : \hat{\mu}_t^j = \hat{\mu}_1^{NS}$ for treatment j , round t) or the theoretical prediction ($H_0 : \hat{\mu}_t^j = \mu_t^{*j}$). ‡—Joint test of stationary cutoffs across the supergame ($H_0 : \hat{\mu}_1^j = \hat{\mu}_2^j = \hat{\mu}_3^j$ for treatment j); §—Joint test of PBE cutoffs in supergame ($H_0 : 0 = \hat{\mu}_1^j - \mu_1^{*j} = \hat{\mu}_2^j - \mu_2^{*j} = \hat{\mu}_3^j - \mu_3^{*j}$).

stationarity for average behavior in our NS sessions, nor can we reject the stronger risk-neutral prediction.

Given that aggregate behavior in our control is well-behaved and close to the risk-neutral prediction, we turn to aggregate behavior in the S-treatment where non-stationarity is predicted. The bottom half of Table 1 provides the average first-mover cutoffs across rounds 1–3 for the S-treatment, again breaking the estimates up into those obtained in supergames 11–20 and supergame 21. Our prediction for the S-treatment is that cutoffs are decreasing across the rounds (Hypothesis 2), but where behavior starts out at the NS level in round one (Hypothesis 3).

The data does indicate a decreasing response for the S-treatment. While stationarity of the cutoffs in the control can not be rejected jointly, we strongly reject stationarity in the S-treatment ($p < 0.001$ from a Wald test), where the coefficients indicate a strictly decreasing cutoff. However, although behavior is qualitatively in line with the prediction of a decreasing response, the aggregate cutoffs in S are far from the PBE prediction. As illustrated in Figure 3, participants' behavior does not fully internalize the predicted degree of adverse selection. For supergames 11–20, cutoffs in round three indicate approximately half the predicted effect (increasing to about 64 percent if we look at supergame 21 on its own). Moreover, the attenuation in the response relative to the theoretical prediction becomes more pronounced once we consider that first-mover participants in our experiment start out with lower average cutoffs in the very first round. While the behavioral decrease across the three rounds is significant ($\hat{\mu}_3^S - \hat{\mu}_1^S = -7.65$ in supergame 21, different from zero with $p < 0.001$), the magnitude of the observed cutoff change across the supergame is a third of the theoretical drop of -23.

Behavior in the first-round of the S-treatment jumps out as an anomaly. Despite an equivalent decision for first-round first-movers to NS, the provided first-round cutoffs in the S supergames are significantly lower than both the NS cutoffs ($p = 0.002$) and the risk-neutral prediction ($p = 0.004$). Moreover, this effect becomes more pronounced if we focus just on cutoffs in the last supergame.

RESULT 2 (Treatment Dynamics): *In line with Hypothesis 2, we reject that aggregate behavior in the S treatment is stationary, as the cutoffs have a significant and strictly decreasing trend. However, the dynamic reaction is significantly different from the theoretical prediction.*

RESULT 3 (First Round Non-Equivalence): *We reject Hypothesis 3, as average first-round cutoffs in the S treatment are significantly lower than both the NS control and the risk-neutral prediction.*

In Section IV.A we show that the aggregate patterns from the S-treatment (a negative but shallow slope over rounds, and a lower intercept) are the product of individual heterogeneity. Two behavioral types emerge: (a) Sophisticated participants using decreasing cutoffs across the supergame, starting close to the risk-neutral value in the first round; and (b) Boundedly rational participants using a

stationary response across the supergame, but where the level of their cutoff shift downward due to experienced (unconditional) selection. When averaged together, the two types produce the observed behavior, where the boundedly rational types lower the intercept and attenuate the dynamic response from the sophisticated types. Before analyzing individual heterogeneity in the next section, we briefly outline results from four additional treatments that checked the robustness of the *S*-treatment results.

Summary of Robustness Treatments

We run three treatments that manipulate the information participants receive on others' behavior and one additional treatment that generates dynamic adverse selection in a decision-making environment. Details of these treatments and the results are provided in the online Appendix for interested readers, as the main findings replicate the above. Our intention here is to provide a concise summary of each treatment conducted.

ROBUSTNESS TREATMENT 1 (S-Across): *Additional strategic feedback across supergames.*

Here we replicate the *S*-treatment, but where participants are given complete strategic feedback on other players' choices at the end of each supergame. Looking back to Figure 1(A) in the design, where the *S*-treatment only informed participants on their own choices (the elements in black), in *S-Across* participants are informed of all elements in the figure once the supergame has ended. Results mirror those in the *S*-treatment.

ROBUSTNESS TREATMENT 2 (S-Within): *Additional strategic feedback within the S-treatment supergame.*

This treatment modifies the information structure within the supergame so that participants are informed about others switches within the supergame. Rather than time, the relevant conditioning variable for selection is now the observation of a switch by another participant. In the Figure 1(A) example, the first mover would know that the second-mover had switched when they made their choice in round two. We again find qualitatively similar effects to the *S*-treatment at the aggregate level. Sophisticated participants respond to the appropriate signal (observed switches, not the passage of time), but the size of the response is attenuated.

ROBUSTNESS TREATMENT 3 (S-Peer): *Peer advice on strategy for the S-treatment.*

These sessions are identical to the *S*-treatment, except for supergame 21.¹⁶ Before the final supergame (which is paid with certainty) participants are matched

¹⁶Because the treatment is identical to *S* up to supergame 20, data from this treatment was included in Table 1 for the columns examining Supergames 11–20, but not for results examining Supergame 21.

into aligned-interest groups of three. After the chat, group members are matched with members from other chat groups for a final S -supergame, where the supergame-21 outcome from one of the three chat-group members is randomly selected for the entire group. As such, group-members are given a clear incentive to explain the environment to others.

Even though many groups do have chat members who explain the underlying tensions in the game to the other participants, the aggregate behavior in supergame 21 is not significantly different from that observed in the S -treatment environment.¹⁷

ROBUSTNESS TREATMENT 4 (S-Explicit): *This treatment contains a decision problem with adverse selection across time.*

In the NS control, a single agent makes choices over time, and, because the rematching pool is held constant, there is no adverse selection. In this modification, we provide the same rematching pool in the first round (an equal chance of each of the three unheld balls). In round two, however, the rematching pool has the highest-value ball removed, and becomes selected; and in round three, the second-highest rematching ball is also removed, so the only rematching ball is the worst of the three. This treatment exhibits similar effects to the S -treatment, with a non-stationary response that under-reacts to the adverse selection present.¹⁸

IV. Discussion

A. Participant Heterogeneity and Experience Effects

In Section III we provide evidence that average cutoffs are significantly decreasing across the S -supergames but are stationary in the NS control. To an extent this represent a victory for the theory as a *qualitative* prediction on aggregate behavior, where quantitative differences might be explained by other features of preferences. However, in this section, we show that the average behavior masks substantial heterogeneity in the participants' decision making. While a large minority of participants do use strictly decreasing cutoffs in the S -treatment, the majority use a stationary cutoff within the supergame. In this section, we

¹⁷Chat logs from all S -Peer sessions are included in online Appendix D for interested readers. Example explanations: "As the rounds go on, the chances that the ball the computer is holding has a really small value increases[...] because in previous rounds, if someone had a small value they probably switched and gave it to the computer"; "So here are my thoughts: The chance of you getting a low # that someone else switched out is based on which mover you are and what round it is. Typically I go with ~50 if I am mover 1 or 2 on the first round[...] Then drop down for each subsequent round. Because you get stuck with what you switch to and as time goes on that is much more likely to be a low #".

¹⁸The only difference from the S -treatment is that first-mover, first-round behavior in S -explicit is *not* different than first-round behavior in the NS control. This is likely due to the reduced effects of learning, as subjects in S -explicit encounter fewer instances of rematching to a low-value ball compared to S . We elaborate on this in another version of this paper, which is available upon request.

examine individual-level results to better understand the participant's within-supergame response, and the extent to which they are adjust their behavior across supergames as they gain experience.¹⁹

We define a simple type-schema based on participant's choices in the final supergame.²⁰ Specifically, we dichotomize participants as either *decreasing* or *non-decreasing*, where for the non-decreasing types we further break down the fraction that are *stationary*. A decreasing participant is one whose final-supergame cutoffs satisfy $\mu_1^i > \mu_2^i \geq \mu_3^i$, where a stationary participant satisfies $\mu_1^i = \mu_2^i = \mu_3^i$. In addition to these knife-edge types, we create a parallel family of definitions for $\epsilon > 0$, where an ϵ -decreasing type satisfies $\mu_1^i \geq \mu_2^i + \epsilon$ and $\mu_2^i \geq \mu_3^i$ (tightening the definition), and an ϵ -stationary type satisfies $|\mu_1^i - \mu_2^i|, |\mu_2^i - \mu_3^i| < \epsilon$ (weakening the definition).²¹

Table 2 provides the type compositions in the *NS*- and *S*-environments (including the robustness treatments). Focusing on the type definitions with $\epsilon = 2.5$, we find that all but one participant in the *NS*-control use non-decreasing cutoffs, with a large (slim) majority being ϵ (exactly) stationary. In contrast, pooling across our adverse-selection treatments where we expect decreasing behavior, we find that only a third of participants use decreasing cutoffs. As such, a comparison of the decreasing-type proportion between *NS* and *S* suggests that only one in every three participants correctly conditions on time in the selection environment. The majority of participants are instead better classified as using a stationary cutoff across the supergame.

Though our type dichotomy is based solely on behavior in the session's final supergame, the types are highly predictive of earlier responses. We look at two key measures in earlier supergames: (i) the participant's choice as a first-mover in round one, μ_1 , where no selection has occurred yet; and (ii) a within-supergame measure of the response to time, $\Delta\mu := \mu_{j,1} - \mu_{j,2}$ (regardless of the mover role j). For each variable we conduct random-effects regressions on a set of dummies interacting the participant's type (based on supergame 21 behavior), the game environment (*S* or *NS*), and indicators for each block of five supergames. The regression results are provided in Table 3, where the final column indicates the across-session difference in each coefficient (comparing supergames 16–20 and supergames 6–10).²²

For the first-round cutoffs, rematching leads to an identical lottery in both *S* and *NS*, where the risk-neutral prediction is a normalized zero cutoff. Our

¹⁹In a companion paper we examine the extent to which learning is predictable, using additional experimental data to test a model of mis-specified learning.

²⁰The final supergame represents the point where participants have maximal experience with the task and where we ramp up the monetary incentive by an order of magnitude as this supergame is paid for sure.

²¹Figure A.4.1 in the online Appendix provides the type proportions as we vary ϵ from 0 to 10, illustrating the robustness of across definitions.

²²For the *S*-treatment results we include participants from both *S* and *S-Across* as these supergames are theoretically identical. We do not include data from *S-Within* or *S-Explicit* as the theoretical environments here are distinct; nor do we include data from *S-Peer*, as the type classifications are made *after* the chat rounds, and so the type-definition is contaminated.

Table 2—: Type proportions

Treatment	$N_{Sbj.}$	Decreasing		Non-decreasing			
				Total		Stationary	
		Exact	$\epsilon = 2.5$	Exact	$\epsilon = 2.5$	Exact	$\epsilon = 2.5$
<i>NS-control</i>	33	3.0%	3.0%	97.0%	97.0%	57.6%	75.8%
<i>S-treatment</i>	66	42.4%	37.9%	57.6%	62.1%	36.4%	47.0%
Robustness:							
<i>S-Across</i>	60	28.3%	23.3%	71.7%	76.7%	45.0%	60.0%
<i>S-Peer</i>	72	34.7%	30.6%	65.3%	69.4%	47.2%	55.6%
<i>S-Explicit</i>	36	33.3%	33.3%	66.7%	66.7%	33.3%	47.2%
S+Robustness	234	35.0%	31.2%	65.0%	68.8%	41.4%	53.0%

Table 3—: Behavior by type across the session

Type	Supergames			Δ Session
	6–10	11–15	16–20	
<i>NS-Treatment (All)</i>				
First-round cutoff, μ_1	+1.8 (2.4)	+3.9 (2.4)	+3.7 (2.4)	+2.0 (0.7)
Cutoff change, $\Delta\mu$	+0.2 (0.8)	-1.3 (0.8)	-0.4 (0.8)	-0.5 (0.8)
<i>S-Treatment (Non-Decreasing)</i>				
First-round cutoff, μ_1	-1.7 (1.6)	-2.6 (1.6)	-4.1 (1.6)	-2.4 (0.8)
Cutoff change, $\Delta\mu$	-0.3 (0.5)	-0.5 (0.5)	-0.9 (0.5)	-0.6 (0.5)
<i>S-Treatment (Decreasing)</i>				
First-round cutoff, μ_1	-2.0 (2.3)	-3.8 (2.3)	-2.3 (2.3)	-0.3 (1.1)
Cutoff change, $\Delta\mu$	-6.6 (0.7)	-7.9 (0.7)	-8.6 (0.7)	-2.0 (0.7)

Note: Coefficients (and standard errors in parentheses) derived from two random-effects regressions from 1,125/1,233 observations of the first-mover first-round cutoff/cutoff-change within supergame over 159 participants.

results here indicate that: (i) Participants in the *NS* treatment use cutoffs that are consistent with the risk-neutral prediction across the entire session, though the averages are always slightly above zero, and where the increase across the session is significant ($p = 0.003$). (ii) In contrast for *S*, participants coded as non-

decreasing start out consistent with the risk-neutral response, but decrease over the course of the session until significantly negative ($p = 0.009$). (iii) First-round, first-mover cutoffs for participants coded as decreasing in S are not significantly different from the risk-neutral PBE prediction in any supergame block for the S -treatment, and this does not change substantially over the session.

The second measure provided in Table 3 examines how participants change their within-supergame cutoffs in reaction to time. We find that: (iv) Participants classified as non-decreasing in supergame 21 do not have a significant cutoff change between rounds one and two in prior supergames (p-values for blocks 6-10, 11-15, and 16-20 are, respectively, 0.544, 0.254, and 0.067).²³ This is true in environments both with and without adverse selection. (v) Finally, participants classified as decreasing in supergame 21 show a significant within-supergame cutoff reduction in reaction to time in earlier supergames, even in the 6–10 block ($p = 0.000$ for all three blocks). However, we do find evidence that the magnitude of their within-supergame reaction significantly increases across the session ($p = 0.004$). By the last supergame block the cutoff difference is -8.6, close to the PBE prediction of -11.3.²⁴

From Table 3 we conclude that the type classifications based on supergame 21 are useful for understanding participants' prior behavior. While this result may not be surprising for the non-decreasing participants, it does speak to the stability of their classification. For the decreasing types, the results indicate that these participants understand early on within the session that adverse selection accrues within the supergame, and that their cutoff should be decreasing with time.²⁵ Beyond the choice data, independent coding by two RAs of the chat logs from the S -Peer treatment suggest that approximately a quarter of the participants used the chats to *explain* the dynamic adverse selection mechanic within the game to others. This indicates that decreasing participants that do converge toward equilibrium behavior do so through an *understanding* of the strategic interaction. Though the sophisticated participants do adapt their expectations with respect to the size of the within-supergame dynamic selection effect (the significant negative reduction across-supergames for $\Delta\mu$), their appreciation of the need to condition on time emerges early on the session and remains present until the final supergame.

While the non-decreasing participants use an essentially static cutoff within the supergames, the across-supergame results do indicate a shift with experience. In particular, the non-decreasing participants become more pessimistic on the value of rematching as the S -treatment sessions proceed. This trend is absent in the

²³Non-decreasing participants in supergames 16-20 in the S -treatment and 11-15 in the NS control do have a marginally significant decrease in their cutoff. However, the change is quantitatively small (-0.9 and -1.3, respectively), and insignificant when we look at the joint hypothesis across the three blocks.

²⁴Since participants are equally likely to be assigned first, second or third-mover roles, the expected PBE prediction takes account of the frequency of each mover role, so: $1/3(-16 - 9 - 7) = -11.3$.

²⁵Looking just at decreasing types, 75–80 percent have a cutoff-difference in excess of -2.5 in each of the prior supergame blocks. Of the 29 participants with observed within-supergame changes in all three blocks, 21 are consistently negative.

NS-control, where if anything the *NS*-participants move in the opposite direction as the session proceeds.²⁶ While we do not detect any evidence that a substantial fraction of the participants moves between a stationary and decreasing response, we do find evidence that the stationary types reduce their cutoffs.

Focusing on the first-round first-mover cutoff, we are unable to reject the risk-neutral predictions in the *NS* treatment and for the decreasing participants in *S*. A similar failure to reject the risk-neutral prediction arises in the early supergames (6–10) for the non-decreasing participants in the *S*-treatment. However, by the end of the session, a significant gap arises, where the average non-decreasing participant uses a significantly negative cutoff of -4.1 ($p = 0.009$) (where the reduction across the session is also significant, $p = 0.002$).

Taken together the results in Table 3 provide an explanation for our previous findings with the two types pooled together. Aggregate behavior is a convex combination between two responses: (i) A minority of sophisticated subjects with a strongly decreasing, equilibrium-like response. (ii) A majority of subjects who do not respond to time, with a stationary response within the supergame; however, the level of their stationary response falls as they gain experience. Mixing the behaviors of the two types together, the sophisticated minority are large enough in number to exhibit the qualitative decreasing-over-time response in the pooled data. However, the level of the sophisticated types' decrease is heavily attenuated due to the larger mass of stationary participants. This explains Result 2, with a significant decreasing response to time, but with a quantitatively small effect.

Similarly, for Result 3, our aggregate-level data indicates a significant difference in first-round behavior between the *S*- and *NS*-treatments, despite decisions that are predicted to be theoretically identical decisions. This effect is driven by the stationary participants, where their longer-run response in reaction to experienced selection across the session is to lower their cutoffs, but maintain stationarity. This is in direct contrast with the learning effects in the *NS* treatment without selection, where if anything we see the opposite pattern, with subjects increasing their stationary cutoff across supergame as they gain experience. The aggregate-level effects for Result 3 are therefore driven by a combination of an adaptation of the response in reaction to experienced bad outcomes, but where the adaptation fails to identify the mechanic for the selection effect and so maintains a stationary response.

We summarize this finding as:

RESULT 4 (Heterogeneity): *Our aggregate level results are explained by a mixture of two types:*

- 1) *A minority of participants (approximately one-third) who understand the dynamic adverse selection within the environment, and use a decreasing*

²⁶Given the between-subject identification, we should clarify that about 30 percent of the participants classified as non-decreasing in *NS* would be expected to be decreasing types were they counterfactually placed in our *S* environment. However, we cannot separately identify sophisticated participants in the *NS* environment.

response that is close to the theoretical best response.

- 2) *A majority of participants (approximately two-thirds) are best described as using a stationary response, but where continued exposure to the adverse selection environment causes them to use a lower time-invariant cutoff, even for the first round where there is no adverse selection.*

V. Conclusion

We use a novel experiment to examine a common-value matching environment; one where adverse selection is dynamic, growing over time. Though substantially simplified, the core strategic tensions are similar to those present in labor, housing, and mating markets. The equilibrium prediction is that participants recognize the growing adverse selection, conditioning their responses on time, with a greater willingness to retain low-value object at later time periods.

A prior literature on failures of contingent thinking identifies sequentiality of the decision as a key predictor for understanding adverse selection, and suggests optimism for the equilibrium predictions in our sequential setting. However, while a substantial minority do respond to the adverse selection in a sophisticated way with a close-to-best response, the majority exhibit no change at all in their valuations over time: a stationary response to a non-stationary problem. While sequential, our environment does have substantial underlying uncertainty. Coupled with a complementary result in Martínez-Marquina, Niederle and Vespa (2019), which shows much greater strategic sophistication in a simultaneous setting with the uncertainty removed, the strategic sophistication shown in previous sequential studies seems to be driven by removal of uncertainty within the sequential setting.

While the modal stationary response places a cloud over the equilibrium predictions for our dynamic environment, we do see a silver-lining: Our sophisticated minority seem to understand the equilibrium introspectively. Their valuations are decreasing with time from the first supergames where we can observe this response. Most tellingly, in written advice to others they can explain the game's selection mechanic to others, speaking to a deeper understanding of the strategic features. When we think of professionals operating in dynamic markets—for example, finance, insurance, and labor markets—selection forces would seem to make the behavior of our sophisticated *minority* more representative.

That said, outside of professional settings where expert decision-makers are more likely to introspectively understand the strategic forces, our results point to a more ubiquitous misunderstanding of the selection effects. Moreover, while we do not observe any clear 'eureka' moments where participants move between a stationary and decreasing response, we do observe significant adaptation to experience. Future research focused on learning with misspecified models may prove very useful to understanding and predicting limit behavior, even if the conditions for convergence to the standard perfect Bayesian equilibrium are met.

REFERENCES

- Araujo, Felipe A., Stephanie W. Wang, and Alistar J. Wilson.** 2021. “Replication data for: The Times They Are a-Changing: Experimenting with Dynamic Adverse Selection.” *American Economic Association* [publisher], Inter-university Consortium for Political and Social Research [distributor]. <http://doi.org/10.3886/E119467V2>.
- Carrillo, Juan D, and Thomas R Palfrey.** 2009. “The compromise game: two-sided adverse selection in the laboratory.” *American Economic Journal: Microeconomics*, 1(1): 151–81.
- Chang, Briana.** 2018. “Adverse selection and liquidity distortion.” *Review of Economic Studies*, 85(1): 275–306.
- Chapman, Jonathan, Erik Snowberg, Stephanie Wang, and Colin Camerer.** 2018. “Loss Attitudes in the U.S. Population: Evidence from Dynamically Optimized Sequential Experimentation (DOSE).” National Bureau of Economic Research Working Paper 25072.
- Charness, Gary, and Dan Levin.** 2009. “The Origin of the Winner’s Curse: A Laboratory Study.” *American Economic Journal: Microeconomics*, 1(1): 207–36.
- Daley, Brendan, and Brett Green.** 2012. “Waiting for News in the Market for Lemons.” *Econometrica*, 80(4): 1433–1504.
- Enke, Benjamin.** 2020. “What you see is all there is.” *Quarterly Journal of Economics*, 135(3): 1363–1398.
- Esponda, Ignacio.** 2008. “Behavioral equilibrium in economies with adverse selection.” *American Economic Review*, 98(4): 1269–1291.
- Esponda, Ignacio, and Emanuel Vespa.** 2014. “Hypothetical thinking and information extraction in the laboratory.” *American Economic Journal: Microeconomics*, 6(4): 180–202.
- Esponda, Ignacio, and Emanuel Vespa.** 2018. “Endogenous sample selection: A laboratory study.” *Quantitative Economics*, 9(1): 183–216.
- Eyster, Erik, and Matthew Rabin.** 2005. “Cursed equilibrium.” *Econometrica*, 73(5): 1623–1672.
- Frederick, Shane.** 2005. “Cognitive reflection and decision making.” *Journal of Economic Perspectives*, 19(4): 25–42.
- Fuchs, William, and Andrzej Skrzypacz.** 2015. “Government Interventions in a Dynamic Market with Adverse Selection.” *Journal of Economic Theory*, 158: 371–406.

- Gershkov, Alex, and Motty Perry.** 2012. “Dynamic Contracts with Moral Hazard and Adverse Selection.” *Review of Economic Studies*, 79: 268–306.
- Guerrieri, Veronica, and Robert Shimer.** 2014. “Dynamic Adverse Selection: A Theory of Illiquidity, Fire Sales, and Flight to Quality.” *American Economic Review*, 104(7): 1875–1908.
- Hendel, Igal, Alessandro Lizzeri, and Marciano Siniscalchi.** 2005. “Efficient Sorting in a Dynamic Adverse-Selection Model.” *Review of Economic Studies*, 72(2): 467–497.
- Ivanov, Asen, Dan Levin, and Muriel Niederle.** 2010. “Can Relaxation of Beliefs Rationalize the Winner’s Curse?: An Experimental Study.” *Econometrica*, 78(4): 1435–1452.
- Jehiel, Philippe.** 2005. “Analogy-based expectation equilibrium.” *Journal of Economic Theory*, 123(2): 81–104.
- Jehiel, Philippe, and Dov Samet.** 2007. “Valuation equilibrium.” *Theoretical Economics*, 2(2): 163–185.
- Jehiel, Philippe, and Frédéric Koessler.** 2008. “Revisiting games of incomplete information with analogy-based expectations.” *Games & Economic Behavior*, 62(2): 533–557.
- Jin, Ginger Zhe, Michael Luca, and Daniel Martin.** forthcoming. “Is No News (Perceived as) Bad News? An Experimental Investigation of Information Disclosure.” *American Economic Journal: Microeconomics*.
- Kagel, John H.** 1995. “Cross-game learning: Experimental evidence from first-price and English common value auctions.” *Economics Letters*, 49(2): 163–170.
- Kagel, John H, and Dan Levin.** 2002. *Common Value Auctions and the Winner’s Curse*. Princeton, NJ:Princeton University Press.
- Levin, Dan, John H Kagel, and Jean-Francois Richard.** 1996. “Revenue effects and information processing in English common value auctions.” *American Economic Review*, 442–460.
- Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa.** 2019. “Failures in Contingent Reasoning: The Role of Uncertainty.” *American Economic Review*, 109(10): 3437–74.
- Ngangoué, Kathleen, and Georg Weizsäcker.** forthcoming. “Learning from unrealized versus realized prices.” *American Economic Journal: Microeconomics*.