An Experimental Analysis of Observational Learning with Payoff Externalities

David Owens

March 26, 2010

Abstract

Studies of observational learning typically assume that decision-makers consider past choices only for their informational content, which constitutes an information externality. This paper adds payoff externalities, by which observed decisions can directly affect the payoffs of decision-maker's available actions. In our experimental design, two subjects act sequentially and receive a private, informative signal drawn from a uniform distribution. The two signals jointly determine the state of the world, which subjects must guess. An over-reliance on private information in a baseline treatment shows that subjects' behavior is under-responsive to information externalities. Further treatments reveal that decisions are highly responsive to both positive and negative payoff externalities. To account for randomness in behavior, we employ a Quantal Response Equilibrium (QRE) analysis, which affirms subjects' sensitivity to payoff externalities.

1 Introduction

Prior to making a decision, economic agents often observe others making a similar decision. Consider a hungry patron choosing between several restaurants on the same block. She may have some private information as to which restaurant would suit her best. She may also look in the windows to see how many diners have already decided to patronize each establishment. The decisions of the other diners reveal to our patron, to some extent, the private information on which they were based. A full restaurant reveals that others' had positive information about the restaurant. In this way, early diners bestow a positive *information externality* on later diners. On average, later diners make better choices due to the decisions that they observe. This process, by which a decision-maker uses previous decisions to inform her own choice, is called *observational learning.*

The majority of the observational learning literature[1] assumes that payoffs are related across agents *only* through the information externalities described above. That is, observed decisions can increase a decision-maker's expected payoff only through the information that those decisions reveal. Clearly, such a model is not appropriate if a decision-maker has an explicit preference for taking an action similar to, or different from, those that she observes. Our diner, for example, may prefer to eat in lively, crowded establishments, or she may prefer relative solitude. Such a setting involves *payoff externalities*, or network effects, meaning that observed choices change the underlying value of alternatives. In the classical observational learning model, observed choices do not change the value of available choices, they simply reveal information about their value.

We explore a setting where information and payoff externalities exist simultaneously. Our research interests are twofold. First, our test of the classical observational learning process is more precise than that of existing studies. Second, we explore the effect of payoff externalities on this process, and the ability of agents to account for both influences in their decisions.

Our baseline treatment follows Çelen and Kariv (2004a) (CKa). Each player i receives a private signal s_i, drawn from a $U[-1,1]$ distribution. Each player's task is to take an action $d_i \in \{A, B\}$, which amounts to guessing the state of the world $\omega \in \{A, B\}$. $\omega = A$ if the sum of all private signals is positive, and $\omega = B$ if the sum of all private signals is negative. Therefore, higher values of s_i suggest that $\omega = A$, while lower values of s_i suggest that $\omega = B$. Correctly guessing the

[1] See section 2 for a literature review.

state of the world ($d_i = \omega$) yields a payoff of $Y > 0$, while an incorrect guess yields a payoff of 0. Players act in sequences of two, and player 2 observes d_1 before choosing her own action.

In this design, observational learning is possible for player 2. As ω is the same for both players, payoffs are perfectly correlated. Further, $d_1 = A$ makes A a more attractive action for player 2, and $d_1 = B$ makes B more attractive, due to what each reveals about s_1. As the action chosen by player 1 becomes more valuable, in expectation, to player 2, we say that this observational learning setting includes a *positive information externality*. Importantly, in this baseline treatment, the only incentive motivating players' actions is to correctly guess the state of the world. Therefore, one can intuit $d_i = A$ as player i's belief that $\omega = A$ is more likely than $\omega = B$, and the opposite for $d_i = B$.

Two additional treatments add payoff externalities to the baseline treatment described above. In these treatments, an amount X is added to the payoff of both players if both take the same action (i.e. $d_1 = d_2 = A$ or $d_1 = d_2 = B$). This payoff externality gives player 2 an added incentive to choose the same (different) action as player 1 if X is positive (negative). Thus, treatments in which $X \neq 0$ are no longer classical observational learning settings. In particular, it can not be assumed that a player's action is her best guess as to the state of the world. For marginal decisions, the d_2 that maximizes player 2's expected payoff may *not* correspond to the most likely state of the world if it is consistent with receiving $X > 0$ or avoiding $X < 0$. Across treatments, we keep the payoff for correctly guessing the state of the world, ω, constant at $Y = 2$. The payoff externality, X, varies across the three treatments, taking on values of -1, 0 and 1.

The moderate value of the X relative to Y across treatments is crucial. As a result, a rational player 2 always benefits by following extreme private signals. Specifically, player 2 should take action A for s_2 approaching 1, and B for s_2 approaching -1, for all employed values of X. Extreme values of X less than $-Y$ or greater than Y would make the game one of coordination rather than observational learning, in which player 2's private information would become irrelevant to her optimal action. We therefore include only moderate values of X to create an environment where both information and payoff externalities are important determinants of players' actions.

In each treatment, player 1 chooses $d_1 = A$ for positive s_1 and $d_1 = B$ for negative s_1 in equilibrium.[2] As discussed above, player 2 follows her private infor-

[2]Section 3 includes a formal analysis of player 1's decision problem.

mation for s_2 close to the endpoints of $[-1, 1]$ in equilibrium. Player 2's optimal action for moderate signals, however, depends on player 1's action. Due to the existence of positive information externalities, player 2 prefers to imitate player 1's action for moderate s_2. The *imitation set* is the set of private signals for which player 2 imitates player 1's action rather than follow her private signal. It is no surprise that, in equilibrium, imitation sets increase with the value of X. Predicted imitation sets are $[-.25, .25]$, $[-.5, .5]$ and $[-.75, .75]$ for $X = -1$, 0 and 1, respectively, corresponding to imitation rates of 62.5%, 75% and 87.5%.[3] It warrants repeating that only the only for $X = 0$ does this correspond to player 2's best guess for ω. In treatments with $X \neq 0$, d_2 differs from player 2's best guess for marginal s_2.

While other studies employ longer sequences of players, we chose sequences of only two. As such, our investigation of the observational learning process focuses on player 2's decision. A two-player sequence grants a clear theoretical benchmark, through Bayesian Nash Equilibrium (BNE), with which to compare our experimental results. Longer sequences of players have multiple equilibria and allow for complicated beliefs to rationalize a wide range of behavior. Payoff externalities force players to consider the actions of their followers as well as their predecessors.[4] This further clouds a decision process already complicated by the existence of information and payoff externalities. As the purpose of this paper is to assess the influence of the two types of externalities, sequences of only two subjects is most appropriate.

We test the predictions of our model in an experimental setting. Following CKa, subjects in our experimental sessions assign an action to each possible signal on the entire signal space of $[-1, 1]$ before viewing their private information. Specifically, subjects are asked to report a *cutoff strategy*. Before viewing s_i, each subject chooses $\hat{s}_i$, the lowest s_i for which they choose action $\omega = A$. Importantly, player 2 views d_1, but not $\hat{s}_1$, before choosing $\hat{s}_2$. This cutoff elicitation method affords a complete measure of subjects' beliefs, as they are forced to choose an action d_i for the entire $[-1, 1]$ continuum, not only the s_i that is drawn.

The first finding of our experiments is that player 2's guesses show a significant amount of observational learning. However, there is significantly *less* observational learning than predicted by BNE, as the observed imitation sets are smaller than predicted. Our second finding is that second-movers' cutoffs are *more* strongly

[3] These imitation rates classify all cases where $d_1 = d_2$ as imitation, even if d_2 is warranted by s_2.

[4] This tension exists, to a limited degree, for player 1. Section 3 explains that it does not effect player 1's equilibrium strategy for any of the three treatments.

influenced by payoff externalities than BNE predicts.

Combined, these results imply that negative payoff externalities can have significant adverse affects on the decisions of those who observe the choices of others. A diner with a mild aversion to crowded restaurants, for example, may eat more terrible meals than she should in pursuit of solitude. The results hold ambiguous predictions for positive payoff externalities. In our setting, positive externalities improve the quality of player 2's decisions, as subjects in the role of player 2 in the $X = 1$ treatment correctly guess ω with higher frequency than their counterparts in the $X = 0$ treatment. This may be surprising from a theoretical standpoint, as payoff externalities introduce an incentive that competes with the task of guessing the state of the world for marginal s_2. Our findings are consistent with subjects in the $X = 1$ treatment increasing their imitation set in response to the positive externality, and inadvertently benefiting from the information externality as well. One could easily imagine a different scenario where agents over-responsiveness to positive externalities causes them to imitate too much and make bad choices, in a manner similar to that predicted for negative externalities.

Our inclusion of a quantal response equilibrium (QRE) analysis serves two purposes. First, it is an equilibrium model which allows for randomness in behavior, an attractive feature for the analysis of experimental data. Second, it addresses the confound that player 2's expected payoff is much flatter as a function of stated cutoffs in the $X = -1$ treatment than in the other two treatments. QRE assumes that subjects respond *only* to payoff differences, and finds that subjects are least responsive to payoff differences when $X = -1$. In other words, subjects make not only the biggest mistakes under $X = -1$, but the *costliest* mistakes.

The remainder of the paper is organized as follows: Section 2 details existing studies of the three mentioned components of observational learning, while section 3 derives theoretical predictions for our model. Section 4 describes the experimental procedures in detail, and section 5 summarizes the experimental results. Section 6 presents the results of the QRE analysis, while section 7 discusses various implications of our findings and section 8 concludes.

2 Related Literature

Bikchandani, Hirschleifer and Welch (BHW, 1992) and Banerjee (1992) develop the classical observational learning model, a setting where externalities across agents are purely informational. Their theoretical framework involves agents acting in an exogenously-determined order, making once-in-a-lifetime decisions and observing all predecessors. General predictions include uniformity of behavior, even when private information differs across players. In addition, the observational learning process tends to retard information aggregation. Agents' incorporation of observed decisions into their own choices makes those choices less informative to followers[5]. Intuitively, as players ignore their private information in favor of information contained in decisions that they observe, they also fail to reveal their private information to those observing *them*.

Anderson and Holt (1997) use a laboratory setting to test the BHW model. They find that subjects are generally responsive to information contained in observed actions, including in circumstances in which they must act contrary to their private information. Several other experimental studies have reaffirmed subjects' willingness to learn from observed actions[6].

Hung and Plott (2001) conducted the first experimental study of the observational learning process when accompanied by externalities. They find that positive externalities increase subjects' tendencies to imitate actions that they observe. Recent work by Drehmann, Oechssler and Roider (2007) uses an internet design to explore an observational learning setting with externalities [7]. They find evidence that subjects respond to externalities, and that they are myopic, meaning that they tend to ignore the effect of their followers on their actions.

The studies mentioned above find that subjects tend to replicate the actions of those that they observe when it is profitable to do so. All use a binary-information, binary-action design, which affords only a course measure of subjects' tendency to learn from observed actions. In Anderson and Holt's (1997) setup, subjects receive one of two possible signals choose between actions. In their design, the payoff ratio between the two actions can take on three values: $\frac{1}{2}$, 1 and 2. Therefore, when subjects benefit from ignoring their private information, the payoff to doing so is twice as high as the payoff to not doing so. The project therefore yields a binary measure of whether subjects learn from observation when the benefits thereto are very high. Other studies incorporate similar binary measures of learning.

Our continuous decision space allows us to discern a subjects' choice between actions for the entire continuum of private signals. As each private signal corre-

[5]This result requires assumptions on the information and action space that is likely very realistic in real world observational learning settings. A sufficient condition is that the action space is courser than the signal space

[6]Hung and Plott (2001) and Drehman Oechssler and Roider (2007) are examples.

[7]They employ positive externalities, negative externalities and a unique setting where subjects receive a negative externality for those that they follow and a positive externality for those that follow them

sponds to a different payoff for each action, we are able to extrapolate choices for a continuum of expected payoff ratios. This allows us to specify, for each action, the payoff premium necessary to entice subjects to ignore their private information. We consider this a more complete measure of observational learning than has been done to date.

3 Theory

3.1 The Model

In our model, two players act in sequence, with the second observing the action chosen by the first. Each player i receives a private, informative signal s_i. It is private in the sense that it is not observed by the other player, and informative in that it improves the accuracy of beliefs about state of the world. s_i is independently drawn from a uniform distribution on the interval $[-1, 1]$ [8].

The task of each player i is to choose an action $d_i \in \{A, B\}$. s_1 and s_2 determine the state of the world, $\omega \in \{A, B\}$. $\omega = A$ if $s_1 + s_2 \geq 0$, and $\omega = B$ if $s_1 + s_2 < 0$. Player i receives a reward of Y if $d_i = \omega$ and no reward if $d_i \neq \omega$. Thus, each player's task is to correctly guess the state of the world.

In addition to the reward described above, each player receives a payoff externality in the amount of X *iff* $d_1 = d_2$. This paper explores both positive and negative values of X. X adds a layer of complication to the observational learning problem. Subjects are impelled to choose the correct state of the world, but have an additional motivation to choose the same (different) action as their partner if $X > 0$ ($X < 0$).

We restrict our attention to moderate values of the payoff externality, such that $-Y < X < Y$. Such a constraint of the values of X is critical in preserving the project as one of observational learning. Moderate values ensure that observed actions influence decisions through the private information that they reveal *and* the payoff externality that their imitation will yield. As a result, players benefit from following private signals close to -1 and 1, and player 2 benefits from following d_1 for moderate s_2. Thus, the problem remains one of observational learning and does not degenerate into a coordination game.

[8] Appendix A.1 explores the optimal strategies for a generic distribution of s_i.

3.2 Decision Problems

We now consider the decision problem faced by each player. We first derive Bayesian Nash Equilibrium predictions for the classical observational learning case, where externalities are purely informational. Then, we consider more general case with payoff externalities.

Subjects in the experiments received signals uniformly distributed on $[-10, 10]$. To simplify the analysis in this section, we normalize the signal space to $[-1, 1]$, and utilize the value of the reward, Y that was used across treatments. *For the remainder of the analysis,* $Y = 2$.

3.3 Classic Observational Learning Setting: $X = 0$

When payoff externalities accompany the observational learning process, player 1 must consider his actions' influence on player 2's behavior. As player 1's payoff changes by X if $d_1 = d_2$, player 1 may have a motivation to employ strategies that increase or decrease the probability that player 2 takes the same action.

In the classic observational learning setting, where $X = 0$, player 1 does not need to consider the effect of her actions on player 2's actions. $\pi_{1k}(s_1)$ is player 1's expected payoff from choosing action k as a function of s_1. Equations 3 and 4 show the equations for these expected payoffs, which is simply $Y \times Prob(s_1 + s_1 \geq 0|s_1)$ and $Y \times Prob(s_1 + s_1 < 0|s_1)$ for actions A and B, respectively.

$$\pi_{1A}(s_1) = 2 \times \frac{s_1 + 1}{2} = s_1 + 1 \tag{1}$$

$$\pi_{1B}(s_1) = 2 \times \frac{1 - s_1}{2} = 1 - s_1 \tag{2}$$

Figure 1 plots these expected payoffs against s_1. The solid line shows player 1's expected payoff for the choice of action A, and the dotted line the expected payoff of action B. Very low signals make action B more profitable, and high signals make action A more profitable. Player 1 receives a higher payoff choosing B for $s_1 < 0$ and A for $s_1 \geq 0$.

We now define a *cutoff strategy.* A cutoff strategy is one in which player i takes one action for all s_i below some cutoff level $\hat{s}_i$, and the other for all s_i greater than or equal to $\hat{s}_1$. Above, we show that player 1 employs a cutoff strategy of $\hat{s}_1 = 0$ in equilibrium, choosing action B for negative s_1 and A for positive s_1. Let $\hat{s}_1^*$ denote player 1's equilibrium cutoff strategy.

Upon viewing d_1, player 2 can update her expectations. Player 2 knows that

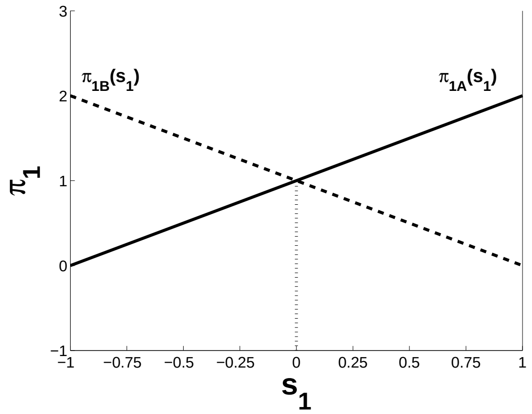

Figure 1: $\pi_{1A}(s_1)$ and $\pi_{1B}(s_1)$

$s_1 \geq 0$ if $d_1 = A$ and $s_1 < 0$ if $d_1 = B$. We define π_{2kl} as player 2's expected payoff for choosing action k after observing player 1 choose action l. Player 2's updated payoffs are illustrated in figure 2. As player 2 may find herself at two different information sets, following actions of A and B by player 1, these two decisions must be analyzed separately. Figure 2(a) shows the payoffs updated for $d_1 = A$, and figure 2(b) for $d_1 = B$. Again, the solid line shows the expected payoff associated with the choice of action A, and the dotted line with the choice of action B.

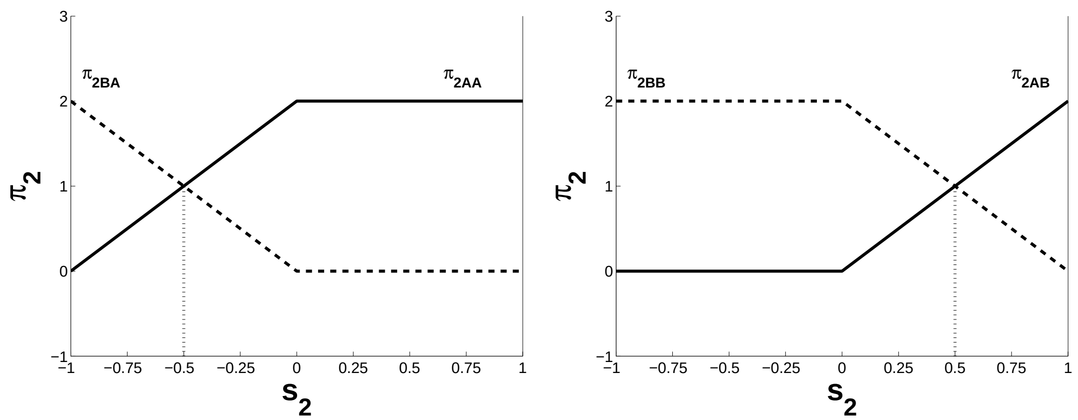

(a) $\pi_2(s_2)$

(b) $\pi_2(s_2)$

Figure 2: π_2 for $X = 0$

First, notice the flatness of player 2's payoff functions for negative values of s_2 in figure 2(b) and for positive values of s_2 in figure 2(b). This occurs because such signals, coupled with d_1 perfectly reveal the state of the world[9]. Further, player 2 employs a cutoff strategy at both information sets. Let s_2^A be player 2's cutoff following $d_1 = A$, and s_2^B following $d_1 = B$. Further, s_2^{A*} and s_2^{B*} denote player 2's optimal cutoffs. For $X = 0$, figure 2 shows that $s_2^{A*} = -.5$ and $s_2^{B*} = .5$

This discussion of the $X = 0$ case reveals that player 1, in theory, suppresses no information. Her action is as informative as it could be, in that it minimizes the variance of player 2's expectation of s_1 given d_1. In the process of maximizing her payoff, player 2 suppresses some of her private information. An omniscient observer learns less after observing d_1 and d_2 than she would if player 2 set $\hat{s}_2 = 0$. For a large subset of her signals $[-.5, .5]$, player 2 simply mimics the action of player 1. This is the *imitation set*: the set of s_2 for which player 2 takes the same action as player 1 unconditionally.

3.4 The Addition of Payoff Externalities: $X \neq 0$

The addition of externalities complicates the analysis. We begin by arguing that both players still employ cutoff strategies in equilibrium. Equations 3 and 4 show player 1's profits for choosing A and B for *any* strategy employed by player 2 (in particular, we do not assume that player 2 uses a cutoff strategy). Pr_{jk} is the probability that, given player 2's strategy, she chooses action j after observing $d_1 = k$.

$$\pi_{1A} = s_1 + 1 + Pr_{AA} \times X \tag{3}$$

$$\pi_{1B} = 1 - s_1 + Pr_{BB} \times X \tag{4}$$

As π_{1A} is increasing in s_1 and π_{1B} is decreasing, there will be some $\hat{s}_1$ such that $\pi_{1A} \geq \pi_{1B}\ \forall\ s_1 \leq \hat{s}_1$ and $\pi_{1A} \leq \pi_{1B} \forall s_1 \geq \hat{s}_1$. In other words, *in any BNE, player 1 employs a cutoff strategy.*

Proposition 1. *In any BNE, player 1 employs a cutoff strategy $\hat{s}_1 = 0$.*

Proof. See appendix A.2 □

The complete proof has been relegated to appendix A.2, which we summarize here. For any X, $\hat{s}_1 = 0$ maximizes the probability that player 1 receives the reward, $Y = 2$. The probability that player 1 receives the payoff externality X, however, varies with the value of X, as player 2 is more likely to imitate d_1 for

[9]This is a departure from the BHW model, where conditional on the state of the world, private signals are independent. In our model, $\omega = A$ makes high s_2 more likely for a given s_1, and the converse foe $\omega = B$.

higher values of X. Player 1 must therefore consider player 2's reaction to her choices in choosing $\hat{s}_1$. Appendix A.2 derives player 2's reaction function to d_1, the result of which is the combination of equations 5 and 6.

$$s_2^{A*} = -\frac{1}{4} \times \big[(2+X) + \hat{s}_1 \times (X-2)\big] \quad (5)$$

$$s_2^{B*} = \frac{1}{4} \times \big[(2+X) + \hat{s}_1 \times (X-2)\big] \quad (6)$$

Taking into account player 2's response to $\hat{s}_1$ as presented in equations 5 and 6, equation 7 shows the derivative of player 1's profit as a function of $\hat{s}_1^2$.

$$\frac{\partial \pi_1(\hat{s}_1)}{\partial \hat{s}_1^2} = -\frac{1}{2} + \frac{X}{16} \times \big[X-2\big] \quad (7)$$

For $X > -Y$, this number is always negative, proving that $\hat{s}_1^* = 0$ across treatments. Intuitively, the $-\frac{1}{2}$ on the left-hand side of equation 7 reflects the fact that setting an extreme cutoff of -1 or 1 reduces player 1's probability of receiving $Y = 2$ from $\frac{3}{4}$ to $\frac{1}{2}$, relative to a cutoff of 0. The portion on the right, $\frac{X}{16} \times \big[X-2\big]$ reflects the expected change in payoff resulting from the altered probability of receiving the externality. This number is negative for $0 < X < 2$, reflecting that the decreased probability of receiving the payoff externality is detrimental to player 1 when $X > 0$. The number is positive, but less than $\frac{1}{2}$ when $-2 < X < 0$. This reflects the fact that under negative payoff externalities, the decreased probability of imitation helps player 1, but not enough to justify the decreased probability of receiving $Y = 2$. Thus, $\hat{s}_1^* = 0$.

The fact that $\hat{s}_1^* = 0$ makes the analysis straightforward for player 2 for the relevant values of X. Our experiment involves three different values of X: 1, 0 and -1. Player 2's expectations of s_1 given d_1 are unchanged from the $X = 0$ case. Thus, π_{2AB} and π_{2BA} are unchanged from the $X = 0$ case illustrated in figures 2(a) and 2(b), respectively. π_{2AA} and π_{2BB} are both simply shifted vertically by X, as the payoff to choosing the same action as player 1 changes by X from the $X = 0$ case. Figure 3 represents player 2's payoffs for $X = 1$ and $X = 0$. The graph shows that $s_2^A = -.75(-.25)$ and $s_2^B = .75(.25)$ for $X = 1(-1)$, respectively. Analytically, substituting $\hat{s}_1^* = 0$ into equations 5 and 6 yields the following equations for $\hat{s}_2^A$ and $\hat{s}_2^B$, across treatments:

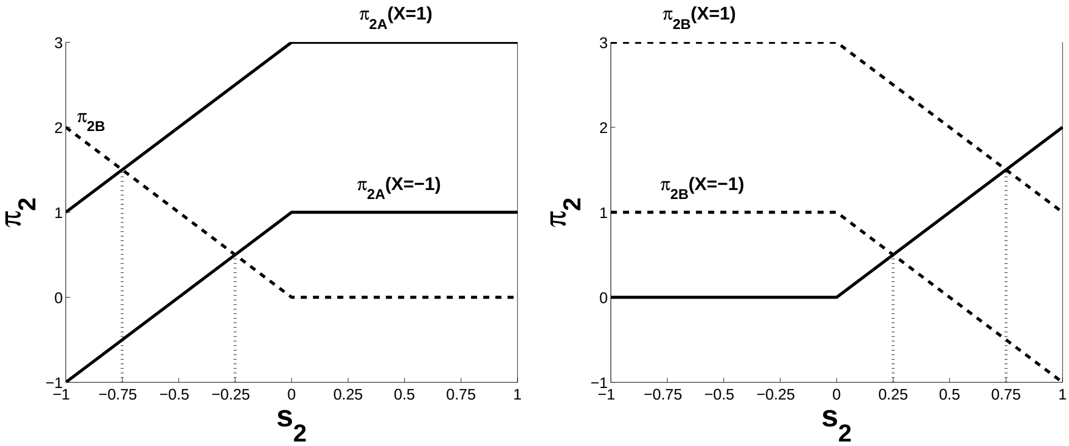

(a) π_2 for $d_1 = A$: $X \neq 0$

(b) π_2 for $d_1 = B$: $X \neq 0$

Figure 3: π_2 for $X \neq 0$

$$s_2^{A*} = -.5 - .25 \times X \tag{8}$$
$$s_2^{B*} = .5 + .25 \times X \tag{9}$$

	Treatment		
	X=1	**X=0**	**X=-1**
$\hat{\mathbf{s}}_1^*$	0	0	0
$\hat{\mathbf{s}}_2^{\mathbf{A}*}$	-.75	-.5	-.25
$\hat{\mathbf{s}}_2^{\mathbf{B}*}$	.75	.5	.25
$Prob(d_1 = d_2)$	.875	.75	.625

Table 1: Predicted Cutoffs

Note that the game is symmetric, in that $\hat{s}_1^* = 0$ and $s_2^{A*} = -s_2^{B*}$. Further, the cutoffs vary in sensible directions. Positive externalities increase the size of player 2's information set. Importantly, unconditional imitation is not predicted. Player 2 benefits from imitating player 1's action for many private signals, but should always follow her private information in the case of extreme s_2.

According to BNE predictions, the information aggregation process is impeded by positive externalities in our two-subject design. An observer gains a noisier measure of ω from observing d_1 and d_2 as X increases. The intuition is clear. Positive externalities allow for more extreme s_2 to be overridden by the observation of d_1. Under negative externalities, marginal s_2 are no longer overturned, allowing player 2 to rely more on, and therefore reveal more completely, her private information.

4 Experimental Design

The experiment took place at the University of California at Berkeley's Experimental Social Sciences Laboratory (XLab). A total of 161 subjects were used, comprised of staff and students of UC Berkeley. Each session involved between 10 and 20 subjects, and subjects were not allowed to participate in more than one session. Each session involved 1 treatment only. A total of 30, 30 and 36 subjects participated in the conditions of $X = 1$, 0 and -1, respectively. Separate sessions were conducted with a computer playing the roles of player 1 in each treatment. 17, 25 and 23 subjects participated in these sessions. Subjects made all decisions on computers, and were divided by cubicles to limit interaction and ensure anonymity of decisions and private information. Each session lasted between sixty and ninety minutes. Before participating, subjects read a detailed set of instructions which was then read out loud to them by an instructor[10]. Included in the instructions was a description of the payoffs specific to each treatment. Subjects were paid \$5 as a show-up fee, and their subsequent earnings depended on their performance, according to the payoff schedule described in section 3. The average earnings totaled roughly \$21. During the experiment, earnings were described in terms of "experimental tokens", and their exchange rate into dollars (4 tokens = 1 dollar).

To avoid granting disproportionate salience to the payoff externality, we presented payoffs in as neutral a manner as possible. To this end, subjects were not told that they received any additional cost or benefit for following the action they observe. Rather, payoffs were presented as illustrated in table 2[11], for the $X = 1$ treatment when $s_1 + s_2 < 0$.

Sessions consisted of 30 rounds each. At the beginning of each round, subjects were randomly divided into anonymous groups of two and randomly assigned to the role of player 1 or player 2. The computer then drew signals s_1 and s_2 from a $U[-10, 10]$ distribution, which players 1 and 2, respectively, used to make their decision. Each subjects' task was to choose $d_i \in \{A, B\}$. Before making her deci-

[10]Copies of the experimental instructions are available at http://rogilla30.googlepages.com/experimentalinstructions.

[11]There was an analagous table for $s_1 + s_2 > 0$.

If $\mathbf{s_1 + s_2 < 0}$:

		Others' Choice	
		A	**B**
Your Choice	**A**	1	0
	B	2	3

Table 2: Payoff Presentation: $X = 1$ for $\omega = B$

sion, player 2 was informed by the computer of player 1's decision.

Rather than observing their signal and then making their decision, subjects were asked to make a decision rule before *viewing their signal.* Specifically, each subject was asked to choose a cutoff $\hat{s}_i$, such that $d_i = A$ if $s_i \geq \hat{s}_i$ and $d_i = B$ if $s_i < \hat{s}_i$[12]. This cutoff elicitation method was employed to extract a richer measure of learning than a binary action space can provide.

After both subjects made their choices, they were informed of the value of each signal, and the payoff that they received. This process was repeated for each of the thirty rounds in each session, and subjects were paid anonymously upon their exit.

5 Results

The application of our experimental results to the research questions will rely on two types of summary statistics: observed choice outcomes, $d_i \in \{A, B\}$, and observed cutoffs, $\hat{s}_i \in [-1, 1]$. We introduce the following definitions for use in the discussion of d_i: *Imitation* occurs when $d_2 = d_1$. *Divergence* occurs when $d_2 \neq d_1$.

Before employing the observed $\hat{s}_i$ in our analysis, it is instructive to discuss its interpretation. Based on private information, subjects should choose action A for $s_i < 0$ and B for $s_i > 0$, or equivalently set $\hat{s}_i = 0$. $\hat{s}_i < 0$ means that player i chooses $d_i = A$ for some $s_i < 0$, which shows some inclination for action A. Conversely, $\hat{s}_i > 0$ can be interpreted as an inclination for action B. Following Smith and Sorensen (2000), we use the term *cascade* to describe an inclination sufficiently extreme to choose the same action for any private signal. $\hat{s}_i = -1$ signifies a players' decision to choose action A for any private signal, while $\hat{s}_i = 1$ means a subject chooses action B for any private signal. Therefore, cutoffs of -1 and 1 in our design are cascades.

For player 2, cascades can take on two forms. Player 2 can choose to imitate

[12] Player 2 viewed d_1 before choosing $\hat{s}_2$.

d_1 for any s_2, or to diverge for any s_2. $s_2^A = -1$ or $s_2^B = 1$, which represent unconditional imitation, shall be termed *positive cascades.* $s_2^A = 1$ or $s_2^B = -1$, which represent unconditional divergence, will be referred to as *negative cascades.*

5.1 Observational Learning

Our first research interest involves the efficiency of the observational learning process. As mentioned in section 2, previous experimental work has largely supported success of BHW's model. Our $X = 0$ treatment is a richer test of the observational learning model, in that we elicit choices for a continuum of expected payoff differentials. We begin by exploring whether player 2 does, in fact, learn from player 1's action. Recall that in the $X = 0$ treatment, subjects have no incentive other than to choose ω. Table 3 shows the rate at which subjects did so in the experiment. Subjects performed better in the role of player 2, as evidenced by the

	$\mathbf{X = 0}$ (Observations)	
	Player 1 (450)	**Player 2** (450)
Success	.691 (.750)	.756 (.875)

Table 3: Choosing ω when $X = 0$

6.5% increase in the success rate. In this way, we can say that player 2 did learn from the action of player 1. Player 2 performed worse relative to theory, however, falling short of the prediction by 11.9% versus 5.9% for player 1.

In quantifying player 2's learning, it is useful to look at the observed $\hat{s}_2$. Table 4 shows the mean of s_2^A and s_2^B for $X = 0$. Two regularities are evident. First, we find further evidence that player learns from d_1. s_2^A and s_2^B are both different from $\hat{s}_1$, and in the direction predicted by theory, $s_2^B < \hat{s}_1 < s_2^A$. This means that player 2 shows a preference for the action chosen by player 1. Second, player 2 learns less than predicted by theory, in that $s_2^A > s_2^{A*}$ and $s_2^B < s_2^{B*}$. In other words, player 2 shows a weaker preference for d_1 than predicted. Further, $\hat{s}_1 > 0$ and $\hat{s}_2^A > -\hat{s}_2^B$, which are consistent with positive-skewed cutoffs, as observed for player 1 in figure 4.

	X=1		
	$\mathbf{\hat{s}_1}$	$\mathbf{s_2^A}$	$\mathbf{s_2^B}$
Observed Cutoff	.136 (.399)	-.188 (.511)	.324 (.434)
Predicted Cutoff	0	-.5	.5

Table 4: Observed Cutoffs for $X = 0$

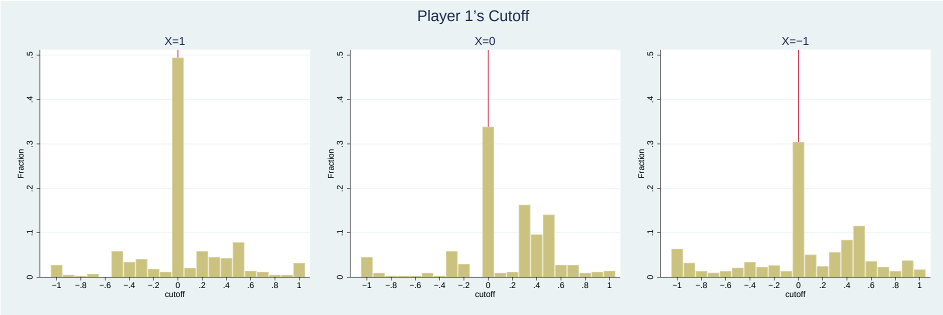

Figure 4: $\hat{s}_1$

Table 4 presents evidence that player learns less from d_1 than predicted by BNE. This prediction, however, assumes that $\hat{s}_1 = 0$, which figure 4 shows is often not the case. It is therefore possible that subjects in the role of player 2 are less influenced by d_1 because they correctly believe that player 1's decision-making is unreliable, making d_1 less informative.[13] In order to address this confound, the sessions were conducted with a computer in the role of player 1. All subjects therefore played the role of player 2 for all 30 rounds, and were told that the computer always chose $\hat{s}_1 = 0$. $\hat{s}_2$ is presented in figure 5 for sessions involving humans versus computers in the role of player 1. 5 shows that there is no drastic difference when a computer plays the role of player 1. If anything, it seems that player 2 learns less than when playing with a human player 1. As a result, we conclude that the under-learning observed in $\hat{s}_2$ is not caused by a mistrust of player 1's rationality. We interpret the under-learning as evidence that observational learning is not an easy process, and even subjects in a straightforward experimental environment have difficulty with it.

5.2 Externalities and the Observational Learning Process

In the study of observational learning, uniformity of behavior is an important metric to consider. The prevelance of imitation is an appropriate measure for such uniformity. Table 5 shows the frequency of imitation in both treatments with externalities, and compares them to their theoretical predictions. Predictably, imitation is higher under the positive externality. The difference is larger than predicted by BNE, with imitation nearly twice as frequent under $X = 1$ as under $X = -1$.

[13] Appendix C shows that d_1 is very informative to player 2, even if $\hat{s}_1$ is very noisy. This is largely due to our restriction that players employ cutoff strategies. Still, mistrust in d_1 could change player 2's behavior.

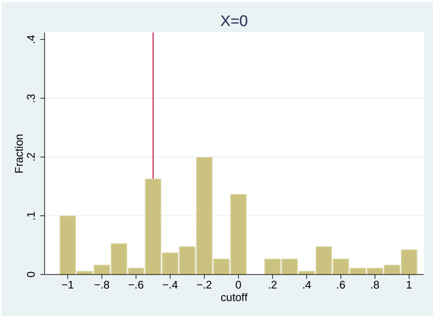

(a) s_2^A: Human Player 1

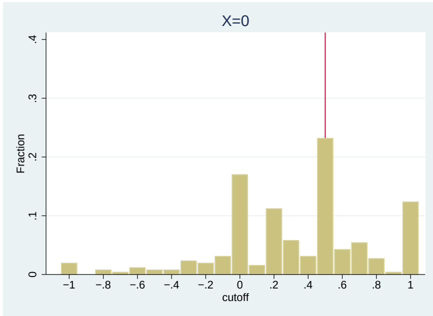

(b) s_2^B: Human Player 1

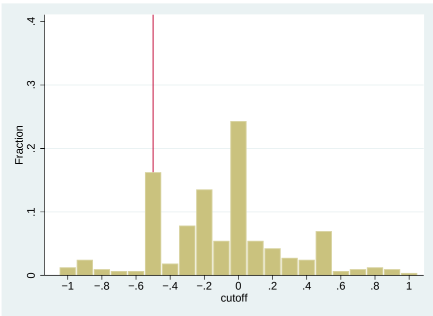

(c) s_2^A: Computer Player 1

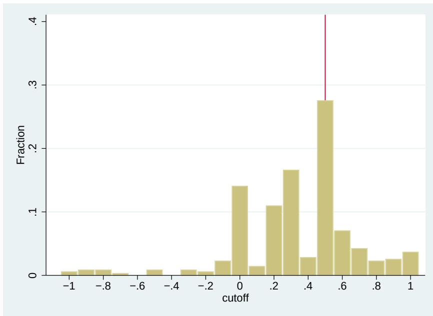

(d) s_2^B: Computer Player 1

Figure 5: $\hat{s}_2$ for human vs. computer player 1: $X = 0$

	Treatment (Observations)	
	X=1 (450)	**X=-1** (540)
Observed Imitation	.798	.430
Predicted Imitation	.875	.625

Table 5: Observed and Predicted Imitation with Externalities

As outlined in table 5, imitation is predicted for 87.5% of s_2 under $X = 1$, and only under 62.5% of s_2 under $X = -1$. It is therefore necessary to bear in mind that, as we compare (d_1, d_2) outcomes to their BNE predictions, that we are comparing the outcomes to different subsets of realized (s_1, s_2). Predicted outcomes are assigned to regions in the (s_1, s_2) plane in figure 6. Table 6 categorizes the observed (d_1, d_2) outcomes for the predictions illustrated in figure 6. The three categories are successful BNE predictions, player 1 error $(d_1 \neq d_1^*)$ and player 2 error $(d_1 = d^*)$, $d_2 \neq d_2^*$. Two important regularities are manifest in the table. First, BNE is more successful in predicting imitation under $X = 1$, and at

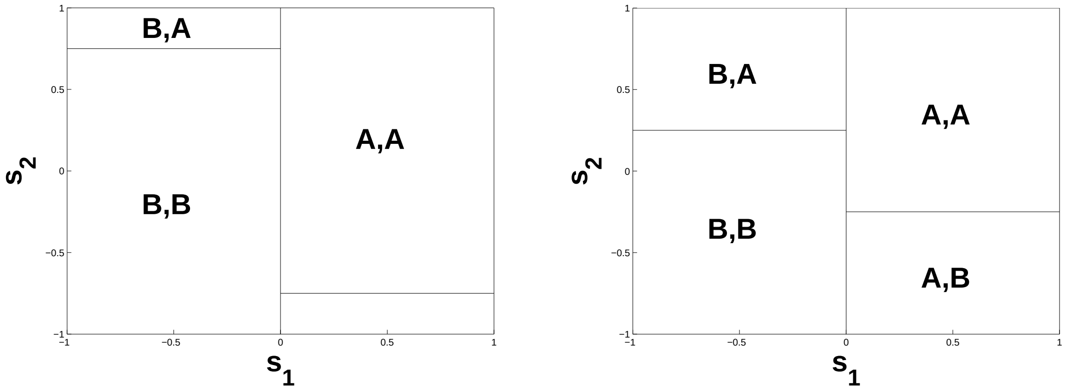

(a) $X = 1$ (b) $X = -1$

Figure 6: Predicted Choices with Externalities

predicting divergence under $X = -1$. Put another way, subjects are more likely to imitate correctly under positive payoff externalities, and to diverge correctly under negative payoff externalites.

The second regularity involves a comparison of player 2's error rate across treatments. Player 2 is far more likely to imitate incorrectly under $X = 1$, and to diverge incorrectly under $X = -1$. In other words, the most common mistake under negative payoff externalities is for player 2 to follow her private information when theory predicts imitation. The most common mistake under positive payoff externalities, on the other hand, is for player 2 to imitate player 1's action, when she should allow her strong private signal to override it.

Each of the two regularities suggests that player 2 tends to imitate under positive payoff externalities and diverge under negative. An interpretation is that subjects are more sensitive to payoff externalities than predicted by BNE. They tend to make 'mistakes' that are consistent with avoiding the negative externality and seeking the positive.

	Imitation Predicted		**Divergence Predicted**	
	X=1	**X=-1**	**X=1**	**X=-1**
Observations	**400**	**334**	**50**	**206**
BNE Success	.730	.480	.580	.680
Player 2 Error	.133	.341	.320	.112
Player 1 Error	.138	.180	.100	.209

Table 6: Error Break-Down by Prediction

Figures 7 and 8 show the observed distributions of s_2^A and s_2^B, respectively, for $X = 1$ and $X = -1$. For the reasons discussed above, results are also shown for the sessions with a computer in the role of player 1, with the knowledge that $\hat{s}_1 = 0$. $\hat{s}_2$ toward the left of figure 7 and to the right of figure C.3 correspond to a strong preference to imitate d_1, while the opposite imply a tendency to diverge. As predicted by BNE, $\hat{s}_2$ shows a stronger preference for imitation under $X = 1$. Further, the majority of $\hat{s}_2$ fall in the vicinity of the BNE predictions of $s_2^{A*} = -.75$ and $s_2^{B*} = .75$. Much of the observed $\hat{s}_2$ under $X = 1$ are in fact classified as positive cascades.

Behavior is less consistent for $X = -1$. $\hat{s}_2$ is spread out, biased slightly toward a preference *not* to imitate d_1[14]. This tendency does not coincide with the BNE prediction, that player 2 should imitate d_1 for 62.5% of s_2. In fact, negative cascades are consistently among the more common cutoff choices for $X = -1$. The observed cutoffs reaffirm the message of choice outcomes, that externalities influence the observational learning process more than predicted by BNE.

5.3 What Determines Cutoff Decisions?

This paper investigates two reasons for player 2 to consider d_1: the information contained therein, and the externality player 2 receives if she imitates player 1. The purpose of this section is to quantify each of these influences on player 2's choice variable, $\hat{s}_2$. Recall equations 5 and 6:

$$s_2^{A*} = -.5 - .25 \times X$$
$$s_2^{B*} = .5 + .25 \times X$$

A regression of the observed s_2^A in all 3 treatments on X yields:

$$s_2^A = -\underset{(.020)}{.202} - \underset{(.024)}{.416} \times X \tag{10}$$

The results of the analogous regression for s_2^B:

$$s_2^B = \underset{(.019)}{.276} + \underset{(.023)}{.359} \times X \tag{11}$$

[14] s_2^B in the computer sessions of $X = -1$ are the exception. A possible explanation is the positive bias in cutoffs across treatments, but we offer no explanation as to why it is so pronounced in this particular case.

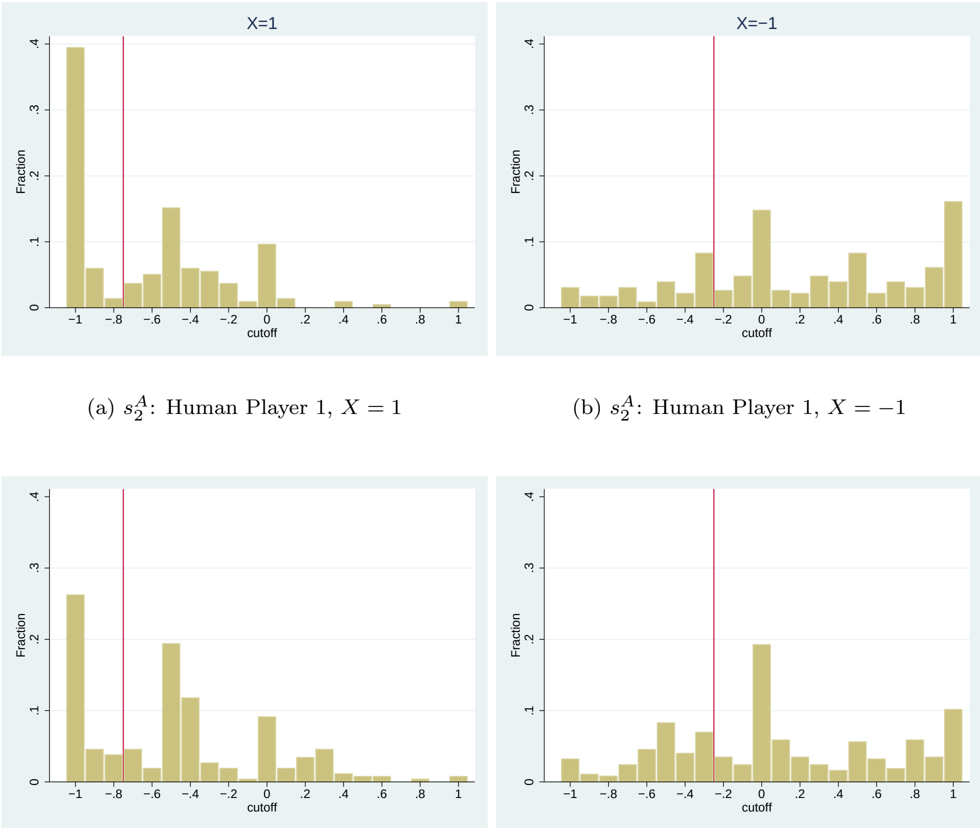

(a) s_2^A: Human Player 1, $X = 1$

(b) s_2^A: Human Player 1, $X = -1$

(c) s_2^A: Computer Player 1, $X = 1$

(d) s_2^A: Computer Player 1, $X = -1$

Figure 7: s_2^A for human vs. computer player 1: $X = 1$ and $X = -1$

The intercept represents player 2's approximate cutoff for $X = 0$. Across treatments, player 2 is sensitive to the information in d_2 (intercept highly significant), but less so than predicted. Player 2 appears to be *more* sensitive to X than predicted by BNE. The positive bias observed in both $\hat{s}_1$ and $\hat{s}_2$ can explain the different coefficients for s_2^A and s_2^B. The bias causes a more extreme intercept for s_2^B simply because s_2^{B*} is positive across treatments.

6 Quantal Response Equilibrium

For the analysis of experimental data, predictive models that allow no randomness in decisions need to be interpreted with caution. Section 5.2 shows that subjects' behavior differs from the theoretical prediction in a systematic fashion.

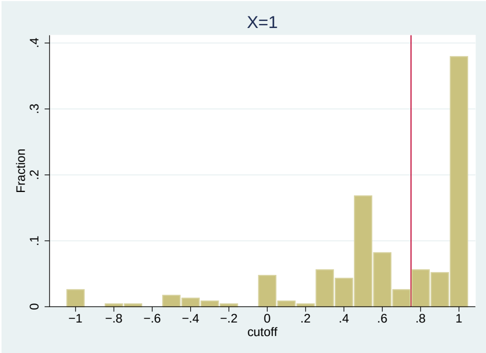

(a) s_2^B: Human Player 1, $X = 1$

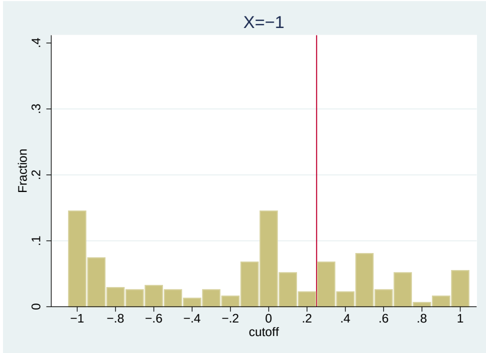

(b) s_2^B: Human Player 1, $X = -1$

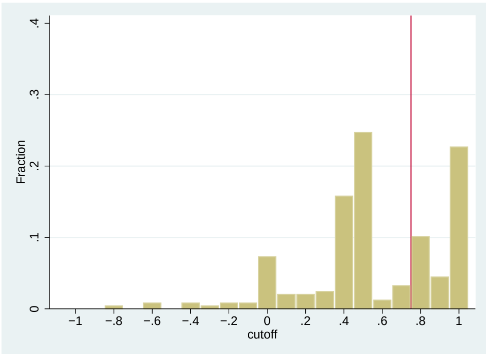

(c) s_2^B: Computer Player 1, $X = 1$

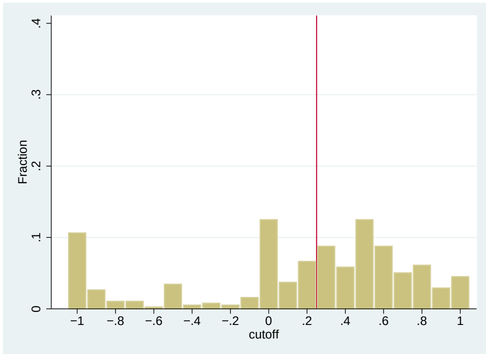

(d) s_2^B: Computer Player 1, $X = -1$

Figure 8: s_2^B for human vs. computer player 1: $X = 1$ and $X = -1$

Specifically, cutoffs are closer to their theoretical predictions when externalities are positive than when they are negative.

A possible confound, however, is the steeper payoff function under positive externalities. In all three treatments, information externalities are positive, in the sense that $d_1 = A$ makes action A a more attractive action for player 2. When $X = 1$, player 2 has an additional motivation to imitate d_1, and is therefore punished severely for not doing so. Under $X = -1$, on the other hand, information externalities motivate imitation, while payoff externalities motivate divergence. Although the information externalities tend to override payoff externalitie, in the sense that player 2 should still imitate for the majoirity of s_2, failing to imitate at least allows player 2 to avoid the negative payoff externality. In this sense, even severe deviations from BNE behavior are not severely punished in terms of payoffs for $X = -1$. Figure 9 shows player 2's expected payoff versus her chosen cutoff, assuming $d_1 = A$. Clearly, similar deviations from the optimal cutoff lead to higher payoff differentials for higher X. The concern is that, while the absolute deviation from the optimal cutoff strategy differs across treatments, the resulting payoff loss may not.

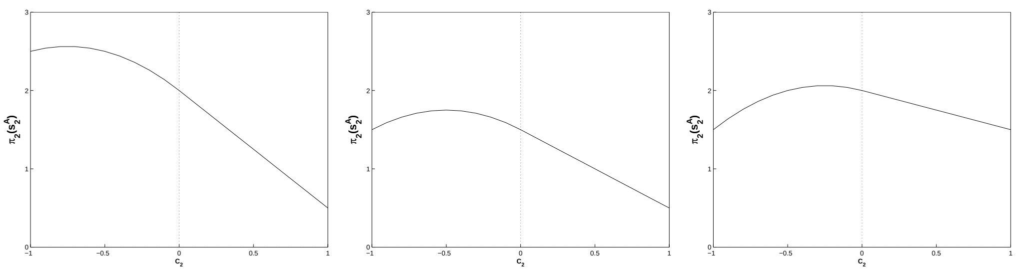

(a) $X = 1$ (b) $X = 0$ (c) $X = -1$

Figure 9: π_2 vs. s_2^A

The problem can be addressed through a Quantal Response Equilibrium (QRE) analysis. QRE assumes that, rather than best-responding, subjects 'better respond', randomizing between all actions but choosing those with higher expected payoffs with a higher probability, determined by a 'sensitivity parameter.' A comparison of this parameter across treatments will tell us if across-treatment differences discussed in the previous section might be a result of the differentially steep payoff functions of the different treatments.

We assume a logistic form of QRE [15]. Specifically, the probability that subject i chooses action j when actions are indexed by k is:

$$Pr(\hat{s}_i = j) = \frac{e^{\beta \pi_j}}{\sum_k e^{\beta \pi_k}} \tag{12}$$

β quantifies subjects' sensitivity to payoff differences. If $\beta = 0$, subjects are perfectly insensitive to payoff differences. Thus, behavior is uniformly random, and subjects choose each cutoff with equal probability. As β approaches ∞, behavior converges to BNE. For intermediate values of the parameter, subjects choose higher-paying cutoffs more often, but choose each cutoff with a positive probability.

[15]As QRE is discrete-choice model, we sorted the cutoffs into 21 bins, just as presented in figures5, 7 and C.3. $\hat{s}_2 \leq -.95$ were grouped together, as were $-.95 < \hat{s}_2 \leq -.85$. etc. When computations were necessary, all grouped cutoffs were assumed to be the midpoint of the range of the group

Table 7 show the β estimates for subjects in each role in each of the three treatments, along with their log-likelihoods. It also presents estimates and log-likelihoods for two alternative models: the random model and a simple non-equilibrium quantal-response (QR) model. The random model assumes that each cutoff is chosen with equal probability [16]. The QR model is simply the quantal response of player 1 to the *observed* cutoffs of player 2 and vice-versa.

	X=1		
	Random	*QR*	*QRE*
β_1	$0\ (-1367)$	$6.62\ (-1234)$	$6.59\ (-1235)$
β_2^A	$0\ (-1475)$	$2.52\ (-1171)$	$2.62\ (-1170)$
β_2^B	0	2.22	2.20
	X=0		
	Random	*QR*	*QRE*
β_1	$0\ (-1366)$	$4.79\ (-1283)$	$4.79\ (-1283)$
β_2^A	$0\ (-1392)$	$1.33\ (-1289)$	$1.58\ (-1288)$
β_2^B	0	3.16	2.91
	X=-1		
	Random	*QR*	*QRE*
β_1	$0\ (-1648)$	$2.49\ (-1608)$	$2.59\ (-1605)$
β_2^A	$0\ (-1692)$	$0.19\ (-1644)$	$0.42\ (-1644)$
β_2^B	0	0.19	0.12

Table 7: QRE Estimates by Treatment

The results presented in table 7 reinforce the cutoff data presented in the previous section. Observed cutoffs differ more from their theoretical predictions for negative externalities. Table 7 addresses the confound mentioned above, that noisier cutoffs may be a result of flatter payoff functions. The presented results suggest that subjects are more sensitive to payoffs when payoff externalities are positive than when they are negative. Errors increase across treatments as the payoff externality decreases, suggesting further that subjects are more responsive to payoff externalities than to information externalities.

A comparison in the values of the log-likelihood functions reaffirms the story. The QRE model improves substantially over the random model for $X = 1$, with the log-likelihood improving from -1475 to -1171 for player 2. There is much less improvement for $X = -1$, with the random model generating a log-likelihood of -1692 and the QRE -1644. The minimal improvement in fit suggests that sub-

[16] Cutoffs between -.9 and .9 are chosen with probability $\frac{1}{20}$, while cutoffs of -1 and 1 are chosen with probability of $\frac{1}{40}$ each, due to the way the continuous cutoffs were sorted into discrete bins.

jects do a poor job of setting the most profitable cutoffs when payoff externalities are negative, so much so that behavior is almost random.

7 Discussion

This paper finds that subjects are relatively insensitive to the information contained in the decisions that they observe, and highly sensitive to accompanying payoff externalities. As evidence for the first finding, we show that second-movers rely on her private information in situations where a Bayesian would imitate first-movers. The second finding is supported by the fact that second-movers' imitation set increases (decreases) in the presence of positive (negative) payoff externalities than predicted by BNE. This section discusses the application to more generalized environments, and addresses several confounds.

We add one layer of realism to the majority of the observational learning studies. We find that subjects react strongly to the additional incentive provided by payoff externalities. It must be said, however, that there is nothing that ties this reaction to the specific form of this incentive. It may be that any separate, salient incentive to choose one action over another affects the observational learning process in the same way, be it by taking the same action as one's predecessor, or an entirely different type of motivation. In other words, observational learning is difficult, and clearer motivations have stronger effects on decisions. The best we can do is to state that observational learning environments with payoff externalities have the payoff structure of the game that we study. Our results therefore pertain to observational learning situations, but may not be unique to them.

The computer treatments addressed confounds of loss aversion (to some degree), social preferences and over-reaction to noisy play by player 1. At least one other simple heuristic can account for the under-responsiveness to observation that we find in $X = 0$ and $X = -1$ treatments. Under-learning corresponds to cutoffs close to 0, which is the median of the decision set. Menu effects have shown that, in the absence of strong preferences, subjects often tend to make choices close to the median of available options[17]. However, such a heuristic would not predict the increase in cascades under negative and (especially) positive payoff externalities.

A major motivation for the study of observational learning is the efficiency of the process. We employ sequences of only two subjects, partially due to the theoretical clarity that such a simple setup affords. Based on the $X = 0$ treatment, we

[17]See Kahneman, Schkade and Sunstein (1998) for an example.

identify two reasons why the information aggregation process is less efficient than predicted by theory. First, player 1's somewhat imprecise behavior retards the available information. More importantly, player 2 does not incorporate all of the information that is nonetheless available in player 1's decision into her own. These two effects combine to make player 2's action even less reflective of the players' combined information than predicted by theory. Therefore, a third mover would find player 2's action relatively uninformative.

However, the typical observational learning model[18] is characterized by perfect information, meaning that DM's observe *all* predecessors. Player 2's action alone is less informative than predicted by BNE specifically because it is less correlated with player 1's action. This makes a combination of player 1 and player 2's actions collectively *more* informative to a sophisticated third mover under perfect information. In other words, the reluctance toward observational learning, in some cases, will counteract the information suppression described in Banerjee (1992) and BHW.

Our results unambiguously suggest that cascades will occur more often with positive payoff externalities than without. Further, the frequency of cascade behavior on the part of player 2 suggests that informational cascades will often reflect only the information of the first mover. In our experiment, the positive externality improves subjects' performance with respect to theory. It is possible, however, that this fact is dependent on our particular experimental design. Subjects can over-adjust their cutoff by a maximum of .25 to positive externalities, and 1.25 to negative (for example, by setting $s_2^A = -1$ and 1 for $X = 1$ and -1, respectively). A different design may find that positive externalities can severely damage the learning process as well.

Our results suggest an unambiguous decline in player 2's performance across a wide variety of environments with negative payoff externalities. The effect on a longer sequence of subjects is unclear. Drehman, Oechssler and Roider (2007) find that subjects tend to divide themselves evenly between actions under negative externalities, suggesting that decisions have little correlation to the state of the world. This piece of evidence suggests that subjects beyond the role of the second-mover continue to have difficulty with observational learning in the presence of negative payoff externalities[19].

[18] Çelen and Kariv (2004b) is an exception.

[19] Again, the complexity of equilibria makes it difficult to determine whether subjects act optimally.

8 Conclusion

It is natural for economic decision-makers to consider similar decisions made in the past. The majority of the existing literature has focused on a model where observation is useful only in revealing private information. In many real-world situations, past choice outcomes may fundamentally change the tradeoff between alternatives, rather than simply revealing information about them. Hung and Plott (2002) and Drehman, Oechssler and Roider (2007) are among the few studies of payoff externalities on observational learning. Their models involve binary action sets which do not allow for a precise measure of the influence that payoff externalities exert on decision-making.

Our continuous signal design, coupled with a cutoff-elicitation technique grants a precise measure of observational learning. Across-treatment comparison permits the comparison of the effects of information and payoff externalities. We find that subjects are more sensitive to payoff externalities, and less sensitive to information externalities, than predicted by BNE. This finding is perhaps best represented in the regressions of equations 10 and 11, which show that choices are under-influenced by d_1, and over-influenced by payoff X, relative to BNE. The results did not change significantly when a computer played the role of player 1, discounting the likelihood of alternative explanations such as social preferences and a mistrust of player 1. QRE analysis shows that our results are robust to payoff differentials.

We conclude that the information aggregation process is muted because it is inherently difficult. The pursuit of externalities is simpler, and is therefore likely to be represented more pronouncedly in decisions. This clearly will hinder second-movers under negative externalities, with ambiguous effects for second-movers under positive externalities. As discussed in section 7, extension to longer sequences of decision-makers is a more complex discourse. Our results do suggest, however, that the dynamics in longer decision-making processes are likely to differ from theoretical predictions, though further study is necessary to determine exactly how.

References

[1] Anderson, Lisa R. and Charles A. Holt, "Information Cascades in the Laboratory," *American Economic Review*, 87(5), 847-862.

[2] Banerjee, Abhijit (1992), "A Simple Model of Herd behavior," *Quarterly Journal of Economics*, 107(3), 797-817.

[3] Bikchandani, Sushil, David Hirschleifer and Ivo Welch (1992), "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascade," *Journal of Political Economy*, 100(5), 992-1026.

[4] Bikchandani, Sushil, David Hirschleifer and Ivo Welch (1998), "Learning From the Behavior of Others: Conformity, Fads and Information Cascades," *Journal of Economic Prospectives*, 12(3), 151-170.

[5] Çelen, Boğaçhan and Shachar Kariv (2004a), "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," *American Economic Review*, 94(3), 484-497.

[6] Çelen, Boğaçhan and Shachar Kariv (2004b), "Observational Learning Under Imperfect Information," *Games and Economic Behavior*, 47(1), 72-86.

[7] Charness, Gary and Matthew Rabin (2002), "Understanding Social Preferences with Simple Tests," *Quarterly Journal of Economics*, 117(3), 817-869.

[8] Drehman, Mathias, Jörg Oechssler and Andreas Roider (2007), "Herding with and without payoff externalities - an internet experiment," *International Journal of Industrial Organization*, 25, 391-415.

[9] Hung, Angela A. and Charles R. Plott, "Information Cascades: Replication and Extension to Majority Rule and Conformity-Rewarding Institutions," *American Economic Review*, 91(5), 1508-1520.

[10] McKelvey, Richard D. and Thomas R. Palfrey (1995), "Quantal Response Equilibrium in Normal Form Games", *Games and Economic Behavior* 10, 6-38.

[11] Smith, Lones and Peter Sørensen (2000), "Pathological Outcomes of Observational Learning", *Econometrica*, 68(2), 371-398.

A Addendum to Theory

A.1 Theoretical Predictions for Generic s_i Distributions

Proposition 2. *Player 2 best-responds to* any *strategy employed by player 1 by using a* cutoff strategy*. A cutoff strategy is one such that there exists some signal $\hat{s_2}$ such that $d_2^* = A$ if $s_2 \geq \hat{s_2}$, and $d_2^* = B$ if $s_2 \leq \hat{s_2}$.*

Proof. Player 1 decides between actions A and B after observing player 1's action. d_1 informs player 2 which action yields the externality, X, as well as allowing player 2 to update her prior beliefs as to the identity of action H. Player 2 uses this information in conjunction with her private signal s_2 in order to decide which action, A or B, yields the higher expected payoff. For now, assume that $d_1 = A$ [20]:

$$\begin{aligned} E\pi_2(A|d_1 = A, s_2) &= X + Y \times Pr(A = H|s_2, d_1) \\ &= X + Y \times Pr(s_1 + s_2 \geq 0|s_2, d_1) \\ &= X + Y \times Pr(s_1 \geq -s_2|s_2, d_1) \\ &= X + Y \times [1 - F_{s_1|d_1=A}(-s_2)] \end{aligned} \tag{13}$$

$F_{s_1|d_1=A}$ is the conditional *cdf* of s_1 given $d_1 = A$. Similarly:

$$E\pi_2(B|d_1 = A, s_2) = Y \times F_{s_1|d_1=A}(-s_2) \tag{14}$$

Player 2 maximizes her expected payoff by choosing action A if:

$$\begin{aligned} E\pi_2(A|d_1 = A, s_2) \geq& E\pi_2(B|d_1 = A, s_2) \\ X + Y \times [1 - F_{s_1|d_1=A}(-s_2)] \geq& Y \times F_{s_1|d_1=A}(-s_2) \\ F_{s_1|d_1=A}(-s_2) \leq& \frac{1}{2} \times \left(1 + \frac{X}{Y}\right) \end{aligned} \tag{15}$$

[20] Note that we do *not* assume symmetry, that $E\pi_2(A|d_1 = A) = E\pi_2(B|d_1 = B)$ and $E\pi_2(A|d_1 = B) = E\pi_2(B|d_1 = A)$. Symmetry is proven below

Similarly, for $d_1 = B$:

$$\begin{aligned} E\pi_2(A|d_1 = B, s_2) \geq & E\pi_2(B|d_1 = B, s_2) \\ Y \times [1 - F_{s_1|d_1=B}(-s_2)] \geq & X + Y \times F_{s_1|d_1=B}(-s_2) \\ F_{s_1|d_1=B}(-s_2) \leq & \frac{1}{2} \times \left(1 - \frac{X}{Y}\right) \end{aligned} \tag{16}$$

As *cdf*'s are increasing functions, $F_{s_1|d_1=A}(-s_2)$ is decreasing in s_2, meaning that if $E\pi_2(A|d_1, s_2) > E\pi_2(B|d_1, s_2)$ for a given value of s_2, it is also true for all higher values of s_2. **In other words, player 2 uses a cutoff strategy for *any* strategy employed by player 1.** □

We now define $\hat{s_2}$, player 2's cutoff. $\hat{s_2}$ is the lowest s_2 for which player 2 chooses $d_2 = A$. Of course, player 2 may behave very differently if she observes $d_1 = A$ versus $d_1 = B$. We therefore divide $\hat{s_2}$ into two classes, classifying it as s_2^A following $d_1 = A$, and s_2^B following $d_1 = B$. Additionally, we define $\hat{s_2}$, player 2's optimal cutoff strategy, the s_2 such that $E\pi_2(A|s_2, d_1) \geq E\pi_2(B|s_2, d_1)$. Note that this is a two-part strategy, s_2^{A*} being player 2's optimal cutoff after observing $d_1 = A$, and s_2^{B*} after observing $d_1 = B$.

$$s_2^{A*} \equiv -F^{-1}_{s_1|d_1=A}\left(\frac{1}{2} \times \left(1 + \frac{X}{Y}\right)\right) \tag{17}$$

and

$$s_2^{B*} \equiv -F^{-1}_{s_1|d_1=B}\left(\frac{1}{2} \times \left(1 - \frac{X}{Y}\right)\right) \tag{18}$$

Proposition 3. *Player 1 best-responds to* any *cutoff strategy employed by player 2 by also employing a cutoff strategy.*

Proof. We now consider player 1's decision problem. To begin with, we assume that player 2 employs an arbitrary cutoff strategy, (s_2^A, s_2^B).

$$\begin{aligned} E\pi_1(A|s_1) =& Y \times [Pr(A = H|s_1)] + X \times [Pr(d_2 = A|d_1 = A)] \\ =& Y \times [Pr(s_1 + s_2 \geq 0|s_1)] + X \times [Pr(s_2 > s_2^A)] \\ =& Y \times [Pr(s_2 \geq -s_1|s_1)] + X \times [Pr(s_2 > s_2^A)] \\ =& Y \times [1 - F_{s_2}(-s_1)] + X \times [1 - F_{s_2}(s_2^A)] \end{aligned} \tag{19}$$

Where $F_{s_2}()$ is the *cdf* of s_2. Similarly:

$$E\pi_1(B|s_1) = Y \times [F_{s_2}(-s_1)] + X \times [F_{s_2}(s_2^B)] \tag{20}$$

and $E\pi_1(A|s_1) \geq E\pi_1(B|s_1)$ if:

$$F_{s_2}(-s_1) \leq \frac{1}{2} \times \left[1 + \frac{X}{Y}[1 - F_{s_2}(s_2^A) - F_{s_2}(s_2^B)]\right] \tag{21}$$

All terms on the right-hand side of equation 21 are independent of s_1. As F_{s_2} is an increasing function, $F_{s_2}(-s_1)$ is decreasing in s_1. **Therefore, player 1 uses a cutoff strategy for any cutoff strategy employed by player 2. Combined with proposition 1, this implies that both players employ a cutoff strategy in any equilibrium** □

We define $\hat{s_1}$ as player 1's cutoff, and $\hat{s_1^*}$ as her optimal BNE cutoff, using equation 21, as follows:

$$\hat{s_1}^* = F_{s_2}^{-1}\left(\frac{1}{2} \times \left[1 + \frac{X}{Y}[1 - F_{s_2}(s_2^A) - F_{s_2}(s_2^B)]\right]\right) \tag{22}$$

A.2 BNE Derivation for $s_i \sim U[-1, 1]$, $X \neq 0$

We begin by arguing that, under $X \neq 0$, both players still employ cutoff strategies in equilibrium. Equation 23 shows player 1's profits for choosing A and B for *any* strategy employed by player 2 (in particular, we do not assume that player 2 uses a cutoff strategy). Pr_{jk} is the probability that, given player 2's strategy, she chooses action j after observing $d_1 = k$.

$$\pi_{1A} = Y \times \frac{s_1 + 1}{2} + Pr_{AA} \times X \qquad \pi_{1B} = Y \times \frac{1 - s_1}{2} + Pr_{BB} \times X \tag{23}$$

As π_{1A} is increasing in s_1 and π_{1B} is decreasing, there will be some $\hat{s}_1$ such that $\pi_{1A} \geq \pi_{1B}\ \forall\ s_1 \leq \hat{s}_1$ and $\pi_{1A} \leq \pi_{1B} \forall s_1 \geq \hat{s}_1$. In other words, *in any BNE, player 1 employs a cutoff strategy.* Note that the proof did not assume that player 2 uses a cutoff strategy. We use this to update player 2's profit function. π_{2jk} is player 2's expected payoff for choosing action j after observing $d_1 = k$, as a function of s_2.

$$\pi_{2AA} = \begin{cases} Y \times \frac{s_2+1}{1-\hat{s_1}} + X & \text{if} \quad s_2 \leq -\hat{s_1} \\ Y + X & \text{if} \quad s_2 > -\hat{s_1} \end{cases} \qquad \pi_{2BA} = \begin{cases} Y \times \frac{-\hat{s}_1 - s_2}{1-\hat{s_1}} & \text{if} \quad s_2 \leq -\hat{s_1} \\ 0 & \text{if} \quad s_2 > -\hat{s_1} \end{cases}$$

$$\pi_{2AB} = \begin{cases} 0 & \text{if } s_2 \leq -\hat{s}_1 \\ Y \times \frac{s_2+\hat{s}_1}{1+\hat{s}_1} & \text{if } s_2 > -\hat{s}_1 \end{cases} \quad \pi_{2BB} = \begin{cases} Y + X & \text{if } s_2 \leq -\hat{s}_1 \\ Y \times \frac{1-s_2}{1+\hat{s}_1} + X & \text{if } s_2 > -\hat{s}_1 \end{cases}$$

We can solve for player 2's optimal strategy:

$$s_2^{A*} = -\frac{1}{2} \times \left(1 + \frac{X}{Y}\right) + \frac{\hat{s}_1}{2}\left(\frac{X}{Y} - 1\right) \quad s_2^{B*} = \frac{1}{2}\left(1 + \frac{X}{Y}\right) + \frac{\hat{s}_1}{2}\left(\frac{X}{Y} - 1\right) \tag{24}$$

As player 2's optimal strategy is a function of $\hat{s}_1$, player 1 must consider this influence in her own strategy. Player 1's expected payoff:

$$\pi_1 = Pr(d_1 = \omega) \times Y + Pr(d_1 = d_2) \times X \tag{25}$$

$$Pr(d_1 = \omega) = \frac{3 - \hat{s}_1^2}{4} \tag{26}$$

$$Pr(d_1 = d_2) = Pr(d_1 = A|\hat{s}_1) \times Pr_{AA} + Pr(d_1 = B|\hat{s}_1) \times Pr_{BB} \tag{27}$$

$$= \frac{1 - \hat{s}_1}{2} \times \frac{1 - s_2^A}{2} + \frac{1 + \hat{s}_1}{2} \times \frac{s_2^B + 1}{2} \tag{28}$$

Substituting equations 24, 26 and 28 into equations 25 yields:

$$\pi_1 = Y \times \frac{3 - \hat{s}_1^2}{4} + X \times \left[\frac{3}{4} + \frac{1}{4}\left[\frac{X}{Y} + \hat{s}_1^2\left(\frac{X}{Y} - 1\right)\right]\right] \tag{29}$$

Meaning:

$$\frac{\partial \pi_1}{\partial \hat{s}_1^2} = \frac{1}{4}\left[-Y - X \times \left(1 - \frac{X}{Y}\right)\right] \tag{30}$$

As this number is negative for all $-Y < X < Y$, $\hat{s}_1^* = 0$ for all such values of X. As for player 2's optimal strategy, substituting $\hat{s}_1 = 0$ into equation 24 yields: $s_2^{A*} = -\frac{1}{2} \times \left(1 + \frac{X}{Y}\right)$ and $s_2^{B*} = \frac{1}{2}\left(1 + \frac{X}{Y}\right)$. Thus, with $Y = 2$, $s_2^{A*} = \{-.75, -.5, -.25\}$ and $s_2^{A*} = \{.75, .5, .25\}$ for $X = \{1, 0, -1\}$, respectively.

B Over-Reacting to X?

As discussed in section 3, the theoretical prediction for each treatment is symmetric, as $\hat{s_1}^* = 0$, and $\hat{s_2^A}^* = -\hat{s_2^B}^*$. This facilitates a straightforward transformation of the data that makes data presentation more manageable. Player 1 always faces the same decision, and her decision thus defines her strategy. Player 2, on the other hand, reveals her decision for only one of her two information sets (the observed $d_1 \in \{A, B\}$). We impose symmetry on her strategy, assuming that $s_2^A = -s_2^B$. We construct $\hat{s_2}$ as follows:

$$\hat{s_2} \equiv \begin{cases} s_2^A & \text{if } d_1 = A \\ -s_2^B & \text{if } d_1 = B \end{cases} \tag{31}$$

The slope terms in equations 10 and 11 imply that player 2 attaches too much importance to the externality, X. This is true, to some degree. If player 2 were to extract all of the available information from player 1's action, this adjustment is too strong. It is argued in section 5, however, that player 2 extracts too little information from player 1's action.

The analysis that follows takes it as given that player 2 under-reacts to the information contained in player 1's action. It poses the question, *given the limited information that player 2 extracts, does she player 2 react appropriately to the network externality?*

This, of course, requires an assumption about player 2's unobserved belief. Equations 10 and 11 imply that player 2 behaves as though $\hat{s_2} = -.244$ is a best response when $X = 0$. We now generate beliefs that rationalize this behavior. *Any belief for which $\hat{s_2} = -.244$ is a best response can be rationalized.* Equivalently, any distribution of $\hat{s_1}$ for which $\hat{s_2} = -.244$ is a best response is a possible belief.

For lack of a better alternative, we assume that player 2 believes that player 1's strategy is a deterministic cutoff strategy, but her belief of player 1's cutoff changes based on whether she observes $d_1 = A$ or $d_1 = B$:

$$\tilde{f}(s_1|d_1 = B) = \begin{cases} \frac{1}{1.466} & \text{for } s_1 \leq .466 \\ 0 & \text{for } s_1 > .466 \end{cases} \tag{32}$$

Player 2 reacts as though $\hat{s_1} = .466$ if $d_1 = B$, and as though $\hat{s_1} = -.466$ if $d_1 = A$. Substitution into equation 10 (along with $Y = 2$) yields: $\hat{s_2} = -.267 - .367 \times X$.

The empirical weight from equation 10, -.383, suggests that player 2's over-weighting of the network externality may be rational given their under-appreciation of the informational content of player 1's action. Within the context of observational learning experiments, the findings suggest that positive externalities will improve decision-making. On the other hand, decision-making could be increasingly harmed as externalities become more negative.

C Deriving Player 2's Optimal Response to a Noisy Player 1

The analysis above shows that players do not strictly perform as predicted by theory. To that end, it is necessary to consider how a rational decision-maker should respond to such players. A claim of this paper is that subjects in the role of player 2 do not react strongly enough to the information contained in the actions of player 1. If, however, the inherent noise in the decision-making process renders the observation of player 1's action obsolete, then it is perfectly rational for player 1 to under-react to this action (relative to theory). It is the aim of this section to evaluate the decisions of player 2, accounting for the randomness in player 1's behavior.

C.1 A Noisy Player 1

Assume that player 1 is a stochastic player. We will treat her chosen cutoff, $\hat{s_1}$, as the realization of a random process represented by the random variable $\hat{S}_1$. Assume that S_1 has support $[-1, 1]$ and cumulative density function $F(\hat{s_1})$. Without loss of generality, this analysis assumes that player 2 observes $d_1 = A$, meaning that player 1's signal $s_1 > \hat{s_1}$.

The aim of this section is to determine the optimal cutoff for player 2, $\hat{s_2^*}$, given that $\hat{s}_1$ is stochastic. The procedure used here is to first use Bayes' Rule to generate the updated distribution $F(\hat{s_1}|d_1 = A)$, and to use this distribution to find $F(s_1|d_1 = A)$. Finally, these distributions will be used to define player 2's maximization problem, and find $\hat{s_2^*}$.

The role of Bayesian updating is vital in the current problem. Imagine, for example, that $hats_1$ is distributed symmetrically (about 0) and continuously. Ex ante, $f(\hat{s}_1) = f(-\hat{s}_1)$ for any $\hat{s}_1 \in [-1, 1]$. However, if $hats_1 > 0$, $Pr(d_1 = A|\hat{S}_1 = \hat{s}_1) < Pr(d_1 = A|\hat{S}_1 = -\hat{s}_1)$ if $\gamma > 0$. Therefore, the observation that $d_1 = A$ skews the distribution $F(\hat{s}_1|d_1 = A)$ to the left. In other words, because lower cutoffs make the choice $d_1 = A$ more likely, player 2's observation that $d_1 = A$ makes lower $\hat{s}_1$ more likely. Formally,

$$Pr(\hat{S}_1 = \hat{s_1}|d_1 = A) = \frac{Pr(\hat{S}_1 = \hat{s_1}) \times Pr(d_1 = A|\hat{S}_1 = \hat{s_1})}{Pr(d_1 = A)} \tag{33}$$

$$f(\hat{s_1}|d_1 = A) = \frac{f(\hat{s_1}) \times (1 - F_{S_1}(\hat{s_1}))}{\int_{-1}^{1} f(\hat{s_1}) \times (1 - F_{S_1}(\hat{s_1}))d\hat{s_1}} \tag{34}$$

where $F_{s_1}(\hat{s_1})$ is the cumulative density function of S_1 evaluated at, here $F(s_1) = \frac{s_1+1}{2}$. The next step in determining $\hat{s_2^*}$ is to find the conditional probability density function of S_1 for a given value of $\hat{s_1}$. Given $d_1 = A$, there is 0 probability attached to values of $S_1 < \hat{s_1}$. Therefore, the conditional distribution $f(s_1|d_1 = A, \hat{s_1})$ condenses the probability function $f(s_1)$ to values of $s_1 \geq \hat{s_1}$.

$$Pr(S_1 = s_1|d_1 = A, \hat{s_1}) = \begin{cases} 0 & \text{for} \quad s_1 \leq \hat{s_1} \\ \frac{Pr(S_1 = s_1)}{Pr(S_1 \leq \hat{s_1})} & \text{for} \quad s_1 > \hat{s_1} \end{cases} \tag{35}$$

$$f(s_1|d_1 = A, \hat{s_1}) = \frac{f(s_1|d_1 = A, \hat{s_1})}{\int_{\hat{s_1}}^{1} f(s_1|d_1 = A, \hat{s_1})} \tag{36}$$

We now have the conditional distributions for both $\hat{s_1}$ and s_1, which will be used to evaluate player 2's maximization problem. Figure C.1 illustrates player 2's problem for values of $\hat{s_1} \leq -\hat{s_2}$. Player 2 can alter the probability of her four outcomes $Y + X$, Y, X and 0 by adjusting $\hat{s_2}$. Due to the assumption that $f(\hat{S_1})$ has a support of $[-1, 1]$, there is a positive probability that $\hat{s_1} > -\hat{s_2}$, and as player 2 does not observe the realization of $\hat{s_1}$, this possibility must be accounted for in the analysis. Figure C.1 shows the regions as defined for low values of $\hat{s_1}$.

The state represented in figure C.1 is less likely than that in figure C.1, as $d_1 = A$ makes low values of $\hat{s_1}$ more likely and high values less likely than in the unconditional distribution $f(s_1)$. The probability calculation must be done in two parts: the case involving $\hat{s_1} \leq -\hat{s_2}$, and that involving $\hat{s_1} \geq -\hat{s_2}$. The following as a general calculation for the probability of receiving each of the payoffs for a general $\hat{s_2}$, over which we will then optimize to find $\hat{s_2^*}$. We begin with the probability of a given $\hat{s_2}$ resulting in the payoff $\pi_2 = Y + X$, defined the irregular pentagon in the top-right of figure C.1 or the rectangle in the top-right of figure C.1.

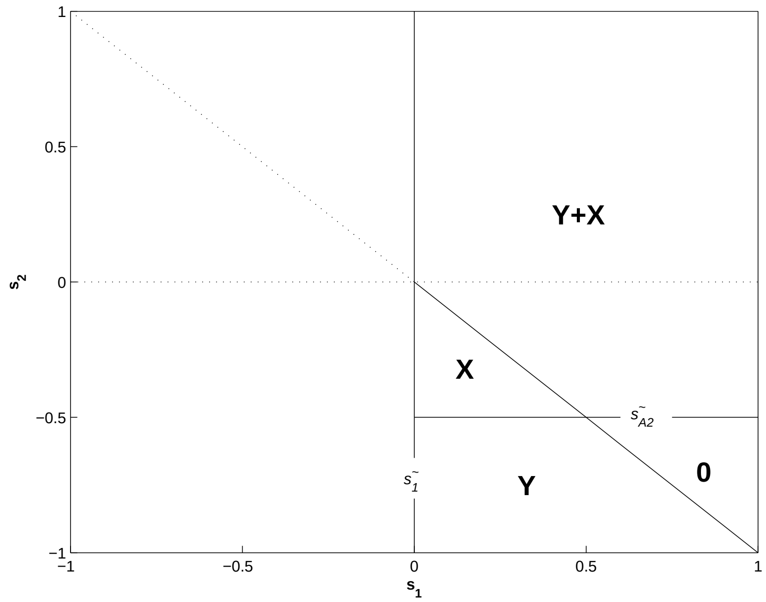

Figure 10: Assigning π_2 to Regions in the s_1, s_2 Plane for $\hat{s_1} \leq -\hat{s_2}$ and $d_1 = A$

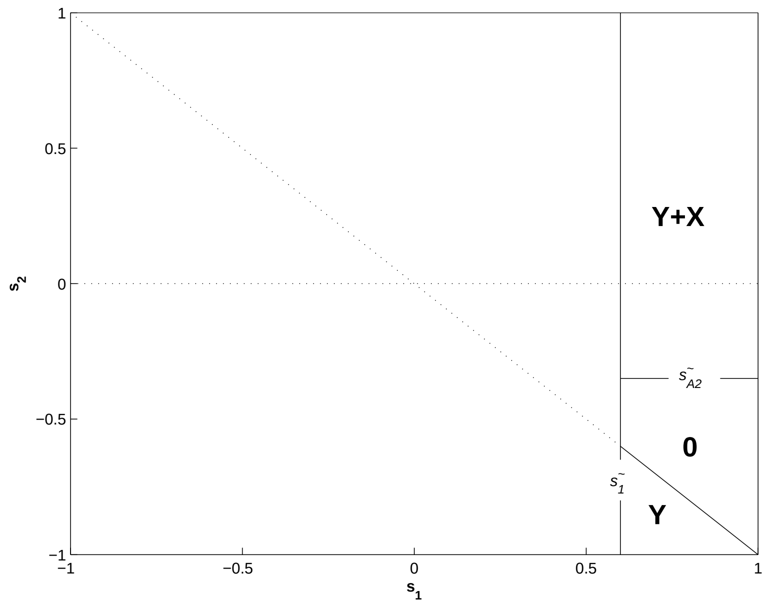

Figure 11: Assigning π_2 to Regions in the S_1, S_2 Plane for $\hat{s_1} \geq -\hat{s_2}$ and $d_1 = A$

$$
\begin{aligned}
Pr(\pi_2 = Y + X | d_1 = A) = & \int_{-1}^{-\hat{s_2}} \int_{-\hat{s_1}}^{1} \int_{\hat{s_1}}^{1} f(s_1 | d_1 = A, \hat{s_1}) ds_1 f(s_2) ds_2 f(\hat{s_1} | d_1 = A) d\hat{s_1} \\
& + \int_{-1}^{-\hat{s_2}} \int_{-\hat{s_1}}^{\hat{s_2}} \int_{-S_2}^{1} f(s_1 | d_1 = A, \hat{s_1}) ds_1 f(s_2) ds_2 f(\hat{s_1} | d_1 = A) d\hat{s_1} \\
& + \int_{-\hat{s_2}}^{1} \int_{\hat{s_2}}^{1} \int_{\hat{s_1}}^{1} f(s_1 | d_1 = A, \hat{s_1}) ds_1 f(s_2) ds_2 f(\hat{s_1} | d_1 = A) d\hat{s_1}
\end{aligned} \tag{37}
$$

where

$$
f(s_1 | d_1 = A, \hat{s_1}) = \frac{f(s_1 | d_1 = A)}{\int_{\hat{s_1}}^{1} f(s_1 | d_1 = A) ds_1} \tag{38}
$$

The function $f(s_1|D_1 = A, \hat{s_1})$ is the conditional distribution of s_1 given that $d_1 = A$ for a *given* $\hat{s_1}$ (as the *pdfs* are then integrated over all values of $\hat{s_1}$). $f(s_1|d_1 = A, \hat{s_1})$ shifts the all of the density of $f(s_1|d_1 = A)$ to possible values of s_1, those greater than the given $\hat{s_1}$.

The first two lines of equation 37 correspond to the pentagon from the top-right of figure C.1, while the third line represents the rectangle at the bottom-left of figure C.1. Notice that the definition of the regions corresponding to high values of $\hat{s_1}$ integrate with respect to $\hat{s_1}$ from $-\hat{s_2}$ to 1, while that corresponding to low values of $\hat{s_1}$ integrate with limits of -1 and $-\hat{s_2}$. Similarly, the probability that $\pi_2 = X$, Y, and 0 are computed as follows:

$$Pr(\pi_2 = X|d_1 = A) = \int_{-1}^{-\hat{s_2}} \int_{\hat{s_2}}^{-\hat{s_1}} \int_{\hat{s_1}}^{-s_2} f(s_1|d_1 = A, \hat{s_1}) ds_1 f(s_2) ds_2 f(\hat{s_1}|d_1 = A) d\hat{s_1} \tag{39}$$

$$\begin{aligned} Pr(\pi_2 = Y|d_1 = A) = & \int_{-1}^{-\hat{s_2}} \int_{-1}^{\hat{s_2}} \int_{\hat{s_1}}^{-s_2} f(s_1|d_1 = A, \hat{s_1}) ds_1 f(s_2) ds_2 f(\hat{s_1}|d_1 = A) d\hat{s_1} \\ & + \int_{-\hat{s_2}}^{1} \int_{-1}^{-\hat{s_1}} \int_{\hat{s_1}}^{-s_2} f(s_1|d_1 = A, \hat{s_1}) ds_1 f(s_2) ds_2 f(\hat{s_1}|d_1 = A) d\hat{s_1} \end{aligned} \tag{40}$$

$$\begin{aligned} Pr(\pi_2 = 0|d_1 = A) = & \int_{-1}^{-\hat{s_2}} \int_{\hat{s_2}}^{-1} \int_{-s_2}^{1} f(s_1|d_1 = A, \hat{s_1}) ds_1 f(s_2) ds_2 f(\hat{s_1}|d_1 = A) d\hat{s_1} \\ & + \int_{-\hat{s_2}}^{1} \int_{\hat{s_1}}^{1} \int_{-s_1}^{\hat{s_2}} f(s_2) ds_2 f(s_1|d_1 = A, \hat{s_1}) ds_1 f(\hat{s_1}|d_1 = A) d\hat{s_1} \end{aligned} \tag{41}$$

The four probabilities calculated above are functions of $\hat{s_2}$. Player 2's problem, then, is to choose the value of $\hat{s_2}$ that maximizes her expected profit, using the probabilities calculated above.

$$E\pi_2 = (Y + X) \times Pr(\pi_2 = Y + X) + X \times Pr(\pi_2 = X) + Y \times Pr(\pi_2 = Y) + 0 \times Pr(\pi_2 = 0) \tag{42}$$

C.2 $\hat{s_1}$ Uniformly Distributed

For illustrative purposes, we now impose a simple structure on $\hat{s_1}$. Specifically, we assume that $\hat{s_1}$ is randomly drawn from a uniform distribution on the interval $[-1, 1]$, meaning that $f(\hat{s_1}) = \frac{1}{2} \forall \hat{s_1}$. Recall further that, as s_1 and s_2 are also uniformly distributed, $f(s_i) = \frac{1}{2}$ and $F(s_i) = \frac{s_i+1}{2}$ for $i = 1, 2$. Substituting these values into equations 34 and 36, we can retrieve the conditional *pdf*s $f(\hat{s_1}|d_1 = A)$ and $f(s_1|d_1 = A, \hat{s_1})$. As the probability function will be calculated by integrating with respect to s_1 before $\hat{s_1}$, $f(s_1|d_1 = A, \hat{s_1})$ should *not* take into account the distribution of $\hat{s_1}$.

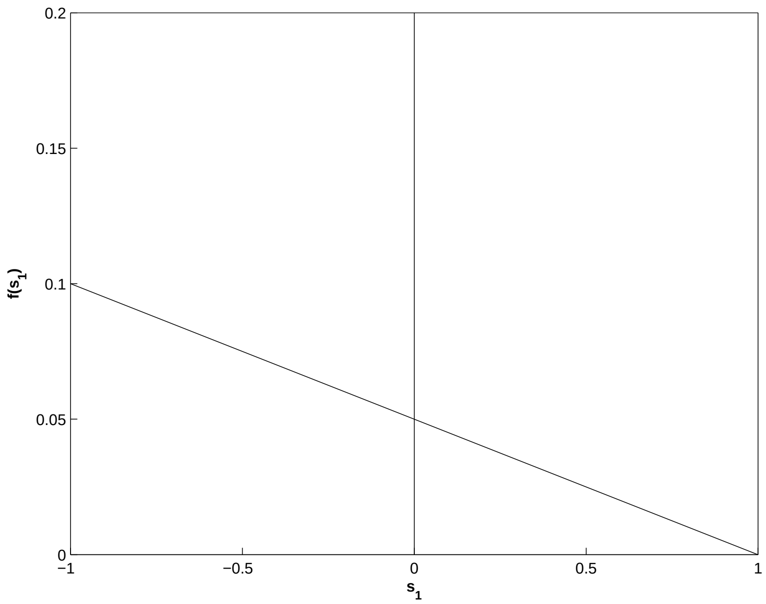

Figure 12: $f(\hat{s_1}|d_1 = A)$ for $\hat{S_1} \sim U[-1, 1]$

$$f(\hat{s_1}|d_1 = A) = \frac{\frac{1}{2} \times \frac{1-\hat{s_1}}{2}}{\int_{-1}^{1} \frac{1}{2} \times \frac{1-\hat{s_1}}{2} d\hat{s_1}} = \frac{1-\hat{s_1}}{2} \tag{43}$$

$$f(s_1|d_1 = A, \hat{s_1}) = \frac{\frac{1}{2}}{\int_{\hat{s_1}}^{1} \frac{1}{2} ds_1} = \frac{1}{1-\hat{s_1}} \tag{44}$$

As seen in figure C.2, $f(\hat{s_1}|d_1 = A)$ is skewed to the right, with low values of $\hat{s_1}$ more likely than high values, as expected. Figure C.2 shows the an example of the distribution $f(s_1|d_1 = A, \hat{s_1})$, specifically $f(s_1|d_1 = A, \hat{s_1} = -.5)$. The distribution is still flat, as it does not account for the fact that $\hat{s_1}$ is known.

Solving the integrals in equations 37 through 41, and substituting these values into player 2's objective function, 42, we can retrieve a profit-maximizing $\hat{s_2}$, $\hat{s_1}$, for each treatment. Notice that noise associated with the noisy play of player 1 does not change player 2's optimal action substantially.

1. $X = 1$: $\hat{s_2} = -(\sqrt{2} - 1)$
2. $X = 0$: $\hat{s_2^*} = -(\sqrt{3} - 1)$
3. $X = -1$: $\hat{s_2^*} = 0$

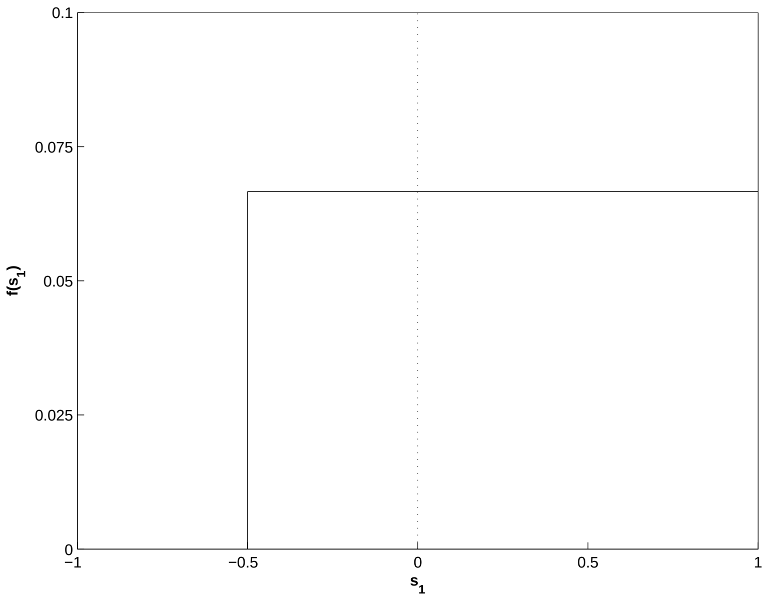

Figure 13: $f(s_1|d_1 = A, \hat{s_1} = -.5)$

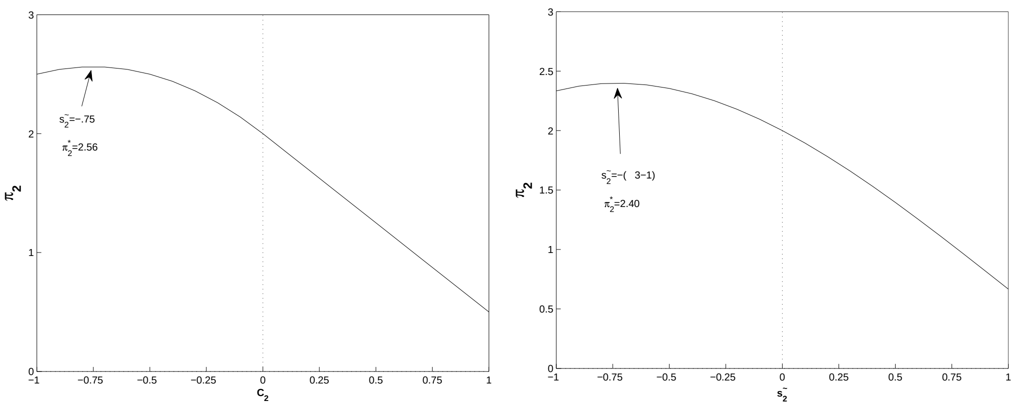

(a) $E\pi_2(\hat{s_2})$ when $\hat{s_1} = 0$

(b) $E\pi_2(\hat{s_2})$ when $\hat{s_1} \sim U[-1, 1$

Figure 14: $X = 1$

C.3 Best-Responding to Observed $\hat{s_1}$

Appendix C.2 argues that, even for a very noisy $\hat{s_1}$, player 2 can extract a great deal of information from observing player 1's action. Specifically, if $\hat{s_1}$ is commonly known to be uniformly distributed on the interval $[-1, 1]$, player 2's optimization problem changes very little. This section analyzes player 2's optimization problem when facing the *observed* $\hat{s_1}$.

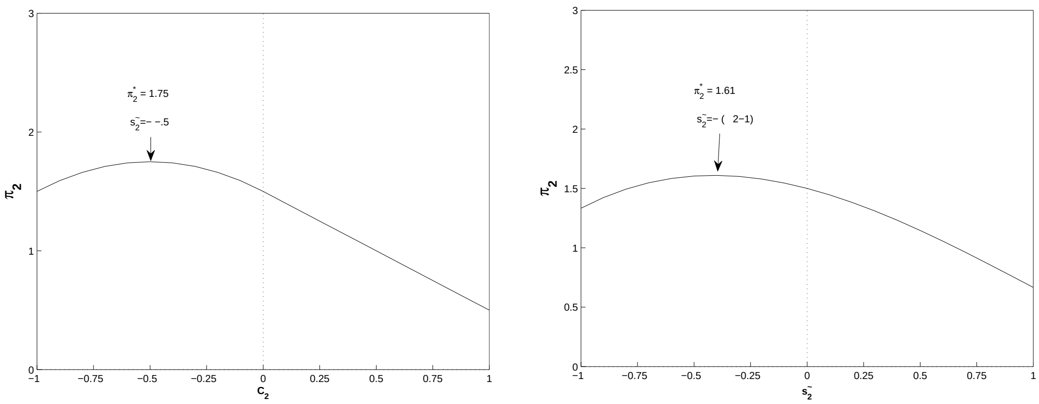

(a) $E\pi_2(\hat{s_2})$ when $\hat{s_1} = 0$

(b) $E\pi_2(\hat{s_2})$ when $\hat{s_1} \sim U[-1, 1$

Figure 15: $X = 0$

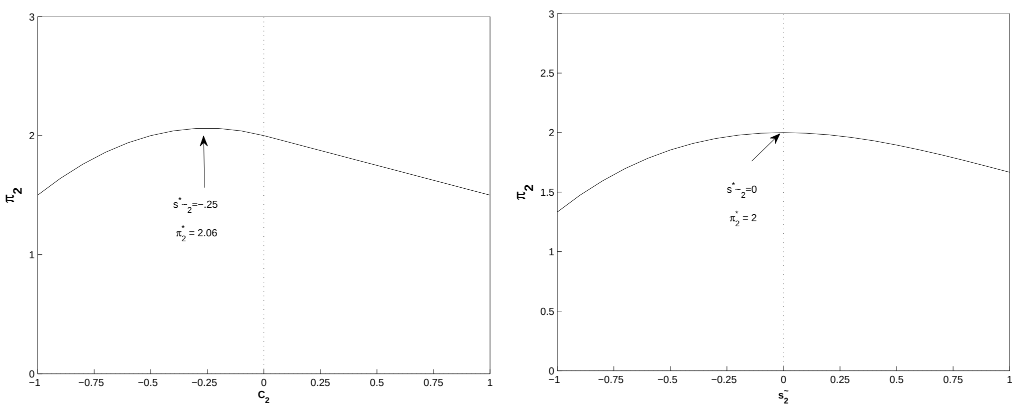

(a) $E\pi_2(\hat{s_2^*})$ when $\hat{s_1} = 0$

(b) $E\pi_2(\hat{s_2})$ when $\hat{s_1} \sim U[-1, 1$

Figure 16: $X = -1$

To achieve this goal, we consider $\hat{s_1}$ to be a random variable. $\hat{s_1}$ is assumed to take each value on $[-1, 1]$ with the same frequency with which it is observed in the data. We define $\hat{s_2}^o = (s_2^{Ao}, s_2^{Bo})$ as the strategy that maximizes player 2's payoff in response to the observed distribution of $\hat{s_1}$.

First, let $\hat{\mathbf{s_1}}$ be the $1 \times n$ vector of observed $\hat{s_1}$ for a each treatment, where n is the number of observations for $\hat{s_1}$ in that treatment. Then, allow $E\pi_2(\hat{\mathbf{s_1}}, \hat{s_2}|d_1)$ to be a vector of functions defining the expected profit for any $\hat{s_2}$ for each element of

$\hat{\mathbf{s_1}}$ and d_1. Recall that each element of $E\pi_2(\hat{\mathbf{s_1}}, \hat{s_2}|d_1)$ will be a two-part function, discontinuously differentiable at $\hat{s_2} = -\hat{s_1}$.

A complication arises from the fact that, for any $\hat{s_1} \neq 0$, choices A and B are no longer equally likely. Given that player 2 observes $d_1 = A$ or B before making her decision, we must allow her to update her priors. $\eta(\hat{\mathbf{s_1}}|d_1 = A) = \frac{1-\hat{\mathbf{s_1}}}{2}$ and $\eta(\hat{\mathbf{s_1}}|d_1 = B) = \frac{1+\hat{\mathbf{s_1}}}{2}$ are functions that weight each observed $\hat{s_1}$ according to its updated probability given $d_1 = A$ and $d_1 = B$, respectively. Then $\eta(\hat{\mathbf{s_1}}|d_1 = A) \cdot \pi_2(\hat{\mathbf{s_1}}, \hat{s_2}|d_1 = A)'$ is a function defining profit for any $\hat{s_2}$ after observing $d_1 = A$, and $\eta(\hat{\mathbf{s_1}}|d_1 = B) \cdot \pi_2(\hat{\mathbf{s_1}}, \hat{s_2}|d_1 = B)'$ defines the analogous expected payoff given $d_1 = B$. Maximizing this function for the observed $\hat{s_1}$ in each treatments yields the values in table 8. Note that we no longer assume the optimal strategies to be symmetric, as the observed $\hat{s_1}$ often differ from zero. As figures 7 and show, there is a positive bias in $\hat{s_1}$ under $X = -1$ and $X = 0$. In table 8, $\hat{s_2^*}$ denotes the

	Externality		
	$X = 1$	$X = 0$	$X = -1$
$\hat{s_2^*}$	$-.75$	$-.5$	$-.25$
s_2^{Ao}	-.789	-.560	-.153
s_2^{Bo}	.702	.399	.069

Table 8: Best-Responses to Observed $\hat{s_1}$

theoretical prediction for player 2's cutoff strategy $\hat{s_2}$. s_2^{Ao} and s_2^{Bo} are player 2's optimal responses to the observed $\hat{s_1}$ for $d_1 = A$ and B, respectively.

The negative externality treatment exhibits the greatest difference between the theoretically predicted optimal response, and the optimal response to the observed $\hat{s_1}$. The reason for this is two-fold. First, observed play by player 1 is the noisiest in this treatment, making the information held therein the least reliable. Second, player 2's payoff function is the flattest in this treatment to begin with, allowing smaller informational differences to alter the optimal cutoff more significantly.

Table 8 also shows that the optimal strategy is asymmetric. In the optimal strategy, $s_2^{Ao} < -s_2^{Bo}$. For $X = 1$ and $X = 0$, d_1 contains nearly as much information as predicted by theory, but is more revealing for $d_1 = A$. This is a result of the positive bias observed in $\hat{s_1}$. Table 9 explores whether subjects do, in fact, conform more when observing the more informative action $d_1 = A$.

Table 9 does not support the hypothesis that player 2 properly adjusts $\hat{s_2}$ for the positive bias in $\hat{s_1}$. In the positive externality treatment, $s_2^A \approx -s_2^B$. For negative and zero externality treatments, conformity is stronger for $d_1 = B$ than

	Treatment		
	$X = 1$	$X = 0$	$X = -1$
$\hat{s_2}$	.75	.5	.25
$\overline{s_2^A}$	-.623 (.408)	-.188 (.511)	.208 (.586)
$\overline{s_2^B}$	.608 (.471)	.324 (.434)	-.103 (.603)

Table 9: Actual Response by d_1

for $d_1 = A$. This coincides with a positive bias in s_2^A and s_2^B for these treatments. The best-response to the positive bias in $\hat{s_1}$ would be to decrease s_2^A, but subjects tend to *increase* s_2^A in these cases. If player 2's behavior is a reaction to a belief about noisy play on the part of player 1, this reaction is payoff-decreasing.

D Player 2 with Social Preferences

As player 2's behavior affects the probability that player 1 receives the externality X, a player 2 with social preferences must consider player 1's payoff in setting $\hat{s_2}$. We first discuss the predictions of Rawlsian and utilitarian preferences for each treatment using the framework of Charness and Rabin's (2002) model, then derive parameters necessary to describe the observed behavior.

D.1 The Model

Charness and Rabin's (2002) two-person model of social preferences allows for an agent to place a different weight the other's payoff based on whether the other's payoff is higher or lower than the agent's. Applied to this experiment, player 2's preferences are:

$$u_2(\pi_1, \pi_2) = \begin{cases} \rho\pi_1 + (1-\rho)\pi_2 & \text{if } \pi_2 \geq \pi_1 \\ \sigma\pi_1 + (1-\sigma)\pi_2 & \text{if } \pi_2 \leq \pi_1 \end{cases} \tag{45}$$

We will analyze predictions of two different social preferences, using the framework of this model. For simplicity, we will assume that player 2 makes his decision knowing that $\hat{s_1} = 0$[21]. This is a strong assumption based on the observed behavior of player 1, but not one with extreme behavioral implications for player 2.

Recall that for values of $s_2 \leq s_2^B$, $d_1 = d_2 = B$, causing both players to earn

[21] It is possible that player 1 has social preferences as well. Applying Charness and Rabin's model to our game, player 1 would rationally choose $\hat{s_1} = -1$ or 1 for highly negative σ. This corresponds to player 2 sabotaging her own decision in order to destroy the informational content passed to player 2

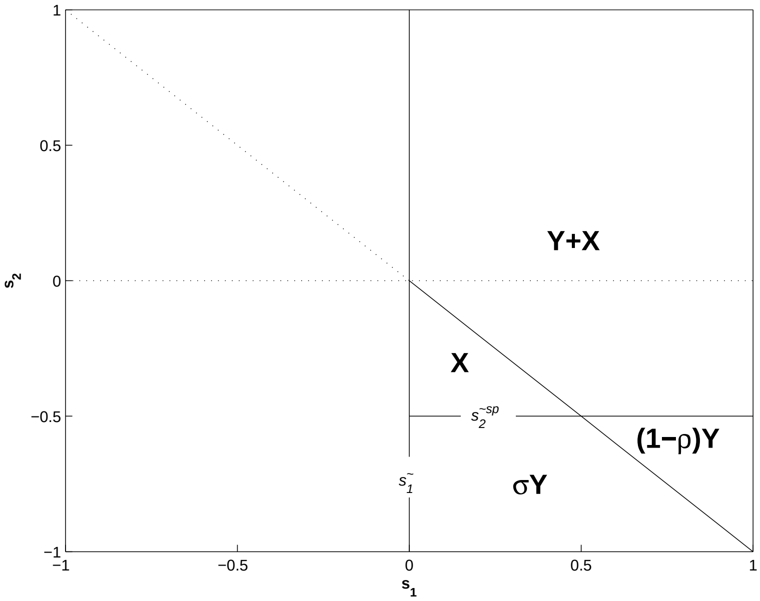

Figure 17: $u_2(\pi_1, \pi_2)$ for Charness-Rabin Social Preferences

the same payoff. Therefore, $u_2(\pi_1, \pi_2) = \pi_2$ for such values of (s_1, s_2). For values of $s_2 > \hat{s_2}$, players 1 and 2 receive different payoffs. If $\pi_2 = Y$, then $\pi_1 = 0$, and $u_2(\pi_1, \pi_2) = (1 - \rho)Y$. If $\pi_2 = 0$, then $\pi_1 = Y$, and $u_2(\pi_1, \pi_2) = \sigma Y$. Player 2's Charness-Rabin preferences are shown in figure 17. We now discuss the predictions of two specific types of preferences, using the framework of Charness and Rabin's model.

- **Rawlsian Preferences**: $\rho = 1, \sigma = 0$

An agent with Rawlsian or maximin preferences increases the lower of the two profits at any cost to the higher. Importantly, we assume that player 2 is Rawlsian in terms of outcomes, rather than in a probabilistic sense. In other words, a Rawlsian player 2 seeks to maximize the expectation of the minimum rather than the minimum of the expectations. Figure 17 guides intuition as to the optimal $\hat{s_2^r}$. Clearly, for $X > 0$, $\hat{s_2^r} = -1$, and for $X < -Y$, $\hat{s_2^r} = 1$.

For $-Y < X < 0$, Equating the beneficial and harmful tradeoffs of an increase in $\hat{s_2^r}$, $\hat{s_2^r} = -1 \times \frac{Y+X}{Y}$. For the negative externality treatment of this paper, that makes $\hat{s_2^r} = -.5$. Notice that the Rawlsian player 2 reacts to the negative externality *less* than the self-interested player. The cause of this reaction is that a Rawlsian feels significantly worse about $\pi_2 = Y$ than does a self-interested player 2. Therefore, she imposes a negative externality on both players 1 and 2 to increase the probability that *both* players receive Y, even though it decreases both profits in expectation. Rawlsian preferences perform moderately well in predicting behavior when $X = 1$, but quite poorly when $X = 0$ and $X = -1$.

π_2	π_1	$u_2^r(\pi_1, \pi_2)$	$u_2^u(\pi_1, \pi_2)$
$Y + X$	$Y + X$	$Y + X$	$Y + X$
Y	0	0	$\frac{1}{2}Y$
X	X	X	X
0	Y	0	$\frac{1}{2}Y$

Table 10: Rawlsian and Utilitarian Social Preferences

- **Utilitarian Preferences**: $\rho = \sigma = \frac{1}{2}$

An agent with utilitarian preferences maximizes the sum $\pi_1+\pi_2$. If $X \geq \frac{Y}{2}$, player 2 would always prefer both players to receive the externality than for herself to receive Y, so $\hat{s_2^u} = -1$. If $X < -\frac{Y}{2}$, $\hat{s_2^u} = 1$, as the sum of the negative externality to both players outweighs the payoff Y to player 2 herself. If $-\frac{Y}{2} < X < \frac{Y}{2}$, there is an interior solution, $\hat{s_2^u} = -\frac{(Y+2X)}{2Y}$.

- **Social Preferences Fitting the Results**

Rawlsian and Utilitarian fail to explain the results that we observe. We now apply Charness and Rabin's model to the mean action of player 2, $\overline{\hat{s_2}}$ in each treatment to account for this behavior. Player 2 behavior in the zero externality case cannot be explained by social preferences, as it remains an individual decision problem. Behavior in the positive and negative externality treatments are vulnerable to influence by social preferences. The following are the optimal cutoffs in the positive and negative externality treatments:

$$\begin{aligned} &X = 1 && X = -1 \\ &\hat{s_2^{sp}} = -\frac{(2\sigma - 3)}{2(\rho + \sigma - 2)} && \hat{s_2^{sp}} = -\frac{(2\sigma - 1)}{5(\rho + \sigma - 2)} \end{aligned} \tag{46}$$

Of course, as each equation has two unknowns, it is impossible to find a best-fit ρ and σ for the two treatments independently. However, using both treatments, we can solve the two simultaneously:

$$
\begin{aligned}
&X = 1 && X = -1 \\
&.615 = \frac{5(2\sigma - 3)}{\rho + \sigma - 2} && -.148 = \frac{5(2\sigma - 1)}{\rho + \sigma - 2} \qquad (47) \\
&\rho = 0 && \sigma = .69 \qquad (48)
\end{aligned}
$$

It is more commonly found that $\rho > \sigma$, as it is often found that subjects are more generous to those with less than them. Therefore, social preferences seem an unlikely explanation for player 2's behavior.

E Individual Decisions Over Time

Section 5 and details average behavior of subjects in each treatment. A different question altogether is the behavior of individual subjects. This section analyzes the behavior of individual subjects, both on average and the evolution of this behavior over time. Table E shows how the average behavior of subjects in the role of player 2 in each treatment evolves over time.

	Externality		
	$X = 1$	$X = 0$	$X = -1$
Obs. per Bin	75	75	90
Periods 1-5	-.39 (.58)	-.20 (.56)	.22 (.64)
Periods 6-10	-.57 (.45)	-.41 (.43)	.10 (.69)
Periods 11-15	-.62 (.39)	-.24 (.50)	.15 (.63)
Periods 16-20	-.68 (.41)	-.24 (.44)	.23 (.60)
Periods 21-25	-.71 (.34)	-.22 (.51)	.12 (.57)
Periods 26-30	-.73 (.36)	-.29 (.34)	.08 (.51)

Table 11: Average $\hat{s_2}$ Over Time

Table E shows that there is significant aggregated learning over time under the positive externality. $\hat{s_2}$ gradually approaches the theoretical prediction of -.75 by the final period. For the zero and negative externality treatments, there is no clear trend (although in both cases $\hat{s_2}$ is closer to the theoretical predictions in the final five rounds than in the first five). The variance does fall somewhat from the early to the later rounds, but not significantly.

With the exception of $X = 1$, there is little movement in average behavior over time. At the individual level, however, it is still desirable to isolate any predictive patterns in behavior over time. We now turn our attention to individual performance over the 30 periods. Figures 18, 19, 20 and 21 show selected examples of four different classes of behavior. They are not representative of the behavior or the entire pool of subjects, but rather examples of divergent behavioral patterns that are not manifest in the pooled data.

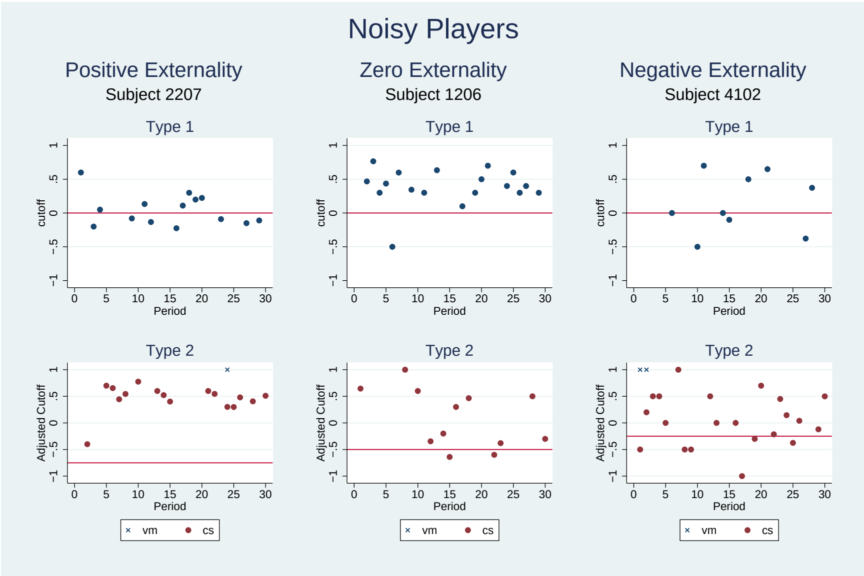

Figure 18: Players Whose Choices Appear Random

Figure 18 shows examples of players whose behavior is noisy. Even so, there are patterns in the behavior of the subjects displayed. Subject 2208's $\hat{s_2}$, although in flux, hovers around -.5. Subject 1206's $\hat{s_1}$ is noisy, but consistently positive[22]. Even for subjects whose behavior is the noisiest, there are distinguishable patterns.

As seen in figure 19, some subjects follow theory very closely. Subject 1204 chose the exact theoretical prediction in every period. The behavior of subjects 2209 and 4104 began somewhat noisily, but was consistently in line with theory by the end of the treatment.

This paper finds that subjects are reasonably sensitive to externalities and insensitive to information relative to theory. Figure 20 shows examples of players in each treatment that take this conclusion to the extreme. These subjects react absolutely to externalities. In their absence, they follow their own private information. They do not react at all to information contained in d_1. Subject 2110 follows this behavior by conforming absolutely to the action he observes in the presence of a positive externality. In the case of subject 1213, there is no externality. Therefore, the subject sets $\hat{s_2}$ each time, neglecting to use the information

[22] Recall that $\hat{s_1}$ was, on average, positive when $X = 0$.

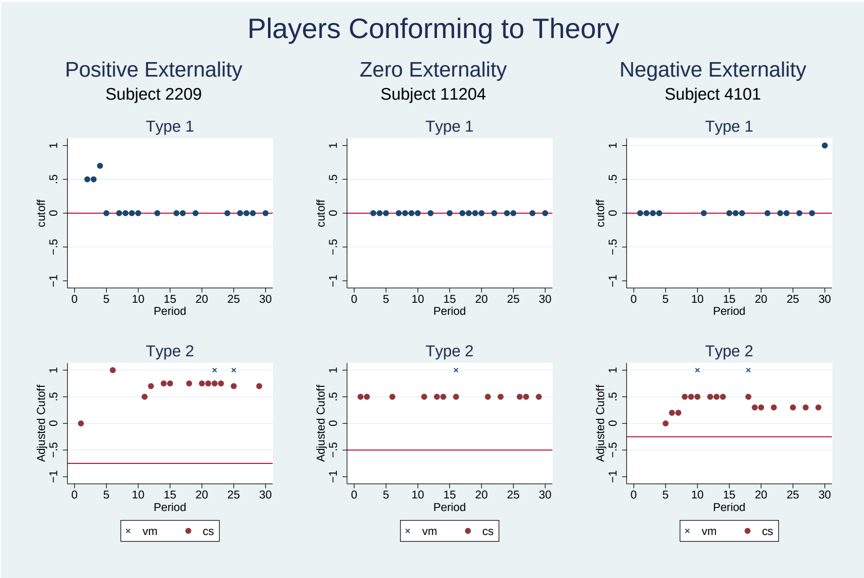

Figure 19: Players Behaving According to Theory

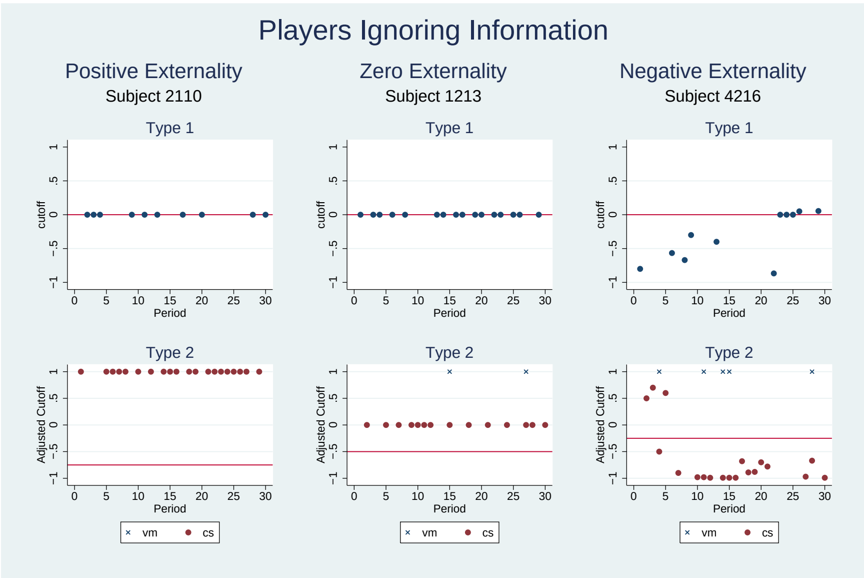

Figure 20: Players Ignoring Information

inherent in player 1's action. Subject 4216 behaves noisily in both roles in the beginning of the experiment, then settles into the strategy of following the theoretical prediction in the role of player 1 ($\hat{s}_1 = 0$) and following a (near) negative cascade in the role of type 2 ($\hat{s}_2 \approx 1$).

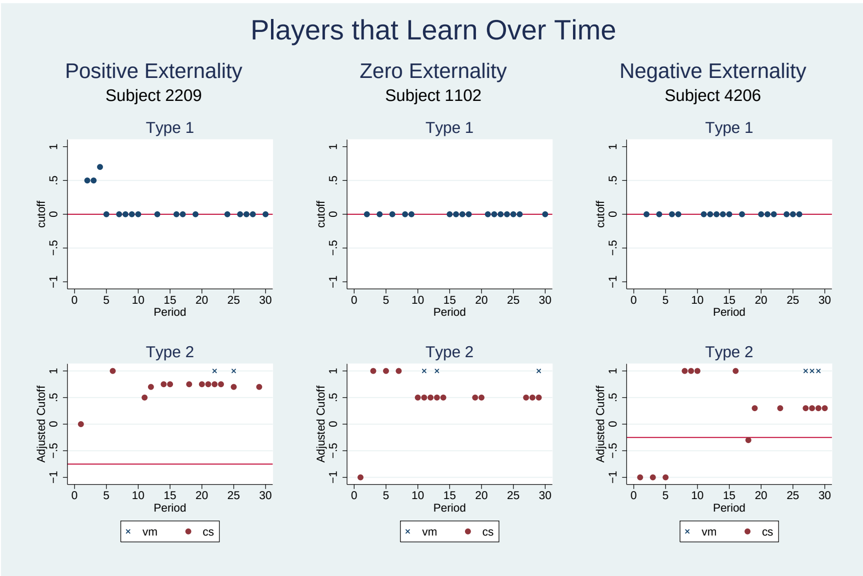

Figure 21: Players that Learn over Time