Conservative Reporting and Cooperation

Paul Fischer
University of Pennsylvania

Hong Qu
Penn State University

Preliminary and Incomplete – Comments Welcome

Abstract

We hypothesize that a conservative reporting system, relative to a full information system, can serve to reduce the likelihood of inefficient uncooperative run-like behavior by pooling a state that a susceptible to run-like behavior with a state that is not susceptible to such behavior. Theoretical analysis predicts that conservative reporting is more likely to be effective in thwarting inefficient runs when the likelihood of the state not subject to run-like behavior is high relative to the state that is subject to run-like behavior. We conduct a number of preliminary experiments to assess the validity of that prediction and find that conservative reporting can reduce inefficient runs, but not as effectively as the standard theoretical analysis would suggest.

I. Introduction

Conservative reporting practices, which force firms with poor prospects to reveal their status in a timely fashion while making it difficult for firms with good prospects to credibly reveal their status in a timely fashion, have been present since the inception of formal accounting systems. The enduring value of conservative reporting has been attributed, among other things, to efficiency gains arising from more effective debt contracting, more efficient employee incentive contracts, and more effective tax planning.[1] We consider another potential source of value by considering whether conservative reporting practices can facilitate cooperation between firm stakeholders.

Failures of cooperation arise when all stakeholders would be better off if they could commit to take a particular individual action, but that action is not individually optimal for any one of them.[2] The classic game theoretic framework illustrating such a failure is the well-known two-person prisoner's dilemma game. In that game, each player decides whether to cooperate or cheat. The payoffs are structured so that the players are better off if they both cooperate than if they both cheat, but each player's dominant strategy is to cheat. As a consequence, cheating is the only Nash equilibrium even though both players would be better off if they could commit to cooperate.

For an example of how a failure to cooperate can lead to efficiency losses for a firm, consider an investment firm that raises capital from multiple investors to acquire a portfolio containing some illiquid assets. Furthermore, assume the investors have a right to withdraw funds at some interim date. If too many investors withdraw early, however,

---

1 See, for example, the discussion in Watts (2003).

2 We have used the term "cooperation" in a manner similar to that employed in Palfrey and Rosenthal (1994).

then the most illiquid assets must be sold at a "fire sale" price, which causes all investors to receive less than they otherwise would receive. As a consequence, investors are best off if they collectively maintain their investments as long as the portfolio management is deemed to be adequate. If investor beliefs regarding the portfolio management fall from good to just adequate, however, each investor might individually have an incentive to deviate from the strategy that maximizes their collective expected cash flows. As a consequence, they engage in a run on the firm and all withdraw early, which will lead them all to be worse off.

As another example, consider a start-up firm engaged in a multistage rollout that is financed by multiple investors. If all investors committed to provide subsequent stages of financing as long as some modest objectives have been satisfied, they would achieve higher expected cash flows. If the objectives are subsequently just barely satisfied, however, it can be individually optimal for each investor to cease to invest. As a consequence, the investors end up "pulling the plug" too early on the firm.

As a final example, consider a firm who relies upon and owes money to multiple suppliers. If the firm is potentially distressed, the firm's suppliers are all better off if they collectively continue to provide the firm with supplies because doing so allows the firm to maintain operations, thereby increasing the expected cash flows that each supplier will ultimately receive. An individual supplier, however, might have incentives to defect from such an agreement because, holding the behavior of the other suppliers constant, demanding immediate payment prior to making future deliveries generates the greatest expected cash flows for that supplier. As a consequence the only outcome is one in which the firm's suppliers force the firm out of business prematurely.

Of course, sometimes a firm's performance and prospects are so poor that it should be shut down to avoid further deterioration of its assets, while at other times a firm's performance and prospects are so good that the stakeholders' individual incentives are consistent with cooperation. Hence, only moderate firms are subject to the possibility of a cooperation failure. Our insight is that, when the possibility of uncooperative run-like behavior is possible, a conservative information system that just identifies poor firms (i.e., a firm is either designated as poor or not poor) can be a superior to a more revealing full information system (i.e., firm type is perfectly revealed) because moderate firms are no longer subject to inefficient runs. We present a model and conduct an experiment based upon that model to ascertain if and when a conservative reporting system can alleviate losses associated with inefficient runs.

We employ a three-stage model to illustrate our idea and frame the experiment. In the first stage, two risk averse investors fund a project. In the second stage, some information is received and each investor simultaneously decides whether to withdraw from the project or hold their initial investment. The project payoffs, if at least one party holds, are realized and paid in the third stage.

There are three possible project states that are knowable in the second stage, high, medium, and low. If the investors learn that the state is low in the second stage, the Pareto preferred outcome and the unique Nash equilibrium in the game is for both investors to withdraw. If the state is known to be high in the second stage, the Pareto preferred outcome and the unique Nash equilibrium in the game is for both investors to hold. Hence, the low state reflects a case in which a firm's prospects are so poor it is best to shut down the firm and the high state reflects a case in which a firm's prospects are so good that investors do not have incentives to engage in early withdrawal. If the state is known to be medium, however, a failure of cooperation arises because the Pareto preferred outcome is for both investors to hold but the unique Nash equilibrium is for both investors to withdraw. Within the context of this model, we compare a full information system (i.e., one that reveals the state) with an incomplete or conservative information system in which investors learn only whether the state is low or not.

Within the context of our setting, the conservative information system can outperform the full information system as long as the conservative reporting system thwarts withdrawals in both the medium and high states. With conservative accounting, however, withdrawals may occur for the firm in both the medium and high states for two reasons. First, if the probability of the firm being in the high state contingent upon the firm not being in the low state is small, the medium state's dominant strategy of withdrawing continues to be a dominant strategy.

Second, holding may fail to be sustained as an equilibrium because of a coordination failure, which can arise when holding is an optimal individual strategy for an investor if and only if the other investor holds. In this case, there are two Nash equilibrium, one in which both investors hold and one in which both investors withdraw, and the coordination failure arises when both investors withdraw because they believe the other will do so. The coordination failure can occur when the possibility that the state is high conditional upon the state being not low is large enough make cooperation possible (i.e., each investor optimally holds because they believe the other will do so), but not so high as to rule out the possibility of premature withdrawals as well (i.e., each investor optimally withdraws because they believe the other will do so). Conservative accounting can thwart premature withdrawals arising from a coordination failure if the preferred equilibrium, which is for both investors to hold, is sufficiently focal.

We conduct a number of laboratory experiments to ascertain whether and, more importantly, when a conservative reporting system can alleviate inefficient run-like behavior. We find that, when the state is revealed to be not low in a conservative reporting system, run-like behavior is reduced relative to the behavior exhibited when the state is known to be medium, but not necessarily eliminated except in cases where the probability of the high state conditional upon the state not being low is very high. In fact, conservative reporting fails to eliminate inefficient withdrawals even when theory suggests hold/hold is a unique Nash equilibrium when the state is revealed to be not low. As a consequence, a significant downside of a conservative reporting system is that inefficient runs can occur for both medium and high states, as opposed to just the medium state. Given this tension, our experimental results suggest the probability of the state being high conditional on the state not being low must be quite large for expected investor payoffs to be enhanced by conservative reporting.

Our analysis is related to three streams of literature, the literature regarding accounting conservatism, the literature pertaining to mechanisms supporting cooperative behavior in prisoner's dilemma games such as ours, and the literature pertaining to inefficient runs. We employ a formal definition of conservatism in our model that is consistent with the definitions employed in models by Guay and Verrecchia (2007), Goex and Wagenhofer (2009), and Caskey and Hughes (2012). More specifically, a conservative accounting regime is one in which bad outcomes are revealed in a more timely manner than good outcomes, which is consistent with the idea that a conservative reporting regime forces the bad news out earlier while delaying recognition of good news.[3] Unlike these papers, we focus on the role conservatism can play in mitigating cooperation failures as opposed to making debt contracts more efficient as in Goex and Wagenhofer (2009) and Caskey and Hughes (2012), or supplementing good news voluntary disclosures to enhance the information environment as in Guay and Verrecchia (2007).

The type of setting considered in our study, which is a prisoner's dilemma game, has been studied extensively, both theoretically and experimentally. Early experiments establish that players do sometimes cooperate in Prisoner's dilemma games, which suggests that the Nash criterion applied to the payoff structures alone is insufficient to completely explain actual behaviors. Our analysis, like a significant stream of the literature, focuses on identifying factors that can potentially increase the extent of cooperation, as opposed to explaining why cooperation arises at all. Factors considered in the literature include, but are not limited to, the extent of repeated interactions, allowing communication prior to play, and strategies, such as tit for tat, that foster cooperation by an opponent. In contrast to the literature, we discuss how a conservative information system can help induce cooperative behavior when such behavior cannot be induced explicitly via enforceable contracts.

Given our focus on how alterations in the information environment can thwart run-like behavior faced by a firm, our analysis is also related to the significant literature focusing on runs on a firm, which are modeled as coordination failures as opposed to a lack of

---

3 The formal definition of conservatism we employ, however, differs somewhat from the formal definition employed in studies such as Chen et al. (2007) and Gigler et al. (2009). In those studies, a conservative accounting regime is one in which a favorable accounting signal is more precise than an unfavorable accounting signal because a conservative accounting regime places a greater burden of proof before favorable news can be recognized.

cooperation. Of these studies, the focus in Bigus (2010) is the same as ours in that Bigus (2010) analytically considers the role conservatism can play in mitigating an inefficient run. In contrast to our model, more conservative reporting in Bigus (2010) causes a bad financial report to be a noisier signal of a firm's prospects, which reduces the likelihood of a run when a bad report is released.[4] Because runs are always inefficient in Bigus (2010), conservative accounting as defined in his model can be optimal. Our study is also related to the experimental studies pertaining to runs, which include Madies (2006), Garret and Keister (2009), Schotter and Yorulmazer (2009)), and Anctil et al. (2010). The most relevant of these studies is Anctil et al. (2010), which focuses on the role information can play in implementing a desired set of equilibrium actions when coordination failures can occur. Unlike our study, Anctil et al. (2010) consider a coordination game in which players have private information about the state (i.e., a global game). In their setting finer private information leads to improved coordination only when it makes the efficient equilibrium behavior risk dominant in expected payoffs in an "otherwise identical game" involving common information.

II. Model

Consider a three stage investment game in which a firm is funded in the first stage by two risk averse investors with common utility function U, where $U' > 0$ and $U''(W) \leq 0$. The capital invested is employed to undertake a long-term project yielding an uncertain terminal gross return in the third stage. In the second stage, the investors receive a common signal about the terminal gross return and then decide whether to withdraw or hold their initial investment, which affects the realized terminal gross return to the project.

[4] This definition of conservatism is similar to that employed in Chen et al. (2007) and Gigler et al. (2009).

If both investors withdraw, the firm is liquidated and the game ends. If at least one investor holds in the second stage, the terminal gross return is realized in the third stage and divided equally among investors who held.

To be consistent with our experimental design, we assume there are three possible return states for the firm's project, referred to as low, medium, and high, although the insights derived can be extended to a setting with any number of states. Let L denote the low state, M denote the medium state, and H denote the high state. The prior probability of state L is $\lambda \in (0,1)$. The prior probability of state H is $(1-\lambda)\mu$, where $\mu \in (0,1)$, and the prior probability of state M is naturally $(1-\lambda)(1-\mu)$. Hence, μ denotes the probability of state H conditional upon the state not being state L, which is convenient notation for the purposes of our analysis.

Formally, the payoffs in state S to the two investors are represented via the following table, where the first payoff in each cell corresponds to the investor choosing the action in the leftmost column, the second payoff corresponds to the investor choosing the action in the top row, and the subscript S represents a payoff that is contingent upon the state:

	Withdraw	Hold
Withdraw	w,w	ω,η_S
Hold	η_S,ω	h_S,h_S

To capture the notion the firm has illiquid assets or that withdrawal of financing leads to costly project disruptions, we assume that $\omega > \eta_S$ for all S, which implies that, if one investor withdraws and the other holds, the investor who withdraws receives more than the investor who holds. In addition, we assume $\omega > w$, which implies that an investor who withdraws receives more if the other investor holds as opposed to withdraws, which is also consistent with the notion that the firm's assets are illiquid or that withdrawal of financing leads to costly project disruptions. In addition, we assume that payoffs to holding are highest in the high state and lowest in the low state so: $h_H > h_M > h_L$ and $\eta_H > \eta_M > \eta_L$.

We make specific assumptions regarding the payoffs in each state so that it is best for both investors to withdraw in the low state, and best for both investors to hold in the medium and high states because doing so maximizes their joint payoffs. In addition, the payoffs are structured to make the best decision the unique equilibrium in the low and high states, but not the medium state. In the medium state, the payoffs are structure so the worst decision is implemented in the sense that the joint payoffs are minimized in equilibrium. Hence, if the investors know that the medium state has been realized, they will engage in an inefficient run on the firm. We then consider how altering the information system might allow the investors to avoid an inefficient run in the medium state, while still facilitating the best decisions in the low and high states.

The specific formal assumptions are as follows:

(A1) $h_H > \omega$ and $\eta_H > w$

(A2) $h_M < \omega$ and $\eta_M < w$ and $h_M > w$

(A3) $h_L < \omega$ and $\eta_L < w$ and $h_L < w$

Assumption (A1) implies that the combined payoff in the high state is greatest if both investors hold, and that holding is a best response to the other investor holding or withdrawing. Assumption (A2) implies that the combined payoff of both investors in the medium state is greatest if both hold, and that withdrawing is the best response to the other investor holding or withdrawing. Assumption (A3) implies that the combined payoff of both investors in the low state is greatest if both withdraw, and that withdrawing is the best response to the other investor holding or withdrawing.

A. Full Information

To establish a benchmark for our analysis, we determine state contingent Nash equilibrium when the state is revealed to the investors. One might interpret this information system as a "fair value" system because both good and bad news about the future state are revealed publicly in the interim stage when investors make withdrawal decisions. As we alluded to previously, the medium state is the only one in which the best equilibrium outcome fails to be played.

Observation 1. *Assume the investors know the state at the second stage. Given assumptions (A1), (A2), and (A3), the unique Nash equilibrium for each is as follows:*

State	*Strategy Maximizing Combined Payoffs*	*Unique Equilibrium Strategies:*
Low	*Both investors withdraw.*	*Both investors withdraw.*
Medium	*Both investors hold.*	*Both investors withdraw.*
High	*Both investors hold.*	*Both investors hold.*

As an aside, note that the equilibrium is unique for any state because the payoffs are structured so that withdrawing in the low and medium states are dominant strategies and holding in the high state is a dominant strategy. If the payoffs are modestly altered, say, for the medium state, it is possible for there to be two pure strategy equilibrium – one in which both investors withdraw and one in which both investors hold – and for there to be mixed strategy equilibrium. For the purposes of our illustration, we have avoided such payoff structures to sustain a clear theoretical outcome in the full information case. When we consider a conservative information environment below, multiple equilibrium issues naturally arise.

B. Conservative Information

As noted in the introduction, a hallmark of conservative accounting is that potential loss realizations are more likely to be recognized, whereas potential gain realizations are less likely to be recognized. Hence, investors are more likely to know when the firm is more likely to have poor cash flow realizations and less likely to know when the firm will have excellent cash flow realizations. Within the context of our illustrative game, we mimic a conservative accounting system by imposing an information system that only distinguishes whether or not the realized state is low or is not low. Hence, it is an information system imposing pooling at the top.

When investors learn that the state is low, the unique Nash equilibrium is identical to that in the full information setting – both investors choose to withdraw and these strategies maximize their combined payoffs. The conservative information system, however, leads to pooling at the top (i.e., the M and H states are pooled), which leaves open the potential for a variety of Nash equilibria when investors learn that the state is not low (i.e., either M or H). If the equilibrium implemented has the property that both investors hold when they learn the state is not low, then the conservative information system clearly outperforms the full information system. On the other hand, if both investors withdraw in equilibrium when investors learn that the state is not low, then the full information system outperforms the conservative information system.

The equilibria that are sustained when there is pooling at the top are a function of the model parameters and, in our experiments, we primarily focus on the parameter μ, which is the probability of high state conditional upon the state not being low. A higher value for μ implies that the high state is more likely given that the state is not low, which should have the effect of making it more likely that investors will choose to hold given that holding is a dominant strategy. In contrast, a lower value for μ implies that the medium state is more likely given that the state is not low, which should make it more likely that investors choose to withdraw given that withdrawing is the dominant strategy in the medium state.

We confirm the above intuition by characterizing the equilibrium played as a function of the value of μ. Before doing so, it is convenient to define the following notation: $r_h = \frac{U(\omega)-U(h_M)}{U(h_H)-U(h_M)}$, $r_\eta = \frac{U(w)-U(\eta_M)}{U(\eta_H)-U(\eta_M)}$, and $p = \frac{1}{\frac{(U(h_H)-U(h_M))(r_h-\mu)}{(U(\eta_H)-U(\eta_M))(\mu-r_\eta)}+1}$. With this notation in hand, the formal characterization is as follows:

Observation 2. *Assume a conservative information system. When investors learn that the state is the low state, the unique Nash equilibrium is for both to withdraw. When investors learn that the state is not the low state and $r_h \leq r_\eta$, the Nash equilibria for each value of μ are characterized as follows:*

μ Range	*Nash Equilibria*
$\mu \geq r_\eta$	*Unique equilibrium: both investors hold*
$r_\eta \geq \mu \geq r_h$	*Equilibrium 1: both investors hold* *Equilibrium 2: both investors withdraw* *Equilibrium 3: both investors hold with probability p*
$r_h \geq \mu$	*Unique equilibrium: both investors withdraw*

When investors learn that the state is not the low state and $r_\eta \leq r_h$, the Nash equilibria for each value of μ are characterized as follows:

μ Range	*Nash Equilibria*
$\mu \geq r_h$	*Unique equilibrium: both investors hold*
$r_h \geq \mu \geq r_\eta$	*Equilibrium 1: one investor holds and one withdraws* *Equilibrium 2: both investors hold with probability p*
$r_\eta \geq \mu$	*Unique equilibrium: both investors withdraw*

Observation 2 provides some guidance for our experimental parameter choices. For experimental purposes we restrict attention to parameterizations in which $r_h \leq r_\eta$ under an assumption of risk neutrality. Within that setting, we consider cases in which μ is so large that both investors holding is likely to be a unique Nash equilibrium. We also consider cases in which μ is sufficiently large so that both investors holding satisfies the conditions for a Nash equilibrium, but not so large that both investors withdrawing is also a Nash equilibrium. Hence, the investors are subject to a coordination game and there is a potential for a coordination failure in which investors fail to implement the Pareto superior equilibrium. We do not consider settings in which $r_h \geq r_\eta$ because there is no intermediate range for μ such that both investors holding is one of multiple equilibria when the state is not low. That is, when $r_h \geq r_\eta$, holding arises as the unique equilibrium outcome or holding is not an equilibrium outcome.

We extend our analysis to theoretically refine our thinking for how changes in μ might influence the equilibrium selection when μ is in the range of values where there are multiple equilibria, one of which is both investors holding. Prior research involving coordination games suggests that a risk dominance criteria has substantial predictive power for determining behavior and that, in the class of two-person games such as ours, many equilibrium selection criteria select the risk dominant equilibrium.[5] Within the context of our model, the risk dominance criteria can be intuitively tied to a simple thought experiment. Assume an investor believes the other investor plays each strategy with equal probability. If the investor's unique best response to these beliefs is to hold, then hold is

[5] For a discussion of the predictive power of the risk dominance criteria see Schmidt et al. 2003).

the risk dominant equilibrium and if the unique best response is to withdraw, then withdraw is the risk dominant equilibrium.[6] As formalized in the following observation, when μ is in the range of values so that the investors are engaged in a coordination game, the risk dominant equilibrium is hold/hold when μ is sufficiently large and is withdraw/withdraw when μ is not sufficiently large.

Observation 3. *Assume a conservative reporting system and that* $r_h \leq \mu \leq r_\eta$. *When investors learn that the state is not low, the risk dominant equilibrium is hold/hold if* $\mu > \gamma r_h + (1-\gamma) r_\eta$ *where* $\gamma = \frac{\mathrm{U}(\mathrm{h_H})-\mathrm{U}(\mathrm{h_M})}{\mathrm{U}(\mathrm{h_H})-\mathrm{U}(\mathrm{h_M})+\mathrm{U}(\eta_\mathrm{H})-\mathrm{U}(\eta_\mathrm{M})}$ *and is withdraw/withdraw if* $\mu < \gamma r_h + (1-\gamma) r_\eta$.

Observation 3 suggests that, when the conservative reporting system subjects investors to potential coordination failures, a sufficiently high value for μ is expected to induce investors to play the Pareto superior hold/hold equilibrium.

III. (Preliminary) Experiment

Our experiments stem directly from the previously described model. We first examine behaviors in a full information benchmark setting and then consider behaviors in the less informative conservative information system. Our theoretical analysis suggests that transitioning from the full information system to the conservative reporting system should facilitate a hold/hold equilibrium in the medium and high state as long as μ is sufficiently large to make hold/hold the equilibrium that is implemented when the state is revealed to be not low. Of course, as noted above, the theoretical equilibrium construct we employ, the

---

6 While our intuitive thought experiment characterizes the risk dominant equilibrium in our setting, we should note that the risk dominance criteria, as developed by Harsanyi and Selten (1988), is based upon a more nuanced thought experiment that can be applied to a much larger class of games.

Nash criterion, has been shown to not perfectly capture individual behavior. In particular, in prisoner's dilemma games, individuals sometimes cooperate when doing so is not Nash equilibrium behavior, and, in coordination games, individuals do not necessarily select the Pareto dominant equilibrium.

A. Experiment Payoffs and Probabilities

In all of our experiments, the payoff parameters are characterized for each of the three states in Table 1. Assuming the players know the state, the unique Nash equilibrium for each state is highlighted in bold. Note that these payoffs are consistent with the underlying assumptions in the analytical modeling. Furthermore, the payoffs imply that $r_\eta > r_h$ a risk neutral utility function, $r_\eta = 0.545 > 0.2 = r_h$. Furthermore, the relationship derived in the risk neutral utility case holds for the negative exponential (i.e., CARA) utility function as well.

Table 1: State Contingent Payoff Table

State L			*State M*			*State H*		
	Withdraw	Hold		Withdraw	Hold		Withdraw	Hold
Withdraw	**600,600**	1000, 200	Withdraw	**600,600**	1000, 360	Withdraw	600,600	1000,800
Hold	200,1000	500,500	Hold	360,1000	900,900	Hold	800,1000	**1400,1400**

There are two information conditions: Full Information and Conservative Information. We adopt a between-subject design. Participants only play in one condition. In the Full Information condition, they know exactly which state occurs and they make decisions for each state. In the Conservative Information condition, they receive only partial information about the state, specifically, they are informed whether the low state occurs or not. In the

Conservative Information condition, we vary μ, the probability of state H conditional on the state is not L. We conduct four sessions with μ that equals 20%, 40%, 60%, and 80% respectively. We conduct five experimental sessions in total. Table 2 summarizes our experimental sessions.

Table 2: Experimental Session Summary

Sessions	*# of Subjects*
Full Information	20
Conservative Information (μ=20%)	22
Conservative Information (μ=40%)	22
Conservative Information (μ=60%)	24
Conservative Information (μ=80%)	22

In the Full Information condition, out analysis above predicts that both subjects, which we will refer to as investors, will withdraw in the low, L, and medium, M, states and that both will hold in the high, H, state. In the case of the medium state, the predicted equilibrium actions are inferior to both investors holding, while the predicted actions in the low and high states are not Pareto dominated.

In the Conservative Information condition, investors only know whether the state is L or not L. When they know that state is L, the strictly dominant strategy is the same as in the Full Information setting, which is for both subjects to withdraw. When the state is not L, investors are not sure whether the state is M or H. As outlined in our theoretical analysis, the equilibrium actions that will be induced hinges critically on the conditional probability that the state is H, μ.

Assuming investors are risk neutral, the expected payoff with different μ chosen can be summarized in Table 3 in which any Nash equilibrium is highlighted in bold. We discuss how risk aversion might influence our experimental findings in the next subsection.

Table 3: Expected Payoff Given State is Not L

	μ=20%		*μ=40%*		*μ=60%*		*μ=80%*	
	Withdraw	Hold	Withdraw	Hold	Withdraw	Hold	Withdraw	Hold
Withdraw	**600,600**	1000, 448	**600,600**	1000, 536	600,600	1000, 200	600,600	1000, 712
Hold	448,1000	**1000,1000**	536,1000	***1100,1100***	624,1000	**1200,1200**	712,1000	**1300,1300**

When μ is 20% and investors are risk neutral, withdraw/withdraw and hold/hold are both Nash equilibrium, with hold/hold being the Pareto dominant equilibrium. Withdraw/withdraw is the risk dominant equilibrium, however. When μ is 40% and investors are risk neutral, we again have a setting in which withdraw/withdraw and hold/hold are both Nash equilibrium, with hold/hold being the Pareto dominant equilibrium. In this case, the risk dominant equilibrium is also hold/hold. Finally, when investors are risk neutral and μ is 60% or 80%, hold/hold is a unique Nash equilibrium.

B. Risk Aversion

Given the possibility that investors will behave in a risk averse manner, it is worthwhile to provide some insight how the payoff matrices are affected by risk aversion. To do so, we consider the negative exponential (CARA) utility function, $-(1/\rho)\exp(-\rho w)$ where w is wealth, and we assess how changes in the coefficient of risk aversion, ρ, influence the payoff matrices and associated equilibria. If an investor is certain of the other investor's strategy (i.e., there is no strategic risk), the investor faces uncertainty only if they choose to hold. In particular, the certainty equivalent payoffs for holding when investors are risk averse fall below the expected payoffs. Hence, a risk averse subject will be less incline to hold relative to the risk neutral subject.

In Table 4, we identify a range for the coefficient of risk aversion, ρ, such that hold/hold is an equilibrium, a risk dominant equilibrium, and a unique equilibrium for investors for each parameterization for μ. We consider risk dominance in addition to existence and uniqueness because past literature has found the risk dominance criterion explains behaviors in coordination games similar to the coordination game that can arise with conservative accounting.[7] Finally, note that the wealth levels employed to compute the risk aversion range reflect wealth levels after converting from experimental dollars to US dollars (i.e., if the experimental payoff is 1,400 experimental dollars, the wealth employed for the computation is 1,400/12,000).

Table 4: Risk Aversion Effects

	range for ρ, where $\rho > 0$, such that hold/hold is		
μ	an equilibrium	risk dominant equilibrium	unique pure strategy equilibrium
20%	none	none	none
40%	$\rho < 52.7191$	$\rho < 4.73838$	none
60%	$\rho < 107.969$	$\rho < 30.4622$	$\rho < 12.1306$
80%	$\rho < 192.978$	$\rho < 69.6056$	$\rho < 64.5405$

Because hold/hold is a knife-edged equilibrium when investors are risk neutral and μ = 20% (i.e., each investor is indifferent between withdrawing and holding if the other investor holds), hold/hold is never an equilibrium when investors are risk averse and μ = 20%. When μ = 40%, hold/hold is a risk dominant equilibrium even for a fairly high level of risk aversion. In contrast, if when μ = 60% hold/hold is risk dominant for very high levels

[7] See, for example, Schmidt et al. (2003).

of risk aversion and unique for quite high levels of risk aversion. Finally, when μ = 80%, it is hard to see how withdraw/withdraw would even be an equilibrium for any plausible level of risk aversion.

C. Procedures

We conducted five experimental sessions. A total of 110 subjects participated in the experiment. The majority of subjects were undergraduate students and the rest are graduate students at a university. We conduct one experimental session for each experimental condition. One session is Full Information treatment. Four other sessions are Conservative Information treatment with varying μ, the probability of the H state conditional on the state not L.

In an experimental session, dividers separated participants. The experiments were conducted using Z-tree software. Participants read the computerized instructions and completed quizzes testing their understanding of the experimental instructions. They were required to answer all questions correctly before they could move on to the experimental task. The computer randomly assigned participants into two-person groups. There were 10 periods in each experimental session. Participants were re-grouped in each period in a turnpike design. At the end of the experiment, one of the ten periods was randomly selected to determine their cash payoff. The payoff in the experiment was measured in experimental dollars, which was converted to US dollars at the rate of 12,000 =1 US dollar. Participants also received a $5 show up fee. The average cash payment was $11.09. The experiment lasted from 45 minutes to 1 hour.

In the Full Information treatment, the computer screen showed the payoff table in each state, and participants were asked to input a decision for each state. At the end of each period, they received feedback information about the other subject's decision for each state and their own payoff. In the Conservative Information treatment, the computer screen showed the payoff table in each state, but participants were asked to make two decisions: one decision when state is L and the other decision when state is not L. At the end of each period, they received feedback information about the other subject's decisions and their own payoff. When state is not L, the computer randomly selected 10 states based on μ. For example, if μ=0.2, computer draws state H with probability of 20% and state L with probability of 80%. Then the computer calculated the average payoff over the 10 randomly selected scenarios, which, in theory, should diminish any risk aversion effects associated with state uncertainty that we discussed earlier.[8]

IV. Experimental Results

This section summarizes our (preliminary) experimental results. We analyze subjects' strategies under the Full Information and Conservative Information conditions and compare subjects' welfare between the two conditions.

A. Full Information Condition

In the Full Information Condition, we observe the convergence toward the predicted equilibrium in each state. Figure 1 plots the frequency of hold in each state. When state is H, all subjects choose to hold by period 2. In state H, it is a strictly dominant strategy for both subjects to hold. The experimental evidence suggests that it is easy for subjects to learn to play the dominant strategy. Similarly, in state L, it is a strictly dominant strategy for both subjects to withdraw. All subjects choose to withdraw by period 3. In state M, however, the

[8] Risk aversion also plays an important role for strategic risk, which will arise in cases where the game becomes a coordination game. More risk averse people are generally also more averse to strategic risk (see Heinemann et al. 2009).

dominant strategy is to withdraw, but the combined payoffs of both subjects are highest if both subjects hold. All subjects choose to withdraw by period 6. Initially, 35% of subjects choose to hold. This is consistent with prior experimental findings for the Prisoner's Dilemma game (e.g., Andreoni and Miller 1993).

Figure 1: Frequency of Hold in Full Information Condition

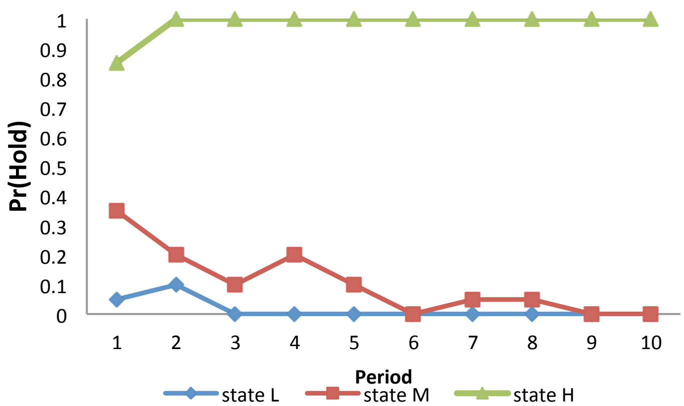

B. Conservative Information Condition

In the Conservative Information Condition, subjects make two decisions: one is for state is L and the other for state is not L. In the four Conservative Information treatments, we vary μ, the probability of H state conditional on state is not L. First, we examine their strategy for state L and state not L. Figure 2 plots the frequency of Hold given state is L or not L in the four treatments. When state is L, the dominant strategy is for both subjects to withdraw and we observe the convergence toward this predicted equilibrium in all four treatments. When state is not L, the equilibrium in the expected game varies across the four treatments with different μ. When μ is 0.2, the probability of holding decreases over time. By the last period, the probability of hold drops to 0.32. Theoretically, we predict that both withdraw and both hold are pure strategy equilibria, but both withdraw is risk dominant, assuming all subjects are risk neutral. When μ is 0.4, the probability of hold stays around 0.64. Theoretically, both withdraw and both hold are also pure strategy equilibria, but both hold is risk dominant. When μ is 0.6 or 0.8, the theoretical prediction is that both hold is the unique equilibrium. The experimental data shows convergence toward the equilibrium when μ is 0.8. When μ is 0.6, the probability of hold stays around 0.83 and does not converge to 1.

Figure 2: Probability of Hold in the Conservative Information Conditions

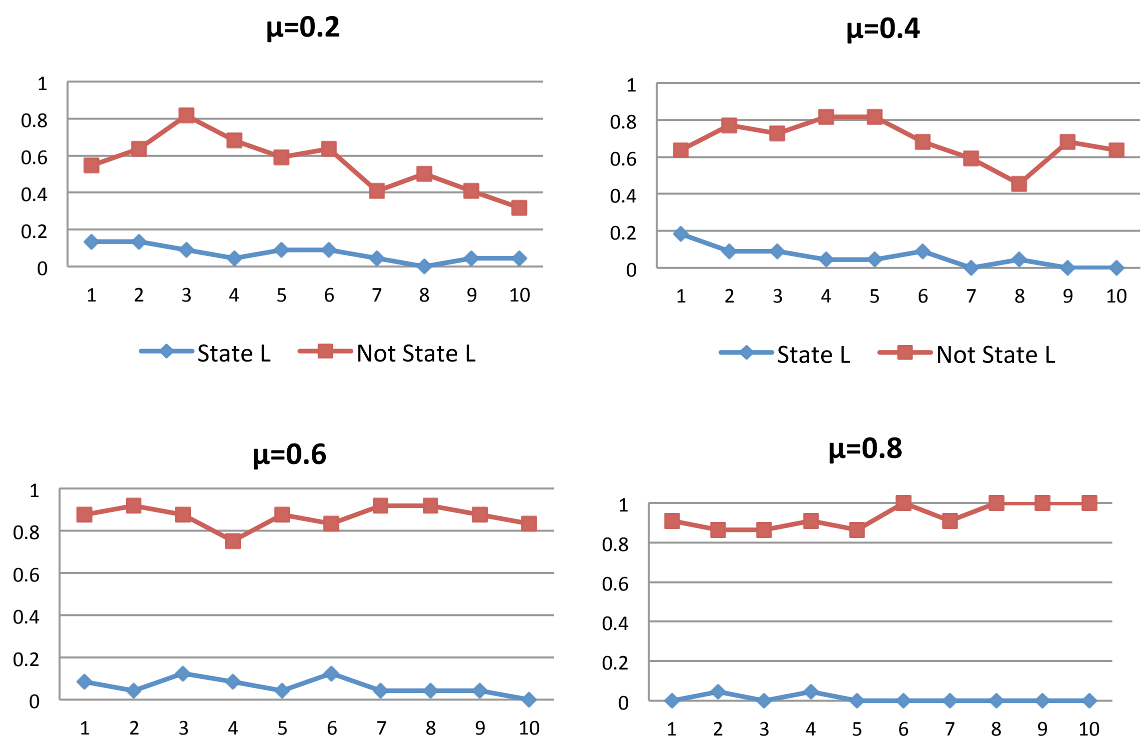

Table 5 reports the frequency of Hold in the four Conservative Information conditions. We report the frequency of Hold in the first period, the last period and the average over all periods.

Table 5: Frequency of Hold Given State is not L

μ	*First Period*	*Average*	*Last Period*	*Risk Neutral Prediction*
0.2	0.55	0.43	0.32	Withdraw Risk Dominant
0.4	0.64	0.64	0.64	Hold Risk Dominant
0.6	0.88	0.85	0.83	Hold Unique
0.8	0.91	0.95	1.00	Hold Unique

When μ = 20%, we find that some players still play hold even though it is unlikely to satisfy the risk dominance criterion, or even be an equilibrium with a minute level of risk aversion. This finding, coupled with the deterioration in the likelihood subjects hold, is consistent with prior experimental results and suggests that at least some subjects have a natural propensity to cooperate that is eroded over time.

When μ = 40%, as predicted, subjects are more incline to hold. On the other hand, given that hold/hold is risk dominant equilibrium for plausible levels of risk aversion (i.e., ρ < 4.7), it is perhaps a bit surprising that subjects do not hold more frequently. This finding is not in line with what prior literature led us to predict. An element of our model not incorporated in prior studies, however, is the fact that the there is uncertainty associated with the (seemingly) risk dominant equilibrium that is not present in the dominated equilibrium. This uncertainty, coupled with the inherent strategic uncertainty (i.e., the uncertainty about the other player's strategic choice) may be sufficient to undermine the dominance criterion.

When μ = 60%, the experimental results again are consistent the prediction that hold/hold is more likely to be played. As in the case when μ = 40%, however, subjects choose hold somewhat less frequently than one might have expected given that hold/hold is the unique equilibrium for risk neutral investors, and even investors who have a coefficient of risk aversion up to $\rho = 12.1$. In particular, we expected hold to be selected somewhat more than 83% of the time. Finally, when $\mu = 80\%$, subjects converged to hold/hold equilibrium 100% of the time, which the theoretical framework predicts given that hold/hold is the unique equilibrium for risk neutral investors, and for investors with a coefficient of risk aversion up to $\rho = 65.4$.

In summary, our experimental results suggest that conservative reporting does promote a greater likelihood that subjects hold relative to subject behavior when the medium state is realized. The extent to which they hold, however, is not overwhelming. In particular, when conservative accounting is associated with a subsequent coordination game, which is when $\mu = 20\%$ or $\mu = 40\%$, subjects do not naturally converge to the Pareto dominant equilibrium even when it is risk dominant. Hence, converting the prisoner's dilemma problem to a coordination problem is insufficient to promote the hold/hold equilibrium in almost all instances. Furthermore, even when conservative accounting appears to induce hold/hold as a unique Nash equilibrium, subjects do not appear to converge to that equilibrium en mass unless the probability of the high state conditional upon the report of not low is quite large. Taken together, these findings suggest that the state uncertainty associated with conservative accounting interacts in unique way with the strategic uncertainty inherent to the two-person investment game.

C. Reporting Efficiency

To assess the relative efficiency of the conservative reporting and full information reporting systems we must establish metrics for making a comparison. We consider two forms of metrics. First, we focus solely on the probabilities that the outcomes fall into the three possible categories: both hold, both withdraw, and one holds and one withdraws. While this approach ignores the game specific payoff specific parameters, it has the benefit of focusing solely on the behaviors induced under the general assumptions that the payoffs happen to satisfy. Second, we consider the expected average payoffs, which incorporates the game specific payoffs as well.

To focus solely on the probabilities of outcomes, we calculate the proportion of outcomes in the last period of each treatment that fall into each possible outcome: both hold, both withdraw, and one holds and one withdraws. Results are reported in Table 6.

Table 6: Probability of game outcomes when State is not L

μ	*Both Withdraw*	*One Withdraw/One Hold*	*Both Hold*
0.2	0.45	0.45	0.09
0.4	0.09	0.55	0.36
0.6	0.00	0.33	0.67
0.8	0.00	0.00	1.00

As summarized in Table 6, the probability of both subjects holding increases in μ. When μ is 0.2, only 1 out of 11 groups has both subjects hold. When μ is 0.8, all subjects hold and the outcome is both hold in all groups. In addition, the probability of both investors withdrawing decreases as μ increases. When μ is 0.2, 5 out of 11 groups have both subjects withdraw. When μ is 0.6 and 0.8, the probability of both investors withdrawing drops to 0. The probability of outcomes in which one withdraw and one hold increases when μ increases from 0.2 to 0.4, but decreases when μ increases to 0.6 and 0.8. The results for 0.4 are suggestive of the subjects having difficulty managing the coordination game because hold/hold is not a sufficiently compelling equilibrium. The results for 0.6, where a third of the subjects play a hold/withdraw outcome, are perhaps the most puzzling because hold/hold is a unique Nash equilibrium even for relatively high levels of risk aversion.

To further evaluate whether the pooling inherent to the conservative information system increases or decreases efficiency, we compare subjects' average expected payoffs for the first and last period in the Full Information treatment and the Conservative Information treatment. In state L, all subjects choose to withdraw in both Full Information and Conservative Information treatment so we only compare their payoffs when the state is not L. When the state is not L, we calculate the average expected payoff in the Full Information treatment by first computing each subject pair's expected payoff by weighting the payoffs for each subject pair's choices in the M and H state by $(1-\mu)$ and μ respectively. We then compute the average payoffs across all subject pairs. In the Conservative Information treatment, subjects choose only one action for both M and H state. To compute the expected payoff, we again weight the subject payoffs in M and H state by $(1-\mu)$ and μ respectively for the single action choice. We then compute the average expected payoff across the subjects. Table 7 reports the average expected payoff per subject in the Conservative Information condition compared to the payoff in the Full Information condition when the state is not L.

A key parameter that affects subjects' payoffs and the relative payoff difference between the Full and Conservative Information treatments is μ. As μ increases, the expected payoff in both conditions increases because the probability of H state increases and the probability of M state decreases. In the Conservative Information condition, the potential benefit of pooling state M and H is to change the behaviors in the M state from both withdraw to both hold, which increases each subject's state M payoff by 900 – 600 =

300. Such benefits only obtain in state M, however, which implies that as the probability of state H conditional upon the state not being L, μ, increases, the expected size of the potential benefit decreases. On the other hand, because increases in μ increase the probability that the investors will choose to hold when they learn the state is not L, subjects are more likely to obtain the benefit from pooling state M and H when μ increases. When this latter effect dominates the former effect of increasing μ, the average expected payoffs to the subjects rises.

Table 7: Average Expected Payoff per Subject when State is not L

Panel A: Period 1

μ	*Conservative Information*	*Full Information*	*Diff*
0.2	791	786	5
0.4	888	902	-14
0.6	1118	1018	100
0.8	1236	1134	102

Panel B: Period 10

μ	*Conservative Information*	*Full Information*	*Diff*
0.2	693	760	-67
0.4	873	920	-47
0.6	1071	1080	-9
0.8	1300	1240	60

When μ is 0.8, all subjects choose to hold in the Conservative information setting, which causes the latter effect to dominate the former effect and their average expected payoff increases. In both period 1 and period 10, the average payoff is higher in the Conservative Information treatment than the Full Information treatment.

At the other extreme when μ is 0.2, subjects are less willing to hold as the probability of the H (M) state conditional upon the state being not low is very low (high). In this case, the high probability of M state causes subjects to withdraw, which leads to inefficient outcomes for state H. As a consequence, subjects' payoffs are lower in the Conservative Information condition compared to the Full Information condition in period 10.

In the cases where μ has an intermediate value, μ is 0.6 or 0.4, subjects' payoffs in the Conservative Information condition are not significantly different from the Full Information condition in period 10. That is, any increase in the probability that investors choose to hold fails to overcome the fact that the expected benefit from holding is smaller.

V. Conclusion

When investors receive more information they are generally expected to make better investment decisions as individuals (i.e., holding the decisions of other investors constant). When investors can engage in suboptimal runs on a firm, however, better individual decision-making need not correspond to better investment decisions for all investors collectively. As a consequence, conservative reporting, which provides less than full information, might be a vehicle for thwarting inefficient runs.

For our illustration we employ a two-investor, three-state investment game in which investors have the option to withdraw early or hold given an interim report, and the returns received by an investor are affected not only by their own withdrawal decision, but also by the withdrawal decision of the other investor. We parameterize the model so that, when the state is revealed perfectly prior to the withdrawal decision, the Nash equilibrium is for both investors to hold in the high state and withdraw in the medium and low states. In the medium state, however, both investors would be better off if that could commit to hold. Hence, the medium state is associated with inefficient runs. Conservative reporting, which is defined in our setting as a reporting system that clearly identifies the state is low or note, can be superior to a full information reporting system because the resulting "pooling at the top" of the state distribution can, for some parameterizations of the model, thwart inefficient runs (i.e., inefficient withdrawal behavior) that would otherwise occur in the medium state.

We conduct a laboratory experiment to ascertain whether behaviors predicted by the theoretical model prevail. We find that, when the state is revealed to be not low in a conservative reporting system, run-like behavior is reduced relative to the behavior exhibited when the state is known to be medium, but not necessarily eliminated. In fact, conservative reporting fails to eliminate inefficient withdrawals even when theory suggests hold/hold is a unique Nash equilibrium when the state is revealed to be not low. As a consequence, a significant downside of a conservative reporting system is that inefficient runs can occur for both medium and high states, as opposed to just the medium state. Given this tension, our experimental results suggest the probability of the state being high conditional on the state not being low must be quite large for expected investor payoffs to be enhanced by conservative reporting.

References

Anctil, R. M., Dickhaut, J., Johnson, C., & Kanodia, C. (2010). Does information transparency decrease coordination failure?. Games and Economic Behavior, 70(2), 228-241.

Andreoni, J., & Miller, J. H. (1993). Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence. The Economic Journal, 103(418), 570-585.

Bigus, Jochen, Accounting Conservatism and Creditor Conflicts (March 1, 2010). Available at SSRN: http://ssrn.com/abstract=1629503 or http://dx.doi.org/10.2139/ssrn.1629503

Chen, Q., Hemmer, T., & Zhang, Y. (2007). On the relation between conservatism in accounting standards and incentives for earnings management. Journal of Accounting research, 45(3), 541-565.

Garratt, R., & Keister, T. (2009). Bank runs as coordination failures: An experimental study. Journal of Economic Behavior & Organization, 71(2), 300-317.

Gigler, F., Kanodia, C., Sapra, H., & Venugopalan, R. (2009). Accounting conservatism and the efficiency of debt contracts. Journal of Accounting Research, 47(3), 767-797.

Göx, R. F., & Wagenhofer, A. (2009). Optimal impairment rules. Journal of Accounting and Economics, 48(1), 2-16.

Guay, Wayne R. and Verrecchia, Robert E., Conservative Disclosure (June 2007). Available at SSRN: http://ssrn.com/abstract=995562 or http://dx.doi.org/10.2139/ssrn.995562

Harsanyi, J.C., Selten, R., 1988. A General Theory of Equilibrium Selection in Games. Massachusetts Institute of Technology Press, Cambridge, MA.

Heinemann, F., Nagel, R., & Ockenfels, P. (2009). Measuring strategic uncertainty in coordination games. The Review of Economic Studies, 76(1), 181-221.

Judson Caskey and John S. Hughes (2012) Assessing the Impact of Alternative Fair Value Measures on the Efficiency of Project Selection and Continuation. The Accounting Review: March 2012, Vol. 87, No. 2, pp. 483-512

Madies, P. (2006). An Experimental Exploration of Self-Fulfilling Banking Panics: Their Occurrence, Persistence, and Prevention. The Journal of Business, 79(4), 1831-1866.

Palfrey, T. R., & Rosenthal, H. (1994). Repeated play, cooperation and coordination: An experimental study. The Review of Economic Studies, 61(3), 545-565.

Schotter, A., & Yorulmazer, T. (2009). On the dynamics and severity of bank runs: An experimental study. Journal of Financial Intermediation, 18(2), 217-241.

Schmidt, D., Shupp, R., Walker, J. M., & Ostrom, E. (2003). Playing safe in coordination games:: the roles of risk dominance, payoff dominance, and history of play. Games and Economic Behavior, 42(2), 281-299.

Watts, R. 2003. Conservatism in accounting, Part I: Explanations and implications. Accounting Horizons 17 (3): 207–221.

Appendix: Experimental Instructions

General Instruction

This is an experiment in decision making under uncertainty. If you follow instructions closely and make decision wisely, you will receive a considerate amount of money. We pay you in cash at the end of the experiment. You will need to sign a compensation receipt before you receive your payment.

Your earnings in the experiment are measured by the experimental dollars. From this point forward all units of account will be in experimental dollars. At the end of the experiment, experimental dollars will be converted to U.S. dollars at the rate of 12,000 experimental dollars for every 1 U.S. dollar. This experiment has multiple rounds and one round is randomly selected to determine your cash payment.

If you have any questions during the experiment, please raise your hand and wait for an experimenter to come to you. Please do not talk, exclaim, or try to communicate with other participants during the experiment. Participants who intentionally violate this rule will be asked to leave without any payment.

Experimental Details

In this experiment, you will be randomly grouped with one of the other participants in this room. You and the other participant play the role of creditors who have invested in a project. Each of you will decide whether to (1) withdraw his or her money before the project is complete, or (2) hold your investment in the project until the project is complete.

Your Payoff

Your payoffs depend on your decision, the other creditor's decision and the project return.

- If both of you withdraw, the project will be liquidated immediately and each creditor gets $600.
- If both of you hold, the project will continue and each of you receives R upon project completion.
- If you withdraw and the other creditor holds, you get $1000 immediately. The project is partially liquidated. The other creditor receives r upon project completion, where r is smaller than R.
- If you hold but the other creditor withdraws, the other creditor gets $1000 immediately and you receive r upon project completion.

Your payoff and the other subject's payoff can be summarized in a table as shown below.

Other's decision

		Withdraw	**Hold**
Your decision	**Withdraw**	600,600	1000, r
	Hold	r, 1000	R, R

How to read the payoff table? First, find the row corresponding to your decision. Second, find the column corresponding to the other investor's decision. In the intersection cell, the first number is your payoff and the second number is the other's payoff. For example, If both of you withdraw, you can find your payoff in the intersection of the second row and second column. Your payoff is 600 and the other creditor's payoff is 600. If you withdraw and the other investor holds (the second row and third column), your payoff is 1000 and the other's payoff is r. If you hold and the other investor withdraws (the third row and second column), your payoff is r and the other's payoff is 1000. If both of you hold (the third row and third column), each investor receives R.

State of Nature

The amount R and r are determined by the state of nature. There are three possible states: state 1, state 2, state 3.

In state 1, r=200and R=500

In state 2, r=360 and R=900

In state 3, r=800 and R=1400

Below is a summary of payoffs for the three possible states.

	State 1			*State 2*			*State 3*	
	Withdraw	Hold		Withdraw	Hold		Withdraw	Hold
Withdraw	600,600	1000, 200	Withdraw	600,600	1000, 360	Withdraw	600,600	1000, 800
Hold	200,1000	500,500	Hold	360,1000	900,900	Hold	800,1000	1400,1400

Public Information about State (μ=0.8)

In this experiment, a public information system reveals whether the realized state is state 1. If the report is “not state 1”, there is a 20% chance that state 2 occurs and 80% chance that state 3 occurs. Your task is to choose whether to hold or withdraw when the public report is “state 1” or “not state 1”.

Decision Execution

Your decision will be executed in the following manner in each period. The computer randomly matches you with one of the other participants in this room. The computer will calculate your payoffs in the following way. Your payoff is calculated based on the payoff table in state 1 if the public report is “state 1”. If the public report is “not state 1”, the computer randomly selects 10 scenarios, specifically, state 2 is drawn with 20% chance and state 3 is drawn with 80% chance. The computer calculates your payoffs in each scenario based on your decision and the other matched participant's decision. Then it calculates the average payoff for the 10 scenarios, which determines your payoff.

Experimental procedures

The experiment has multiple periods. In each period, you submit your decisions and the computer executes your decisions as described above. The computer will randomly group you with another participant in the next period. At the end of each period, you will receive a report summarizing your payoffs when report is "state 1" and the 10 scenarios drawn by the computer when report is "not state 1".

Decision Screen

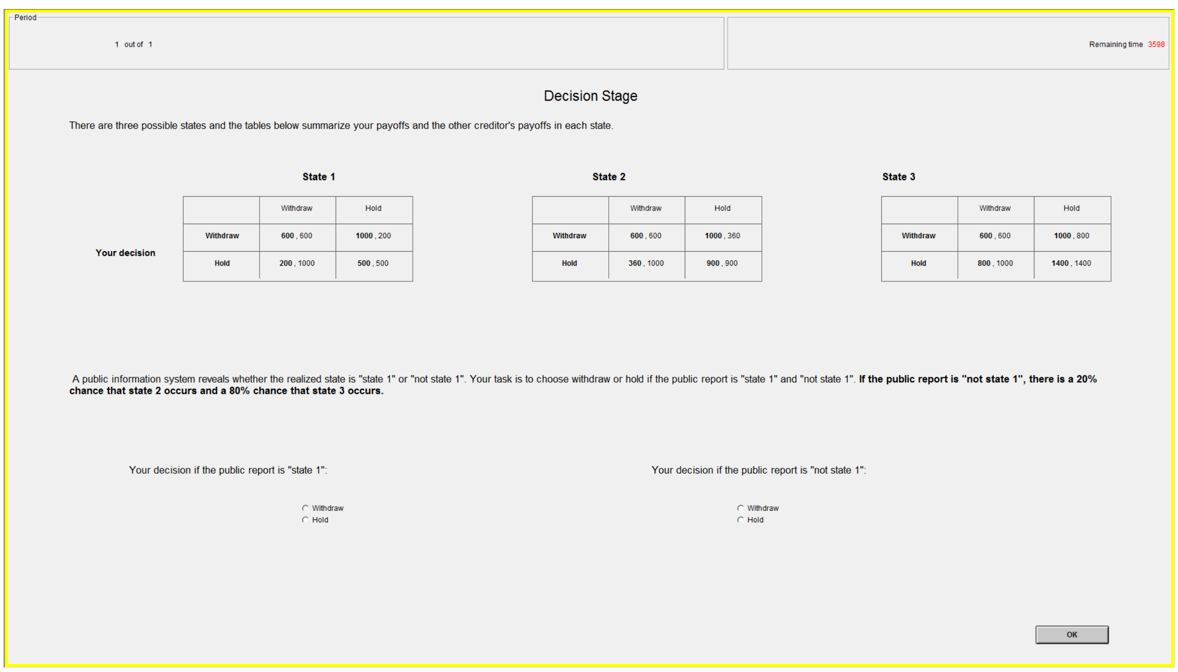

Feedback Screen

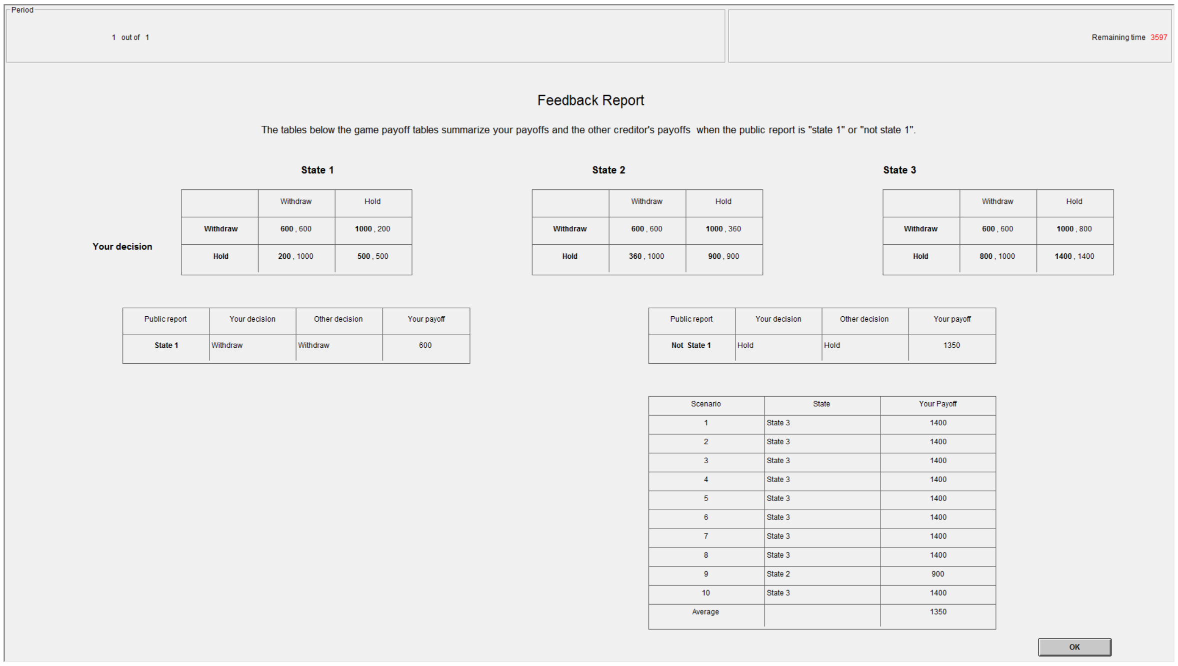

Period

1 out of 1

Remaining time 3597

Feedback Report

The tables below the game payoff tables summarize your payoffs and the other creditor's payoffs when the public report is "state 1" or "not state 1".

Your decision

State 1

	Withdraw	Hold
Withdraw	**600** , 600	**1000** , 200
Hold	**200** , 1000	**500** , 500

State 2

	Withdraw	Hold
Withdraw	**600** , 600	**1000** , 360
Hold	**360** , 1000	**900** , 900

State 3

	Withdraw	Hold
Withdraw	**600** , 600	**1000** , 800
Hold	**800** , 1000	**1400** , 1400

Public report	Your decision	Other decision	Your payoff
State 1	Withdraw	Withdraw	600

Public report	Your decision	Other decision	Your payoff
Not State 1	Hold	Hold	1350

Scenario	State	Your Payoff
1	State 3	1400
2	State 3	1400
3	State 3	1400
4	State 3	1400
5	State 3	1400
6	State 3	1400
7	State 3	1400
8	State 3	1400
9	State 2	900
10	State 3	1400
Average		1350

OK

Proofs of Observations

Proof of Observation 1

Consider first the low state. Withdrawing is a dominant strategy because of assumption (A3). Hence, the unique equilibrium is for both investors to withdraw. Consider next the medium state. Withdrawing is a dominant strategy because of assumption (A2). Hence, the unique equilibrium is for both investors to withdraw. Consider finally the high state. Holding is a dominant strategy because of assumption (A1). Hence, the unique equilibrium is for both investors to hold.

Proof of Observation 2

If $\mu \geq r_h$, holding is the best response to holding so both investors holding is a Nash equilibrium. If $r_\eta \geq \mu$, withdrawing is the best response to withdrawing so both investors withdrawing is a Nash equilibrium. When $\mu \leq r_h$ and $\mu \geq r_\eta$, holding is the best response to withdrawing and withdrawing is the best response to holding so one investor holding and the other withdrawing is a Nash equilibrium. The case in which $(\mu - r_h)(r_\eta - \mu) \geq 0$ derserves a bit more direct analysis to complete the proof. An investor will play a mixed strategy if any only if the investor is indifferent to the outcome attained from either always holding or always withdrawing. If one investor believes the other will hold with probability $p \in [0,1]$, that investor will be indifferent between holding and withdrawing if and only if

$$p\left(\mu U(h_H) + (1-\mu)U(h_M)\right) + (1-p)\left(\mu U(\eta_H) + (1-\mu)U(\eta_M)\right) = pU(\omega) + (1-p)U(w). \text{(E1)}$$

Rearranging (E1) yields

$$p = \frac{\left(\mu U(\eta_H) + (1-\mu)U(\eta_M) - U(w)\right)}{\left(U(\omega) - \mu U(h_H) - (1-\mu)U(h_M)\right) + \left(\mu U(\eta_H) + (1-\mu)U(\eta_M) - U(w)\right)}$$

$$= \quad \frac{(\mu U(\eta_H) - U(\eta_M))(\mu - r_\eta)}{(U(h_H) - U(h_M))(r_h - \mu) + (\mu U(\eta_H) - U(\eta_M))(\mu - r_\eta)}$$

$$= \quad \frac{1}{\frac{(U(h_H) - U(h_M))(\mu - r_h)}{(U(\eta_H) - U(\eta_M))(r_\eta - \mu)} + 1}. \qquad \text{(E2)}$$

The requirement that $p \in [0,1]$ implies that $\frac{(U(h_H) - U(h_M))(r_h - \mu)}{(U(\eta_H) - U(\eta_M))(\mu - r_\eta)} \geq 0$, which can only be true if $\frac{(\mu - r_h)}{(r_\eta - \mu)} \geq 0$ or $(\mu - r_h)(r_\eta - \mu) \geq 0$. Hence, if $(\mu - r_h)(r_\eta - \mu) \geq 0$, there exists a $p \in (0,1)$ such that one investor holding with probability p is a best response to the other investor holding with probability p.

Proof of Observation 3

If one investor believes the other will hold with probability p = .5, that investor will hold if and only if

$$.5\left(\mu U(h_H) + (1 - \mu)U(h_M)\right) + .5\left(\mu U(\eta_H) + (1 - \mu)U(\eta_M)\right) > .5U(\omega) + (1 - .5)U(w). \qquad \text{(E3)}$$

The condition for holding follows directly from rearranging (E3). Similarly, the investor will withdraw if and only if the inequality sign in (E3) flips. The condition for withdrawing follows directly from rearranging the resulting equation.