// NO.19-032 | 08/2019

DISCUSSION PAPER

// BERNHARD GANGLMAIR, ALEX HOLCOMB AND NOAH MYUNG

Expectations of Reciprocity when Competitors Share Information: Experimental Evidence

Expectations of Reciprocity when Competitors Share Information: Experimental Evidence*

Bernhard Ganglmair† Alex Holcomb‡ Noah Myung§

August 16, 2019

Abstract

Informal exchange of information among competitors has been well-documented in a variety of industries, and one's expectation of reciprocity shown to be a key determinant. We use an indeterminate horizon centipede game to establish a *feedback loop* in the laboratory and show that an individual's beliefs about the recipient's intentions to reciprocate matter more than a recipient's ability to do so. This implies that reducing strategic uncertainty about a competitor's behavior has a stronger effect on information flows than reducing environmental uncertainty (about the competitor's ability). We further show results on the formation of beliefs and discuss managerial implications.

Keywords: knowledge diffusion; information sharing; reciprocity; conversation; experimental economics; centipede game

JEL Codes: O33, D8, C72, C91

*We thank Lea Cassar, Carolin Häussler, Ernan Haruvy, Holger Herz, Kyle Hyndman, Wolfgang Luhan, Chris Malloy, Asaf Manela, Michael Rebello, Giorgo Sertsios, Sophie Shive, Jeremy Stein, Alexander Wagner, Joël van der Weele, and seminar and conference participants at the City University London, George Mason University, University of Maryland, University of Texas at Dallas, University of Texas at El Paso, CRC TR 224 - EPoS, the Dublin Colloquium of Behavioral and Experimental Finance 2014, ESA 2014, the Research in Behavioral Finance Conference 2016 (Amsterdam) Society of Experimental Finance (2014, Zurich; 2016, Tucson), and the Virginia Experimental Social Science Conference 2016 for comments and suggestions. Funding by the Center and Laboratory for Behavioral Operations and Economics (CLBOE) at the University of Texas at Dallas and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B2) is gratefully acknowledged. Internal Review Board approval has been obtained at the University of Texas at Dallas (MR 14-127). Declarations of interest: none.

†ZEW Mannheim, MaCCI, and University of Mannheim. Address: ZEW – Leibniz Centre for European Economic Research, L7, 1, 68161 Mannheim, Germany; E-mail: b.ganglmair@gmail.com; Twitter: @ganglmair.

‡Appalachian State University, Department of Finance, Banking and Insurance, 416 Howard Street, Peacock Hall, Boone, NC 28608, USA; E-mail: holcombaj2@appstate.edu.

§University of Virginia, Frank Batten School of Leadership and Public Policy, Charlottesville, VA 22904, USA; E-mail: noah.myung@virginia.edu.

1 Introduction

Economists have for a long time studied the market for ideas and the functioning of the resultant information exchange (e.g., Arrow, 1962). Of particular interest to the economists are the often puzzling cases where competitors share information that is of strategic value to them. Numerous empirical studies provide evidence of such knowledge exchanges in a variety of competitive settings.[1] These studies suggest that an individual's expectation of reciprocity is a motivation to share information, arguing that "potential future reciprocity is weighed against the current loss of competitiveness" (Häussler et al., 2014). Individuals are willing to incur the potential costs of sharing valuable information if they expect to receive something of similar value in return. This form of sharing can lead to a type of *feedback loop*.[2] In this paper, we model information exchange as an indeterminate horizon centipede game and investigate how expectations of reciprocity affect the trading of ideas in a controlled laboratory environment. We argue that a player must have both the *ability* and the *intention* to reciprocate, and we ask how a player's expectations regarding the abilities and intentions of others differentially affect the incentive to share information.

Our formal representation of a feedback loop is based on work by Stein (2008) who develops a model of word-of-mouth communication to characterize the conditions under which competitors have an incentive to share private information. Information is generated through an exogenous random process, and communication occurs via an escalating feedback loop with potential payoffs increasing in the number of interactions. Communication breaks if a player either conceals new information or has no new information to share. Players face a simple tradeoff: concealing new information gives them a relative information advantage that translates into a payoff advantage, but it also breaks the feedback loop for future sharing, and they forego the chance to generate more information which increases absolute payoffs even when there is no disparity in those payoffs. Stein (2008) shows that if the probability

[1]These settings include the steel minimills industry (von Hippel, 1987; Schrader, 1991), the semiconductor industry (Appleyard, 1996), academic research (Bouty, 2000; Häussler, 2011; Häussler et al., 2014), and financial investment (Crawford et al., 2017; Botelho, 2018; Rantala, forthcoming).

[2]This information exchange can be in the form of comments or suggestions (i.e., a critical evaluation of the idea) or in the form of an additional idea or a refinement of a previously shared idea. For the rest of the paper, we refer to the latter kind.

that a player generates new information—and is therefore *able* to share information—is sufficiently high, then a *sharing equilibrium* in which both players always share exists.

We present two simple extensions of this model. First, we introduce a player's expectations of a rival's intentions because in order to reciprocate a rival must both successfully generate new information and be willing to share it. Second, we assume asymmetric abilities so that the distinction between ability and intention is identifiable. With these adaptations, we answer the following questions: First, can we establish a sustainable exchange of information via this model of a feedback loop, and is this exchange more successful as the relative benefits from the exchange increase? Second, how do players' expectations of reciprocity drive their decisions to initiate and maintain the information exchange? Third, what elements of players' experiences affect the formation of their expectations about a rival's intended behavior?

For our laboratory experiment, we use a narrative in which two fund managers exchange ideas for investment opportunities and compete over the uncommitted capital of investors.[3] We choose this finance-related frame because a financial context would be more familiar to our mostly business school student subjects.[4] Further, because we do not consider trading *per se* but rather consider the exchange of ideas between competitors, the ways that expectations of reciprocity incentivize information exchange are similar to those highlighted in the existing empirical literature.[5]

[3]We provide experimental instructions and a detailed description of the game in the Online Appendix that can be downloaded at `https://sites.google.com/site/bganglmair/research/Reciprocity-OApp.pdf`.

[4]A side effect of the finance narrative is the immediate application to an empirical puzzle in the finance literature. It is surprising that—in spite of the large sunk costs that funds incur to ensure informational security and the potentially larger opportunity costs incurred by disclosing valuable investment ideas—there appears to be evidence suggesting that managers circumvent their own safeguards in order to collaborate with rivals. Why *do* competing fund managers share information? Shiller and Pound (1989:47) survey investors and find that "direct interpersonal communications are very important in [their] decisions," and Shiller (2000:155) concludes that "[w]ord-of-mouth transmission of ideas appears to be an important contributor to day-to-day or hour-to-hour stock market fluctuations." More recent empirical evidence, documenting the extent to which financial trades are correlated (Grinblatt and Keloharju, 2001; Hau, 2001; Feng and Seasholes, 2004; Hong et al., 2004; Ivković and Weisbenner, 2005; Brown et al., 2008; Shive, 2010; Pool et al., 2015), suggests that information sharing among investors continues unabated, and that even hedge fund managers in direct competition with one another appear to share investment ideas.

[5]Fund managers compete across potential inventors for funds, and each manager exerts monopolistic control over a fraction of the market (e.g., through lock-in periods (Agarwal and Naik, 2000) or "side pockets" that are frozen by managers so that redemptions do not force the inefficient early liquidation of assets (McCrary, 2002:192)) and competes with their rivals over the remaining portion (e.g., when hedge funds are not closed to new investors after creation and fund managers continue to compete to raise additional capital either from the fund's existing investors or from new investors (Goetzmann et al., 2003)).

In the theoretical model, if the benefits from sharing information (i.e., the long-term payoffs from a feedback loop) are sufficiently high (relative to the immediate gains from concealing information), then the game becomes a coordination game with two equilibria: both players either always share or never share information. We show empirically that, with high net benefits from sharing, subjects are more likely to play the payoff-dominant sharing strategy. In line with existing empirical results, we conclude that we are more likely to observe an exchange of information when a player's expectation of reciprocity (increasing the net benefits from sharing) is higher. We show that the primary determinants for a player's initiation and maintenance of a feedback loop are the recipient's ability *and* her intention to reciprocate. We further find that a player's expectations of the recipient's intentions to share, rather than her ability, have a greater effect on the player's own incentives to share. Thus, strategic uncertainty (via intentions) has a greater effect on the exchange of information than environmental uncertainty (via ability). Last, we show that subjects form complex and cumulative beliefs about their rivals' intentions to share, in part by documenting how negative past experiences (caused by either a rival or by oneself) result in subjects that are less inclined to share information.

Our paper contributes to the literature in a number of ways. In the context of finance, Crawford et al. (2017), Botelho (2018), and Rantala (forthcoming) provide evidence of information sharing among investment professionals. Our results contribute to this literature that studies the mechanisms behind the empirical phenomenon of correlated trading (e.g., Duffie and Manso, 2007; Colla and Mele, 2010; Manela, 2014; Andrei and Cujean, 2017) by providing a more nuanced picture of individuals' incentives and a direct test of one of the key theoretical arguments (Stein, 2008). Beyond these three studies, the relevant literature is limited by the data and is merely suggestive of managers collaborating in the manner we describe (e.g., Hong et al., 2004; Cohen et al., 2008; Pool et al., 2015). Our results on the effects of subjects' past experiences also relate to recent work on the collaborative and reciprocal nature of crowdfunding platforms. Zvilichovsky et al. (2015), for instance, document the effect of an entrepreneur's funding history on financing outcomes. They show that project owners back their backers at a rate that is significantly higher than that for other comparable projects.

Numerous empirical studies covering a variety of industries have highlighted the role of expected reciprocity as a driver of individuals' incentives to share information.[6] von Hippel (1987) and Schrader (1991) report empirical evidence of know-how sharing of competing firms in the steel minimill industry. Bouty (2000), Häussler (2011), and Häussler et al. (2014) present results for knowledge sharing in academic research. Gächter et al. (2010) (modeling knowledge sharing as a coordination game with multiple equilibria) present experimental results for a setting of private-collective innovation (see von Hippel and von Krogh, 2006) in which private investors fund public goods innovation. Ingram and Roberts (2000) find a positive relationship between financial performance and the existence of friendship-networks (for "better information exchange") between managers of competing hotels. We add to this literature by identifying how expectations of reciprocity can underly and sustain these types of observable information exchange.

The structure of our model is akin to a centipede game (Rosenthal, 1981; Binmore, 1987); however, our game is of indeterminate horizon. In the finite-horizon centipede game, a number of articles have studied subjects' choices in a laboratory setting and found that only a small fraction of games ended in the first round (i.e., the equilibrium outcome)—from 0.7% in McKelvey and Palfrey (1992) to 3.9% in Levitt et al. (2011) (with expert chess players as subjects). Our indeterminate-horizon centipede game has at least two (Nash) equilibria. In fact, we calibrate our model so that the unique equilibrium in the finite-horizon version is not the only equilibrium in our model. We choose a calibration so that one of the equilibria is a sharing equilibrium and focus on the determinants that increase the chance that players coordinate on this payoff-dominant equilibrium.[7]

Last, our results relate to the general literature on disclosure of secrets and exchange of information among agents with competing interests. In recent work, Hellmann and Perotti (2011), Guttman et al. (2014), Dziuda and Gradwohl (2015), and Augenblick and Bodoh-Creed (2018) provide theoretical treatments of different aspects of this general

[6] Botelho (2018) provides a detailed review of the literature.

[7] Our version of the centipede game also differs in a second dimension. The previous literature has assumed an exponential increase in subjects' payoffs (McKelvey and Palfrey, 1992; Nagel and Tang, 1998; Kawagoe and Takizawa, 2012) or a linear increase in payoffs (Bornstein et al., 2004; Gerber and Wichardt, 2010). Following the functional form in Stein (2008), we assume players' payoffs are increasing at a diminishing rate.

theme. Ganglmair and Tarantino (2014) extend the model in Stein (2008) by allowing one of the players to hold prior information (a "secret") about the ex post distribution of payoffs.

The paper is structured as follows. In Section 2, we present our theoretical framework of information exchange as an asymmetric version of the model in Stein (2008). In Section 3, we discuss the experimental design and develop our empirical hypotheses. In Section 4, we present our main results. In Section 5, we conclude with a discussion of implications for the design of formal and informal platforms of knowledge exchange within and across organizations.

2 Theoretical Framework

2.1 Summary of the Model

For our theoretical framework, we consider a centipede game with a non-deterministic final period.[8] Instead, in each round, the game ends *by chance* when a player fails to generate a new idea to share[9] or *by choice* if that player decides to conceal the idea. Concealing an idea yields a temporary payoff advantage. If the player decides to share the idea (and thus continues the game), then the rival has the ability to generate and the subsequent option to share a new idea to increase the overall (common) stock of ideas. A player with an idea to share therefore faces a tradeoff between the short-term gains from concealing or potential long-term gains from cooperation. We use this framework to derive a necessary condition for an environment in which both a *sharing equilibrium* and a *non-sharing equilibrium* exist. We use payoff dominance as an equilibrium selection criterion and argue that relaxing the sharing condition results in the sharing equilibrium being selected more often. This eventually allows us to discuss and study the differential effect of a player's *ability* to generate a new idea and her expectations of the rival's *intentions* to share that new idea on the expected equilibrium outcome.

[8] In this section, we present the main results of the model alongside an abridged version of the notation of the model. We provide a full treatment, including formal results, in the Appendix.

[9] In order to preserve the link of the description of our theoretical model to the storyline in the experiment, we refer to this information as *ideas* (e.g., for investment opportunities) where more ideas increase payoffs.

Table 1: Model Notation for Parameters and Decisions

Variable	Description	Calibration
A, B	Players; A moves in odd rounds, B in even rounds	
p_i	*Ability:* Probability that player $i = A, B$ generates a new idea at the beginning of Round t	50% or 90%
β^n	Production costs with decay parameter β and n the number of ideas	$\beta = 3/4$
μ	Market size	$\mu = 400$
θ	Competition intensity	$\theta = 3/8$
$\tilde{\phi}_i$	Expected net benefits from sharing for player $i = A, B$	
σ_j	Probability that player j shares a newly generated idea	
$\tilde{\sigma}_j$	*Expected intentions:* Player i's expectations that player j shares a newly generated idea	
$\pi_j \equiv p_j\tilde{\sigma}_j$	*Expectations of reciprocity:* Player i's expectations in t that she receives a new idea from player j in $t+1$	

2.2 Formal Setup

Our framework is an asymmetric version of the word-of-mouth communication model in Stein (2008). In Table 1, we provide a summary of the model notation. Two players take turns generating and sharing new ideas. Player A moves in odd rounds, and player B moves in even rounds. Player A is endowed with an idea in $t = 1$ and must decide whether to share this idea with player B or keep it to herself. If she decides to conceal the idea, then the game ends. If she shares, then the game continues. Upon observing A's idea, in $t = 2$, player B generates a new idea with probability p_B. If she is successful, she must decide whether to share this idea with A or keep it to herself. If she conceals the idea, the game ends. If she shares, the game proceeds to $t = 3$ in which player A can generate a new idea with probability p_A. The game proceeds in this fashion in all t. We provide a graphical depiction of the game structure in panel (a) of Figure 1.

More ideas increase a player's potential payoffs. The number of ideas n_i a player i can access is the common stock of ideas: the number of ideas she herself has generated plus the number of ideas the other player has shared with her. Stein (2008) uses a simple model in which ideas reduce a player's production costs β^{n_i} with $\beta < 1$. We follow this payoff structure and provide a graphical illustration in panel (b) of Figure 1. Each player faces a market (of size μ) with unit-demand consumers that have a valuation of one. The players compete à la Bertrand in a fraction θ of the market and are monopolists in their $1-\theta$ respective markets. In the monopolistic segment of their market, additional ideas increase profits regardless of how many ideas the other player can access. A player's payoffs from this segment, when holding n_i ideas, are $(1-\theta)\,\mu\,(1-\beta^{n_i})$. In the competitive segment, however, what is important is the relative number of ideas the other player holds. Suppose a player i has access to more ideas than player j so that $n_i > n_j$ and $\beta^{n_i} < \beta^{n_j}$. By Bertrand competition, the price in the competitive segment is β^{n_j}, and player i's profits from that segment are equal to $\theta\mu\,(\beta^{n_j} - \beta^{n_i})$. Conversely, if $n_i < n_j$, then player i's payoffs in the competitive segment are equal to zero.[10] The realized payoffs for player i are

$$\mu\left[(1-\theta)\,(1-\beta^{n_i}) + \theta \max\left\{\beta^{n_j} - \beta^{n_i}, 0\right\}\right]. \tag{1}$$

The payoffs in the model give rise to a simple tradeoff for a player who was successful in generating a new idea in t. Player i must choose between short-term gains from concealing and potential future gains from sharing. Concealing the idea gives her a cost advantage of $\beta^{t-1} - \beta^{t}$ over the other player because $n_i = t$ and $n_j = t-1$. This cost advantage materializes immediately because concealing the idea also means that the game ends and no further ideas can be generated. We refer to such termination of the game as *termination by choice*. A player's payoffs in t in this case are

$$\text{conceal}_t = \mu\left[1 - \beta^{t} - \theta\left(1 - \beta^{t-1}\right)\right]. \tag{2}$$

Sharing the idea, on the other hand, gives the rival j a chance to generate a new idea in $t+1$ (with probability p_j) and then share this idea back with player i, which will give i the chance to continue the game in $t+2$, and so forth. When making this decision to share, player i must form expectations about receiving information back from player j. In order for player i to draw another new idea in $t+2$, player j must generate a new idea in $t+1$ *and*

[10] The structure of the payoffs here can also be seen in the professional investment market described in the introduction and used in the experiment. For a manager's payoffs, both the returns on capital and the amount of capital matter. The monopolistic segment captures the payoffs from existing capital (where more ideas generate higher returns), whereas the competitive segment captures the payoffs from attracting additional capital (where more ideas than the competitors generate higher fund flows and thus higher management fees).

Figure 1: Model Timeline and Payoff Structure

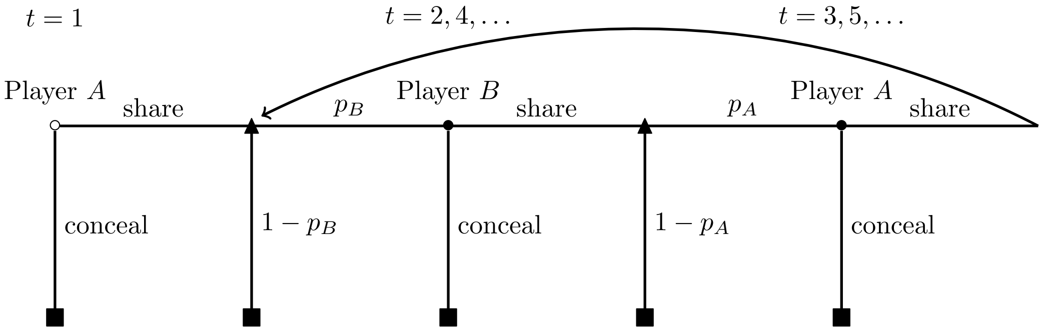

(a) Timeline of Word-of-Mouth Communication

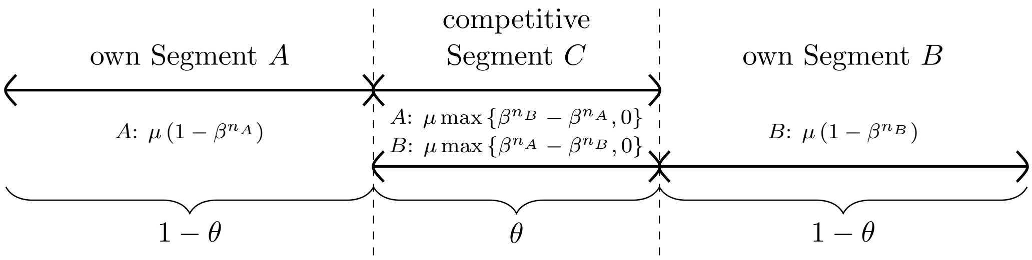

(b) Competition and Payoffs

be willing to share that idea with player i. Suppose that player j shares a newly generated idea with probability σ_j.[11] We can then write player i's expectations about player j sharing as $\tilde{\sigma}_j \equiv E(\sigma_j)$. Combining both components, a player i's expectations of receiving a new idea from player j in $t+1$ upon sharing an idea in t (i.e., the expectations of reciprocity) are $\pi_j \equiv p_j \tilde{\sigma}_j$.

A player i's expected payoffs from sharing in t (and all future $t' > t$), when she expects information in return with probability π_j, is then equal to

$$\text{share}_{it} = \mu\,(1-\theta) \sum_{k=0}^{\infty} p_i^k \pi_j^k \left[(1-\pi_j)\left(1-\beta^{t+2k}\right) + \pi_j\left(1-\beta^{t+1+2k}\right)\right]. \tag{3}$$

Player i shares a new idea in t if $\text{share}_{it} \geq \text{conceal}_t$. After some manipulation, this expression can be rewritten as

$$\tilde{\phi}_i(\tilde{\sigma}_j) \equiv \mu \left[\frac{1+\beta p_i}{1+\beta \pi_j} \beta \pi_j - \theta\right] \geq 0. \tag{4}$$

[11]For instance, assume that σ_j, $j = A, B$ is a time-invariant mixed strategy in t.

This term $\tilde{\phi}_i(\tilde{\sigma}_j)$ denotes a player i's expected net benefits from sharing in t when she expects player j to share a newly generated idea with probability $\tilde{\sigma}_j$ in all future rounds.[12] The net benefits from sharing are increasing in $\pi_j = p_j\tilde{\sigma}_j$ and p_i.

2.3 Equilibrium and Selection

Condition (4) is necessary for a sharing equilibrium – but not sufficient. If the net benefits from sharing $\tilde{\phi}_i(\tilde{\sigma}_j)$ are positive for some $\tilde{\sigma}_j$, then two pure-strategy equilibria exist:[13] First, in a *sharing equilibrium*, all players always share newly generated ideas (so that in equilibrium $\sigma_j = \tilde{\sigma}_j = 1$ for $j = A, B$), and the game continues until it is terminated by chance. Second, in a *non-sharing equilibrium*, neither player ever shares a new idea, and the game is terminated by choice (by player A) in $t = 1$.

We utilize payoff dominance as our equilibrium selection criterion (Harsanyi and Selten, 1988; Cooper et al., 1990) and exploit relative payoff differences for our theoretical predictions of the model (in the Appendix). From the sharing condition (4) we can see that the sharing equilibrium payoff-dominates the non-sharing equilibrium. Anything that relaxes the sharing condition, by increasing the value of $\tilde{\phi}_i(\tilde{\sigma}_j)$, reinforces the payoff dominance and induces players to choose the sharing strategy more often.

3 Experimental Design and Hypotheses

3.1 Experimental Design

We conducted the computerized experiments at the Center and Laboratory for Behavioral Operations and Economics (CLBOE) at the University of Texas at Dallas. The participants were registered with CLBOE and were drawn from a pool of both undergraduate and graduate students. 100 subjects participated across four different treatments. Each subject participated in only one treatment. Each session lasted anywhere from 80 to 120 minutes,

[12] The expected payoffs from sharing in equation (3) reduce to the analogous expression in Stein (2008) when $\tilde{\sigma}_j = 1$ and $p_i = p_j = \pi_j = p$. Similarly for sharing condition (4) that reduces to the simple sharing condition, $\phi_i \equiv \mu\,[p\beta - \theta] \geq 0$ with $\mu = 1$.

[13] In the Appendix, we fully characterize these two pure strategy equilibria (Proposition A.1). We also characterize a mixed-strategy equilibrium in which player i's mixed strategy σ_i is such that $\tilde{\phi}_j(\sigma_i) = 0$ and player j's expectations are consistent with this strategy so that $\tilde{\sigma}_i = \sigma_i$ (Proposition A.2).

depending on the treatment. Payments ranged from \$10 to \$30, averaging \$19.30.[14] Subjects in longer sessions generally had greater earnings.

The number of subjects ranged from 24 to 28 in each session. We randomly divided the subjects into two groups of equal size. Group membership was anonymous, meaning that subjects did not know who else was assigned to a particular group. They were informed that they had been randomly assigned to a group and would be matched only with subjects from the same group.

Each session was divided into two parts: the first part consisted of a Holt-Laury risk-preference task (Holt and Laury, 2002), and the second part consisted of our main experiment. We conducted the Holt-Laury risk-preference task via paper and dice before conducting the main experiment.[15] The main experiment was programmed and executed via zTree (Fischbacher, 2007). We revealed the outcomes of the lottery in the Holt-Laury risk-preference task and the respective payoffs after the computerized experiment at the end of the session. We provided subjects with detailed printed instructions for both the Holt-Laury task and the computerized experiment and conducted a short quiz after the experimenter had read out the instructions.

In the computerized experiment, at the beginning of each period, we randomly matched subjets into pairs without replacement. After the matches were determined, we randomly assigned the subjects the roles of player A and player B with all the parameters and payoffs being common knowledge for the pair. The parameters did not change for the duration of the match. When new matches were established, new roles and their respective values of p_i were revealed as common knowledge. For the instructions of the word-of-mouth communication game, we used a fund-manager narrative in which two players (i.e., fund managers) exchange investment opportunity ideas and compete over the capital of new investors.[16] In each round $t \geq 1$ of a match, after having generated a new idea (with probability p_i), player i takes two actions. First, we surveyed player i's expectations of player j's

[14] These figures include a show-up bonus of \$5 and average payoffs of \$2.5 from a Holt and Laury (2002) risk-preference task.

[15] We control for subjects' risk preferences and many other individual characteristics in robustness results presented in the Appendix.

[16] From the instructions of the computerized experiment: "You are a fund manager. Your goal is to earn as much money as possible. Your earnings can increase in two ways: a) increase the returns from your investments and b) obtain more investors."

intentions, $\tilde{\sigma}_j$.[17] We did so by asking player i to report her expectation (between 0% and 100%) that player j would decide to share an idea in the next round (provided that player j generated a new idea).[18] Second, the player decided whether to share or conceal the idea.

On their decision screens, subjects saw their assigned role (fund manager A or B) and payoffs (for both players) for the current round and the subsequent two rounds for all possible outcomes.[19] Provided that player i shared an idea in round t, once player j generated a new idea with probability p_j in $t+1$, she was shown the decision screen for that round. If, instead, player i decided to conceal the idea, then the match was terminated. Alternatively, if player i failed to generate a new idea, then the match was also terminated. From the current match's payoffs, players were able to infer whether the match had been terminated by chance (no new idea was generated) or choice (a player decided to conceal the new idea).[20] After all matches had been terminated, the subjects observed their payoffs from the current match and their accumulated payoffs from all previous matches. This concluded a match. We then rematched the subjects within their respective groups, and a new game was played.

We had four treatments—only one treatment per session but with multiple groups per treatment. We implemented the game depicted in Figure 1 with the realized payoffs in equation (1). We set $\mu = 400$, $\beta = 3/4$, and $\theta = 3/8$ but varied the success probabilities p_A and p_B by assigning values $p_i \in \{50\%, 90\%\}$ depending on the treatment.[21] Therefore, our four treatments referred to as High, Low, Low-High, and High-Low varied only in p_A and p_B. We summarize the calibrations for the four treatments of the experiment in Table 2.

[17]For all odd rounds, we obtain player A's expectations $\tilde{\sigma}_B$; for even rounds, we obtain player B's expectations $\tilde{\sigma}_A$. Note that, while the truthful reporting of expectations $\tilde{\sigma}_j$ is not incentived, formation of these expectations is fully incentivized within the game itself. This is because, as player j's future actions have a direct effect on player i's payoffs, player i's expected payoffs increase in the accuracy of her expectations $\tilde{\sigma}_j$. We are confident that, on average, these expectations are reported truthfully, possibly with more noise (Trautmann and van de Kuilen, 2015).

[18]We used the following wording: "If, in the next round, the other fund manager successfully generates a new idea (i.e., "chance" does not terminate the match), how likely do you think the other fund manager will share this newly generated idea with you?"

[19]In the printed instructions for the experiment, we provided a table with player A's and player B's payoffs for the first 14 rounds for all possible paths of termination of a match.

[20]If the match terminates with both players having the same payoff, then it was terminated by chance. We make this point explicit in the printed instructions.

[21]For the results presented in the main text, we use a constant degree of competition θ. We provide results on the effect of competition (through a higher value of θ) in the Online Appendix.

Table 2: Calibration and Treatments

	$p_B = 90\%$	$p_B = 50\%$
$p_A = 90\%$	Treatment High	Treatment High-Low
	($\tilde{\phi}_i(1) = 120$, for $i = A, B$)	($\tilde{\phi}_A(1) = 32.73$, $\tilde{\phi}_B(1) = 71.64$)
$p_A = 50\%$	Treatment Low-High	Treatment Low
	($\tilde{\phi}_A(1) = 71.64$, $\tilde{\phi}_B(1) = 32.73$)	($\tilde{\phi}_i(1) = 0$, for $i = A, B$)

In all four treatments, the sharing condition (4) is satisfied for $\tilde{\sigma}_j = 1$ so that $\tilde{\phi}_i(1) \geq 0$. That means, in all four treatments, a sharing equilibrium exists.

3.2 Hypotheses

We propose seven hypotheses related to the effects of a player's and her rival's ability, as well as the player's expectations about the rival's future intentions. The model in Section 2 drives all our hypotheses. We detail formal derivations in Appendix A.

We begin by focusing on the first decision by player A (in Round 1) in each match. This way, the player's expectations are not affected by an earlier decision of the current rival (within the same match). We expect player A to share more often when her expected net benefits $\tilde{\phi}_A$ (provided in Table 2) are higher.

Hypothesis 1 (Higher Expected Net Benefits)**.** *Player A shares more often in $t = 1$ when $\tilde{\phi}_A$ is high. That means she shares (a) more often in treatment* High *relative to* Low-High*, (b) more often in treatment* Low-High *relative to* High-Low*, and (c) more often in treatment* High-Low *relative to* Low*.*

Furthermore, player A's expected net benefits $\tilde{\phi}_A$ are an increasing function of her expectation of receiving new information, $\pi_B = p_B \tilde{\sigma}_B$. Therefore, *ceteris paribus*, player A should share more when π_B increases: first, as the ability of player B, captured by p_B, increases (Hypothesis 2); and second, as player A's expectations that player B will share a newly generated idea, captured by $\tilde{\sigma}_B$, increases (Hypothesis 3).

Hypothesis 2 (Rival's Ability)**.** *Player A shares more often in* $t = 1$ *when* p_B *is high. That means she shares more often in treatment* High *relative to* High-Low *and* Low-High *relative to* Low.

Hypothesis 3 (Expected Intentions)**.** *Player A shares more often in* $t = 1$ *when her beliefs* $\tilde{\sigma}_B$ *about player B sharing an idea in* $t = 2$ *are high.*

While p_B and $\tilde{\sigma}_B$ both increase π_B, comparative statics (utilizing our model results) show that they exhibit different marginal effects on player A's decision to share in $t = 1$. In fact, the effect of p_B is stronger than that of $\tilde{\sigma}_B$ if $p_B < \tilde{\sigma}_B$. If, conversely, $p_B > \tilde{\sigma}_B$, then $\tilde{\sigma}_B$ has a bigger effect than p_B at the margin. This is Hypothesis 4.

Hypothesis 4 (Expected Intentions vs. Rival's Ability)**.** *As the ratio* $p_B/\tilde{\sigma}_B$ *increases, the difference of the marginal effects of* p_B *and* $\tilde{\sigma}_B$ *changes from negative to positive.*

We further expect player A's own ability, p_A, to affect her own decision to share. This effect is two-pronged. Player A's own-success probability has a positive and *direct* effect when deciding to share. This reflects the player's expectations of a mechanical increase in game length (and generated value). Moreover, an increase in p_A comes with a positive and *indirect* effect through the player's expectations of her rival's intentions, $\tilde{\sigma}_B$. The logic for the indirect effect is that, as player A's ability increases, the likelihood that player B will share (through symmetry of Hypothesis 2) also increases. This effect on player B induces player A to increase her expectations $\tilde{\sigma}_B$ that player B will share a newly generated idea. The following hypothesis reflects this joint effect. Our theoretical framework, however, does not inform us as to which of the two effects is stronger. We will address this empirical question utilizing regression results in the next section.

Hypothesis 5 (Own Ability)**.** *Player A shares more often in* $t = 1$ *when* p_A *is high. That means she shares more often in treatment* High *relative to* Low-High *and* High-Low *relative to* Low.

Our next hypothesis draws a comparison between a player's own ability, p_A, and the rival's ability, p_B while holding $\tilde{\sigma}_B$ constant. Accounting for the effect of expectations is important because expectations can change when either player's abilities change. Comparative statics show that the effect of p_A on player A's willingness to share is weaker than the effect of p_B. Intuitively, this difference arises because the outcome of the rival's ability (i.e., generating a new idea) is realized immediately, whereas the outcome of player A's ability is not realized until the following round.

Hypothesis 6 (Own vs. Rival's Ability)**.** *Holding expectations $\tilde{\sigma}_B$ constant, the positive effect of p_A on player A's decision to share in $t = 1$ is weaker than the effect of p_B. That means player A shares more often in treatment* Low-High *than* High-Low.[22]

Our final hypothesis addresses the formation of a player's expectations $\tilde{\sigma}_B$ as a function of abilities, p_A and p_B. Of the two abilities, we expect player A's own ability to have a stronger effect. The intuition is an extension of the idea behind Hypothesis 6. By symmetry, p_A has a stronger effect than p_B on player B's willingness to share. As a consequence, p_A will have a stronger effect than p_B on player A's expectations of player B's intentions, that is, on $\tilde{\sigma}_B$. Hypothesis 7 summarizes this more generally:

Hypothesis 7 (Effect of Abilities on Expectations)**.** *The effect of p_i on $\tilde{\sigma}_j$ is positive and istronger than the positive effect of p_j on $\tilde{\sigma}_j$.*

4 Experimental Results

This section contains our results and is divided into three parts. First, we discuss the data in aggregate and provide summary results of sharing. Next, we provide detailed results related to ability and intention. Last, we focus on game dynamics and the formation of beliefs about intentions.

4.1 Aggregate Data and Sharing

We begin by addressing the first of our three questions: Can we establish a sustainable exchange of information in our model of a feedback loop, and is this exchange more successful

[22]For this comparison of treatments, let $\sigma_A(\text{High}) - \sigma_A(\text{Low-High})$ denote the effect of p_A on player A's sharing behavior, and let $\sigma_A(\text{High}) - \sigma_A(\text{High-Low})$ denote the effect of p_B on her sharing behavior. The effect of p_A is weaker than that of p_B if $\sigma_A(\text{Low-High}) > \sigma_A(\text{High-Low})$, when player A shares more often in treatment Low-High than High-Low.

Table 3: Summary Statistics of Experimental Results

This table provides basic summary statistics for the four main treatments of the experiment (High, Low, Low-High, and High-Low), as summarized in Table 2. We conducted all treatments in one session with two groups of equal size s_g. For the calibration of the treatments, see Table 2. We list the number of subjects per treatment; the number of matches (i.e., the number of pair-wise word-of-mouth communications, $s_g(1 - s_g)$); the average number of rounds each match proceeds; the average earnings per match (in $) for each subject; and the average expected intentions for each of the players.

	Treatment			
	High	Low	Low-High	High-Low
Subjects	24	28	24	24
Matches	132	182	132	132
Average # of rounds (and decisions by a player)	5.62	1.43	2.60	1.70
Average earnings (in $) per match ...				
... for all subjects	1.57	0.75	1.16	0.83
... for player A	1.58	0.91	1.19	0.98
... for player B	1.56	0.59	1.13	0.67
Expected intentions ...				
... by player A (reported $\tilde{\sigma}_B$)	82.4%	56.1%	59.7%	50.9%
... by player B (reported $\tilde{\sigma}_A$)	78.7%	59.8%	56.4%	64.1%
Theoretical expected net benefits $\tilde{\phi}_A(1)$ in Round 1	120.00	0.00	71.64	32.73
Theoretical expected net benefits $\tilde{\phi}_B(1)$ in Round 2	120.00	0.00	32.73	71.64

as the relative benefits from the exchange increase? Basic descriptive statistics paint a general picture of the players' actions and the outcomes in all four treatments. In Table 3, we report the total number of subjects and matches for each treatment, the average duration of each match, the average earnings (per match) for each subject, and how players expected their rivals to behave in each treatment. We also report the expected net benefits (in a sharing equilibrium) of a player continuing in any given period: $\tilde{\phi}_A(1)$ and $\tilde{\phi}_B(1)$.

To start, we observe that treatments with higher expected net benefits $\tilde{\phi}$ tend to have longer matches. Treatment High had an average of 5.62 rounds while treatment Low only had an average of 1.43 rounds. Longer matches are indicative of the players choosing sharing-equilibrium strategies. However, a simple comparison of the duration of matches across treatments is misleading. Direct (first-order) and indirect (higher-order) effects are at play here. For the direct effect, higher success probabilities (by either player A or player B) mechanically increase the duration of a match, as a match is less likely terminated by *chance* (holding $\tilde{\sigma}_i$ constant). For the indirect effect, higher success probabilities are likely to increase the values of $\tilde{\sigma}_i$. We see this effect on expectations of sharing when we compare

Figure 2: Distribution of Match Outcomes

This figure depicts the distribution of match outcomes for all four treatments (with $N = 182$ matches in treatment Low and $N = 132$ in the other treatments). We show five distinct outcomes: a match can be (1) terminated by player A in $t = 1$, (2) terminated by player B in $t = 2$, (3) terminated by player A in any $t \geq 3$, (3) terminated by player B in any $t \geq 4$, (5) terminated by chance because player A fails to generate a new idea in any odd t, or (6) terminated by chance because player B fails to generate a new idea in any even t. Numbers for each treatment may not sum up to 100% due to rounding errors.

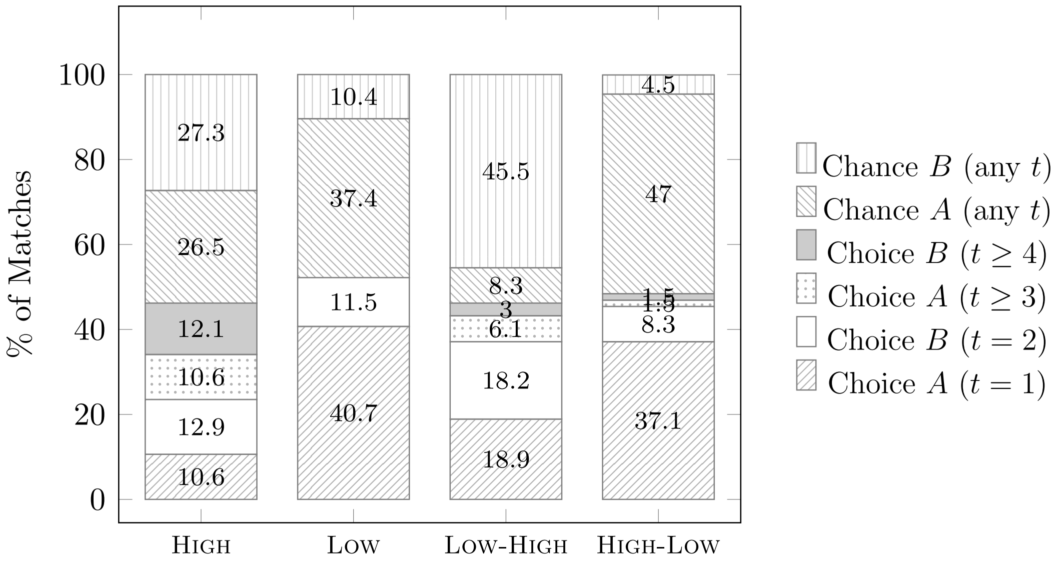

expectations $\tilde{\sigma}_B$ by player A in treatment Low (56.1%) to expectations in treatment High (82.4%). Likewise, we see the same pattern when comparing player B's expectations $\tilde{\sigma}_A$ in treatment Low (59.8%) to the expectations in treatment High (78.7%).

As a first consideration, we take a look at the distribution of causes for termination in Figure 2. For each treatment, we provide the proportion of matches terminated by choice by A in $t = 1$, by B in $t = 2$, by A in rounds $t \geq 3$, by B in rounds $t \geq 4$, or by chance because either A or B failed to generate a new idea. The data suggest that, when the benefits from sharing are higher, players are more likely to play a sharing strategy. We see this in the form of a smaller number of all matches terminated by choice by players in treatment High (46.2%) than in treatment Low (52.1%). When considering player A's decision in Round 1, only 10.6% of the matches were ended by player A in treatment High while a much higher proportion (40.6%) was terminated by player A in treatment Low. This suggests more sharing, implying a payoff-dominant equilibrium outcome.

For player A, the proportion of matches terminated by choice in $t = 1$ is equivalent to player A's termination rate in Round 1—where the termination rate is simply one minus the mean of sharing in $t = 1$. For our first result, we restrict our attention to player A's behavior in Round 1 of each match so that her actions are not affected by the history of that match.

Result 1. *As the expected net benefits of sharing increase, players choose the sharing equilibrium strategies more often.*

In panel (a) of Figure 3, we plot the mean of sharing by player A in Round 1 for each treatment. For Hypothesis 1, we have ranked the treatments by player A's expected net benefits from sharing, predicting more sharing when benefits are higher. From a set of means tests, we can reject the null for Hypotheses 1(a) and 1(b) at the 5% and 1% level, respectively, but we fail to reject the null for Hypothesis 1(c) at the 10% level.[23]

We observe similar patterns for player B's behavior in Round 2—we plot the mean of sharing for each treatment in panel (c) of Figure 3. As expected, player B's termination rate in treatment Low is considerably higher than in treatment High.[24] We return to higher-round behavior by players in our analysis of game dynamics in Subsection 4.3.

4.2 Reciprocity as a Combination of Ability and Intentions

We next present our main results addressing the second question: How do players' expectations of reciprocity drive their decision to initiate and maintain the information exchange? The probability of reciprocity, π_B, is composed of a rival's *ability* p_B to generate information (Result 2) and her expected *intention* $\tilde{\sigma}_B$ to share that information (Result 3).

Results 2 and 3 below establish a positive effect of expected reciprocity on player A's willingness to share. They comport with the collaboration argument in Stein (2008) and Crawford et al. (2017) (using Stein's framework) and confirm the empirical evidence cited earlier: higher expectations of reciprocity increase players' incentives to share.

Result 2. *Player A is more likely to share an idea when the rival player has higher ability.*

We get a first glimpse of this result by comparing the percentage of matches (in Figure 2) terminated by player A in treatment High with treatment High-Low (21.2% <

[23] The average treatment effect for Hypothesis 1(a) is 0.0833 (s.e.: 0.043); for Hypothesis 1(b) is 0.1818 (s.e.: 0.054), and Hypothesis 1(c) (s.e.: 0.055).

[24] The numbers in Figure 2 paint a different picture because they are not proper termination rates but the proportion of matches terminated by B in $t = 2$ in percent of all matches.

Figure 3: Sharing and Expectations in Rounds 1 and 2

This figure plots the average level of sharing in Round 1 by player A (panel (a)), player A's expectations $\tilde{\sigma}_B$ in Round 1 (panel (b)), the average level of sharing in Round 2 by player B (panel (c)), and player B's expectations $\tilde{\sigma}_A$ in Round 2 (panel (d)) for all four treatments. In panels (b) and (d), we provide box plots for the players' expectations $\tilde{\sigma}_j$.

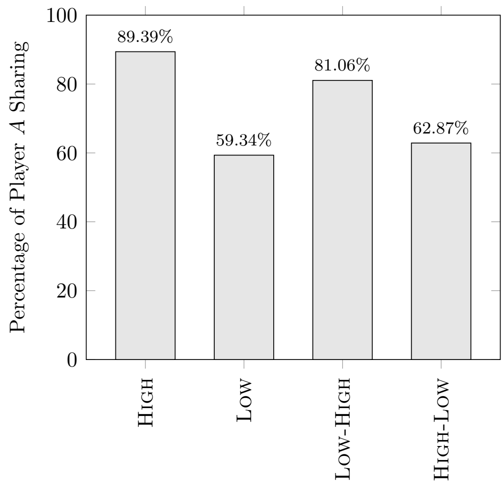

(a) Sharing in Round 1

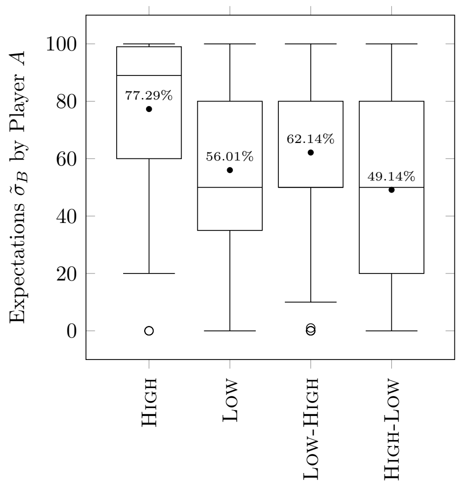

(b) Expectations in Round 1

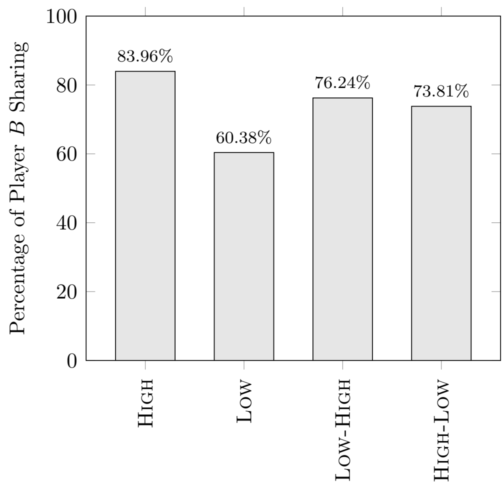

(c) Sharing in Round 2

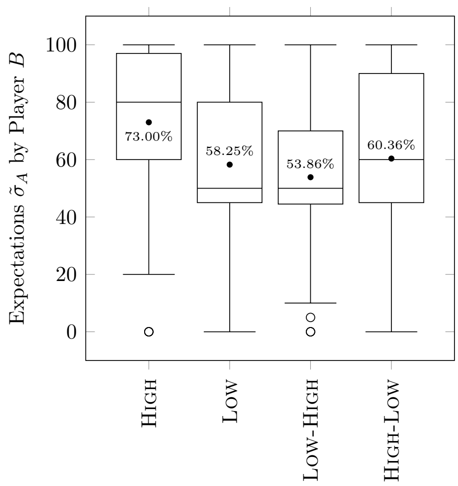

(d) Expectations in Round 2

38.6%), as well as in treatment Low-High with treatment Low (25.0% < 40.7%). Next, we restrict our attention to player A's behavior in Round 1 of each match (panel (a) of Figure 3). We conclude from simple means tests (in Table A.3 in the Appendix) that the average treatment effect is as predicted in Hypothesis 2. We reject the null for the hypothesis at the 1% level.

In Table 4, we also present regression results from probit models. The dependent variable is a dummy variable equal to 1 if player A shares in Round 1 and equal to 0 otherwise. We find a positive marginal effect of the cross-success probability p_B on player A's sharing behavior, further supporting Result 2.[25]

Result 3. *Player A is more likely to share an idea when she has a higher expectation of competitor's intention to share.*

We can reject the null for Hypothesis 3 given the results from simple means tests (in Table A.4 in the Appendix) and a positive and significant marginal effect of expected intentions $\tilde{\sigma}_B$ on player A's sharing across all specifications in Table 4.

The marginal effects reported in Table 4 (supporting our Hypotheses 2 and 3 and collectively supporting Hypothesis 1) imply that player A is 3.4% to 6.1% more likely to share an idea in Round 1 in response to a 10 percentage-point increase in the cross-success probability p_B. Moreover, she is 5.6% to 6.3% more likely to share in Round 1 in response to a 10 percentage-point increase in her expectations $\tilde{\sigma}_B$ that player B will share an idea in Round 2. These effects are also economically significant. A simple metric to measure a real effect is the increase in the generated value of the information exchange that is attributable to player A's sharing behavior.[26] This (lower bound) value increase is 3% in response to

[25]Note that for these results we use minimal specifications but provide robustness results with extended model specifications in the Appendix. Our results concerning how the probability of reciprocity (measured by ability and intentions) affects sharing are robust to a set of variables capturing trust, fairness, and personal connections: all of which have been associated with increased cooperative or pro-social behavior. We report these results in Table A.1 and provide detailed descriptions and summary statistics for these control variables in Table A.2.

[26]For this calculation, we use the expected duration of the information exchange as derived in Proposition refthm:prop3 in the Appendix. For the values of p_i we use the empirical means (across all treatments) of the players' success probabilities. For the values of σ_i we use the empirical means of the players' sharing rates in Round 1 (for A) and Round 2 (for B) and take them to be constant over time (implying that the calculated effects are lower bounds). To calculate the percentage changes in real value, we compare the duration (accounting for 10 percentage-point increase of p_B or $\tilde{\sigma}_B$) without the behavioral response of A (keeping sharing rates constant) to the duration with the behavioral response (adjusting A's sharing rate by the estimated effects from Model (V) in Table reftab:probit1).

Table 4: Baseline Results for the Effects of Ability and Intentions

We report probit results for all four treatments. The dependent variable is a dummy variable = 1 if player A shares in Round 1 and = 0 otherwise. Player A's expectations of receiving information in return are captured by *Cross success:* p_B (player B's cross success probability) and *Expected intentions:* $\tilde{\sigma}_B$ (player A's expectations that player B will share in Round 2). *Own success:* p_A is player A's own success probability. The number of observations is the number of Round 1 decisions by player A. Reported marginal effects are average marginal effects. We report standard errors in parentheses.

	Dependent variable = 1 if player A shares in Round 1 and = 0 o.w.				
	(I) ME	(II) ME	(III) ME	(IV) ME	(V) ME
Cross success: p_B	0.0034 ***	0.0061 ***		0.0060 ***	0.0035 ***
	(0.0008)	(0.0008)		(0.0008)	(0.0008)
Expected intentions: $\tilde{\sigma}_B$	0.0056 ***		0.0063 ***		0.0056 ***
	(0.0004)		(0.0004)		(0.0004)
Own success: p_A				0.0015 *	0.0014 *
				(0.0009)	(0.0008)
Observations	578	578	578	578	578
pseudo R^2	0.2256	0.0645	0.2008	0.0685	0.2299
Log-likelihood	-265.54	-320.78	-274.02	-319.38	-264.05

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

a 10 percentage-point increase in p_B (an elasticity of 0.18) and 4.5% in response to a 10-percentage-point increase in $\tilde{\sigma}_B$ (an elasticity of 0.23).

In Result 4, we summarize the relative effects of the individual components of expected reciprocity (Hypothesis 4):

Result 4. *The effect of strategic uncertainty (captured by $\tilde{\sigma}_B$) is stronger than environmental uncertainty (captured by p_B).*

The marginal effects of the components, as reported in Table 4, exhibit a statistically significant difference.[27] In Table 5, we further report the marginal effects of ability and expected intentions when evaluated at different combinations of p_B and $\tilde{\sigma}_B$ (such that $\pi_B = \bar{\pi}_B = \text{mean}(p_B \times \tilde{\sigma}_B) = 0.4144$). Unlike what is predicted by our theoretical model, we find a negative difference of the marginal effects in all three scenarios in the table.[28] Our results, therefore, do not support the model prediction, and we refute Hypothesis 4. What we have demonstrated, however, is that the effect of expected intentions (representing strategic uncertainty) is stronger than that of ability (representing environmental uncertainty).

[27] A Wald test of simultaneous equality of coefficients for p_j and $\tilde{\sigma}_j$ in the models in Tables 4 rejects the null of equality with at least $p < 0.10$ in all cases.

[28] Hypothesis 4 implies that the difference of the marginal effects ought to be positive in the first column of Table 5 (with $p_B/\tilde{\sigma}_B < 1$), zero in the second column (with $p_B/\tilde{\sigma}_B = 1$), and negative in the third (with $p_B/\tilde{\sigma}_B > 1$). See the Appendix for the formal argument.

Table 5: Differential Effects of Ability and Intentions

We report probit results for all four treatments. The dependent variable is a dummy variable = 1 if player A shares in Round 1, and = 0 otherwise. Player A's expectations of receiving information in return are captured by *Cross success:* p_B (player B's cross success probability) and *Expected intentions:* $\tilde{\sigma}_B$ (player A's expectations that player B will share in Round 2). *Own success:* p_A is player A's own success probability. Marginal effects (ME) for model (V) in Table 4 are evaluated at values of p_B and $\tilde{\sigma}_B$, keeping $\pi_B = p_B\tilde{\sigma}_B$ constant at $\bar{\pi}_B = \text{mean}(p_B \times \tilde{\sigma}_B) = 0.4144$ (i.e., the sample mean probability of reciprocity), while p_A is at the sample mean. In column (2), we evaluate at values of p_B and $\tilde{\sigma}_B$ such that $p_B/\tilde{\sigma}_B = 1$ and the theoretical prediction for the difference of marginal effects is zero. In column (1), $p_B/\tilde{\sigma}_B < 1$; in column (3), $p_B/\tilde{\sigma}_B > 1$. The number of observations is 578; the pseudo R^2 is 0.2299. We report standard errors in parentheses. In the second part of the table, we report the coefficient of a χ^2 test of the difference between p_B and $\tilde{\sigma}_B$ with p-values in parentheses.

	Dependent variable = 1 if player A shares in Round 1 and = 0 otherwise		
	ME evaluated at $p_B = 50\%$, $\tilde{\sigma}_B = \bar{\pi}_B/p_B$	ME evaluated at $p_B = \tilde{\sigma}_B\%$, $1 = \bar{\pi}_B/p_B$	ME evaluated at $p_B = 80\%$, $\tilde{\sigma}_B = \bar{\pi}_B/p_B$
Cross success: p_B	0.0033 ***	0.0040 ***	0.0042 ***
	(0.0010)	(0.0010)	(0.0009)
Expected intentions: $\tilde{\sigma}_B$	0.0054 ***	0.0065 ***	0.0068 ***
	(0.0004)	(0.0006)	(0.0008)
Own success: p_A	0.0014 *	0.0016 *	0.0017 *
	(0.0008)	(0.0009)	(0.0010)
	Test of the difference in coefficients : χ^2		
$p_B - \tilde{\sigma}_B = 0$	5.62 **	4.22 **	3.56 *
	(0.0178)	(0.0400)	(0.0593)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

To illustrate the results from Table 5 and provide additional robustness, we plot (Figure 4) the difference of the marginal effects as a function of the ratio $p_B/\tilde{\sigma}_B$. The solid line depicts the graph for combinations of p_B and $\tilde{\sigma}_B$, keeping the probability of feedback constant at its sample mean (as in Table 5). The difference of marginal effects is negative for *all* values of $p_B/\tilde{\sigma}_B$. To establish robustness of this result, we further calculate and plot the differences of marginal effects holding the probability of feedback constant at various other values as dashed lines.

Although expectations appear more salient to the sharing decision than ability, ability is still a significant predictor of whether a player will share an idea. This is true for both a rival's ability—summarized in Result 2 above—as well as for a player's own ability. We summarize this latter effect in Result 5 below. We also show that the rival's ability has a stronger effect on a player's behavior than the player's own ability. We summarize this in Result 6.

Result 5. *Player A is more likely to share an idea when her own ability is high.*

Figure 4: Strategic Uncertainty and Environmental Uncertainty

We plot the difference of the marginal effect (on player A's decision to share) of ability p_B and intentions $\tilde{\sigma}_B$ as function of the ratio $p_B/\tilde{\sigma}_B$, keeping the probability of reciprocity constant at π_B. The solid line depicts the results for the sample mean of π_B, $\pi_B = \bar{\pi}_B = \text{mean}(p_B \times \tilde{\sigma}_B) = 0.4144$. The dashed lines depict results for $\pi_B \in \{20\%, 30\%, 50\%, 60\%, 70\%, 80\%\}$. For $\pi_B = 20\%$, we plot values for $p_B/\tilde{\sigma}_B$ below 4. The dotted vertical line depicts $p_B/\tilde{\sigma}_B = 1$. We predict a difference of zero for this value and a positive difference for $p_B/\tilde{\sigma}_B < 1$. The results are from probit results for all four treatments (in Table 5), with the dependent variable a dummy variable = 1 if player A shares in Round 1, and = 0 otherwise.

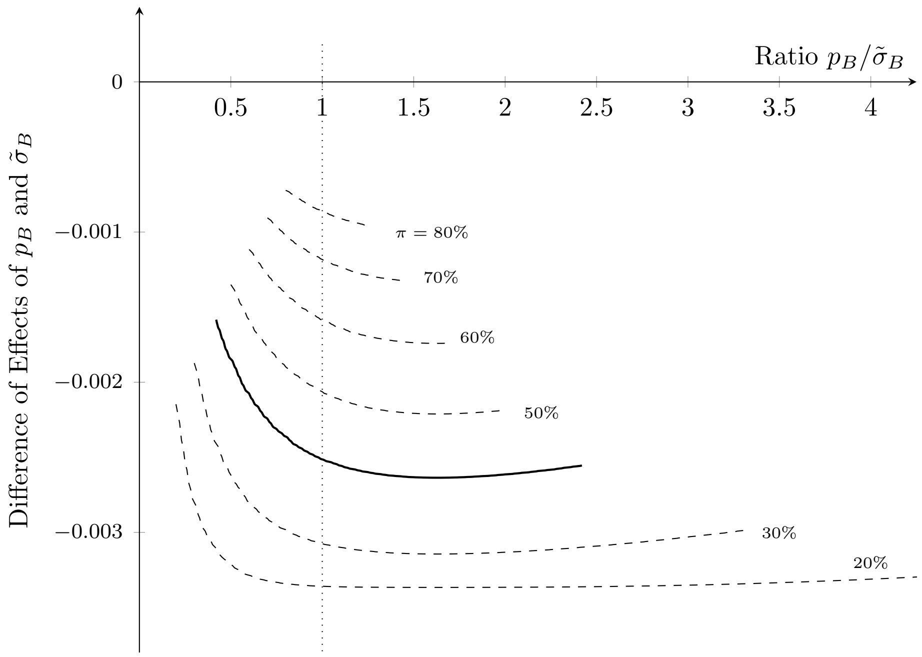

We can reject the null of Hypothesis 5—reporting results from simple means tests in Table A.3 in the Appendix. Moreover, the marginal effects of own-success probability p_A (models (IV) and (V) in Table 4) is positive and significant ($p < 0.10$). These effects imply that player A is about 1.5% more likely to share an idea in Round 1 in response to a 10 percentage-point increase in her own-success probability. The increase of the real value of information exchange is 1.1% (an elasticity of 0.07), and thus roughly a third of the effect of an increase in p_B.

The effect of a player's own-success probability, summarized in Hypothesis 5, is an overall effect, and the reported means tests results are for this overall effect (since we cannot keep expected intentions constant). Our regression results, however, allow us to separately report the direct as well as the indirect effect. The marginal effect of p_A in model (IV) is the overall effect, whereas the effect of p_A in model (V), controling for expected intentions $\tilde{\sigma}_B$, is the isolated direct effect. A comparison of these two suggests that, if there is an indirect effect of p_A on player A's sharing, operating through player A's expectations of player B's intentions, then this effect is, at best, small. We present additional evidence for this in Table 5, where we evaluate the (direct) marginal effect of p_A at different levels of expected intentions. We do not find a significant difference in the coefficients for a player's own-success probability—suggesting that the direct effect is stronger than an indirect effect.

Our preliminary findings of a small indirect effect imply that higher-order beliefs do not play an important role in how one's own ability affects future behavior. We return to this question when we study the factors that determine player A's expectations further below.

Result 6. *Player B's ability has a stronger effect on player A's decision to share than player A's own ability.*

In Hypothesis 6 we posit that, when holding expected intentions $\tilde{\sigma}_B$ constant, the effect of cross-success p_B is stronger than of own-success p_A. In Figure 2, we see preliminary evidence for this when comparing the fraction of matches terminated by player A in treatment High-Low relative to treatment Low-High. Treatment High-Low (with $p_B < p_A$) exhibits shorter matches, and a larger fraction of those matches are terminated by player A than in treatment Low-High (with $p_B > p_A$). We provide results from simple means tests in Table A.3 in the Appendix. Our regression results in Table 4 paint an analogous picture. We find that the marginal effects of p_B are greater than those of p_A in all specifications.[29]

4.3 Players' Beliefs About Rivals' Intentions

In this section, we address the third question: what elements of players' experiences affect the formation of their expectations about a rival's intended behavior? We first extend our theoretical framework by introducing two *types* of players. This will provide for formal guidance when discussing the empirical results.

4.3.1 Game Dynamics

Suppose a player can be one of two types: sophisticated (S) or non-sophisticated (NS). A sophisticated player is a strategic player as introduced earlier in our model. She shares if

[29] To confirm, we perform a Wald test, which rejects the null that the two effects are the same. In model (IV), equality of the coefficients for p_B and p_A (the overall effect of p_A) can be rejected at the 10% level. In model (V), equality of the coefficients for p_B and p_A (the direct effect of p_A because $\tilde{\sigma}_B$ is controlled for) can be rejected at the 1% level.

the sharing condition (4) is satisfied and conceals otherwise. The non-sophisticated type is not strategic and shares at random. We assume a probability of sharing of $1/2$. We further denote by $\gamma_{m,t}$ player i's belief (in Round t of Match m) that player j is the sophisticated type. For the discussion to follow, we assume the sharing condition (4) is satisfied.

Players update their beliefs about their rivals' type along two dimensions. First, as the process progresses, players can update their beliefs *within* a match m. We apply Bayes' rule for updating beliefs. In $t \geq 2$, upon having observed sharing in $t-1$, a player's belief is equal to

$$\gamma_{m,t} = \frac{\gamma_{m,t-2}}{\gamma_{m,t-2} + (1-\gamma_{m,t-2})\, 1/2} = \frac{2\gamma_{m,t-2}}{1+\gamma_{m,t-2}}. \tag{5}$$

For example, let player A's belief at the outset of Match m be $\gamma_{m,1} = 1/2$. When reaching $t = 3$ (and player B has shared in $t = 2$), her belief is $\gamma_{m,3} = 2/3$. In $t = 5$, her belief is $\gamma_{m,5} = 4/5$. This illustrates that, as a player observes sharing by her rival, she becomes more optimistic about the rival being the sophisticated type (who always shares because the sharing condition is satisfied).

The second dimension of belief updating is *across* matches. For this, it is useful to consider a player's belief at the end of a previous match. We denote this end-of-match belief by $\hat{\gamma}_m$. Because a sophisticated type (by assumption) does not conceal an idea, if a match ends by choice (by the rival), then the player's belief that the rival is a sophisticated type is zero. Otherwise (when the match is terminated by chance or by player i herself), the end-of-match belief is equal to player i's belief in equation (5) from a sequence of T rounds:

$$\hat{\gamma}_m = \begin{cases} 0 & \text{if Match } m \text{ is terminated by choice by the rival} \\ \gamma_{m,T} & \text{if otherwise.} \end{cases} \tag{6}$$

For each match, player i is randomly rematched with a new rival from a pool of rivals. Player i's belief that an independently drawn rival for match m is a sophisticated player is equal to the sample mean $\bar{\gamma}_{m-1}$ of her end-of-match beliefs $\hat{\gamma}_m$ of the previous $m-1$

independent draws, with

$$\bar{\gamma}_{m-1} = \sum_{m'=1}^{m-1} \frac{\hat{\gamma}_{m'}}{m}. \tag{7}$$

The prior belief at the beginning of a match is thus equal to $\gamma_{m,1} = \bar{\gamma}_{m-1}$.

Players who experience termination by choice by one of their rivals will ultimately lower their belief. Consider the example above with $m = 1$ and assume the match terminates by chance in $t = 6$. Player i's end-of-match belief is $\hat{\gamma}_1 = {}^4/_5$. This is also the belief at the outset of Match 2, $\gamma_{2,1} = \bar{\gamma}_1 = \hat{\gamma}_1$. Suppose that in Match 2, rival j terminates the match in some t and player i's end-of-match belief is $\hat{\gamma}_2 = 0$. Her belief at the outset of Match 3 is then $\gamma_{3,1} = \bar{\gamma}_2 = {}^2/_5$, and therefore lower than at the outset of Match 2. Note that the negative effect of termination on beliefs weakens as the number of played matches increases.

A player i's expected intentions $\tilde{\sigma}_j$ of player j are a function of her belief about player j's type. This concept of type allows for a discussion of the dynamics of expectations both *within* a match and *across* matches.

4.3.2 Results

In Table 6, we present results detailing the determinants of player i's expectations $\tilde{\sigma}_j$ in Round t. Unlike before, we now consider both player A's and player B's expectations in *all* rounds. This gives us a total of 1,574 observations (i.e., decisions to share or conceal, and the reported expectations $\tilde{\sigma}_j$). We report the results for tobit models with reported expectations as the dependent variable.

We have shown (Result 5) a positive effect of own ability (p_A) on player A's willingness to share. Moreover, the effect of player A's own ability is weaker than the effect of player B's ability p_B (Result 6). We now ask whether these results (on behavior) translate into correct belief formation. Given correct belief formation, the previous results imply that ability p_i is expected to have a stronger effect on $\tilde{\sigma}_j$ (player i's expectations) than player j's ability p_j. We do not find any support for such higher-order belief formation:

Result 7. *There is no conclusive evidence of rational higher-order belief formation.*

Table 6: Determinants of Expectations

We report the results from tobit models for the determinants of a player i's subjective expectations in t about player j's intentions to share in $t+1$ in all treatments. The dependent variable is $\tilde{\sigma}_j \in [0, 100]$ in a given round t of a match. *Cross* p_j is player j's cross-success probability; *Own* p_i is player i's own-success probability; *Match* is the match number; *Round* is the round number, t, in a given match; *Other Terminated* is a dummy variable = 1 if player i has previously had a match partner (either as player i or player j) who terminated that match by choice (i.e., concealed an idea), and = 0 otherwise; *Own Terminated* is a dummy variable = 1 if player i has previously terminated a match by choice (i.e., concealed an idea), either as player i or as player j, and = 0 otherwise; *Other* × *Own Terminated* is an interaction term. Both *Other Terminated* and *Own Terminated* are = 0 in the very first match. *Subject Dummies* indicates whether subject dummies are included to control for subject fixed effects. The number of observations is the total number of decisions by player i in all t. The left-censoring limit for the tobit model is 0; the right-censoring limit is 100. We report standard errors in parentheses.

	Dependent variable: Player i's subjective expectations $\tilde{\sigma}_j \in [0, 100]$ in a given Round t of a match				
	(I) Tobit	(II) Tobit	(III) Tobit	(IV) Tobit	(V) Tobit
Cross p_j	0.4314 ***	0.4236 ***	0.4263 ***	-0.2115	-0.2579
	(0.0443)	(0.0434)	(0.0437)	(0.2405)	(0.2401)
Own p_i	0.1971 ***	0.2062 ***	0.2073 ***	-0.3673	-0.4127 *
	(0.0453)	(0.0443)	(0.0443)	(0.2404)	(0.2400)
Match	-0.9082 ***	-0.1166	-0.1013	-0.0887	-0.1612
	(0.2391)	(0.3451)	(0.3464)	(0.2936)	(0.2934)
Round	1.6110 ***	1.4826 ***	1.4833 ***	0.4001 **	0.4008 **
	(0.2455)	(0.2402)	(0.2402)	(0.1992)	(0.1987)
Other Terminated		2.9797	2.1690	-5.7261 **	-9.5313 ***
		(2.4343)	(2.8949)	(2.2560)	(2.5397)
Own Terminated		-15.7139 ***	-17.2368 ***	-0.4291	-8.4098 **
		(1.7583)	(3.4304)	(2.2003)	(3.2932)
Other × Own Terminated			2.0060		11.7965 ***
			(3.8786)		(3.6334)
Subject dummies	No	No	No	Yes	Yes
Observations	1574	1574	1574	1574	1574
pseudo R^2	0.0211	0.0272	0.0272	0.1009	0.1017
Log-likelihood	-6420.43	-6380.57	-6380.44	-5896.91	-5891.65

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In the specifications in Table 6 without subject dummies, the positive effect of a rival's success probability on expectations (i.e., the effect of player j's probability p_j on player i's expectations) is stronger than the effect of own-success probability p_A ($p < 0.01$). This pattern refutes rational higher-order belief formation, and we therefore fail to reject the null from Hypothesis 7.

Result 7 suggests that players' expectations about their rival's intentions are inconsistent with rational static belief formation of higher order. We next study dynamic belief formation (i.e., updating) of the first order. We first consider a player's updating of *posterior* beliefs (i.e., $\gamma_{m,t}$) *within* a match and study the evolution of a player i's expectations about player j's intentions as their active match continues (Result 8). We then present results on a player's updating of *prior* beliefs (i.e., $\gamma_{m,1}$) *across* matches. We focus on a player's experience in earlier matches (with other rivals) and discuss the effect of that experience on a player's actions and beliefs (Result 9).

Result 8. *A player's expectations about a rival's intentions increase across the length of the interaction.*

This result is consistent with the updating of posterior beliefs (about the other player's intentions through its type) within a given match m. In Table 6, we show the effect of the number of rounds played on player i's expectations $\tilde{\sigma}_j$ in any t of that match. Given that player j has shared in Round $t-1$, player i updates her posterior beliefs $\gamma_{m,t}$ in Round t and is more optimistic about player j's intentions in $t+1$. The result continues to hold when we add subject i fixed effects to the specifications.

One implication of Result 8 is that the history of a match has an effect on a player's decision to share through her expectations of her rival's intentions. This confirms our earlier approach for Results 1 through 6 in which we restricted our analysis to player A's decision and expectations in Round 1. We take the same approach for our last result below. For Result 9, we turn to a player's updating of prior beliefs ($\gamma_{m,1}$) across matches. We consider the effect of past experience on a player's behavior and examine how past experience (in earlier matches) affects that player's beliefs about her rival's intentions.

Result 9. *Negative past experience (both self-inflicted and by a rival) reduces a player's willingness to share new information.*

In Table 7, we present results from probit regressions. The dependent variable is equal to 1 if player A shares in Round 1 and equal to zero otherwise. We use model (V) from Table 4 and capture past experience using two dummy variables. *Other Terminated* is equal to 1 if player A in a previous match had a rival (either as player A or as player B) who terminated that specific match by choice. Likewise, *Own Terminated* is equal to 1 if player A terminated a previous match by choice, either as player A or as player B. In models (IV) through (IX) of Table 7, we use the subsample of players A who vary their decisions across matches. We find that the effects of our reciprocity variables p_B and $\tilde{\sigma}_B$, as well as p_A, are robust in models (I), (II), and (III) (i.e., the full sample) to the inclusion of past experience.

The effects of *Other Terminated* and *Own Terminated* suggest that past experience plays an important role for player A's decision to share. For example, in model (I), if player A in an earlier match faced a rival who terminated the match by concealing an idea, then player A is 8.4 percentage points less likely to share in Round 1 of the given match.[30]

In model (II), if player A herself terminated by choice in an earlier match, then player A is 26.3 percentage points less likely to share in Round 1 of the given match. One possible explanation for this result is that a player's past action in fact captures the player's own type and thus her propensity to conceal instead of to share an idea. Models (VII) and (IX), in which we control for subject-fixed effects, support this explanation. The effects of *Own Terminated* in the models without the subject-fixed effects are stronger than in the models with the fixed effects.[31] If *Own Terminated* were to capture only a subject-fixed effect, then the marginal effects would be nil in these specifications. However, we still obtain a significant effect of *Own Terminated* in model (IX). A possible explanation for this is a player revising her own prior beliefs about the rival's type through an effect analogous to "self-projection" in which a subject "project[s] her known behavior to guess others' behavior" (Lévy-Garboua

[30]The unconditional mean of player A sharing in Round 1 is 72.0% (for the full sample) and 57.4% (for the sample with match-variant behavior in models (IV) through (IX)). The effect of *Other Terminated* in the models with the reduced sample is stronger because the sample includes only player As who have changed behavior at some point during the experiment.

[31]To allow for direct comparison, we use the reduced sample in models (IV), (V), and (VI) without the subject-fixed effects.

Table 7: Effect of Past Experience on Sharing

We report the results from probit models for the effect of player A's previous experience in model (V) in Table 4. The dependent variable is a dummy variable = 1 if player A shares in Round 1, and = 0 otherwise. Player A's expectations of receiving information in return are captured by *Cross* p_B (player B's cross success probability) and *Expect.* $\tilde{\sigma}_B$ (player A's expectations that player B will share in Round 2). *Own* p_A is player A's own success probability. *Other Terminated* is a dummy variable = 1 if player A previously had a match partner (either as player A or player B) who terminated their match by choice (i.e., concealed an idea), and = 0 otherwise; *Own Terminated* is a dummy variable = 1 if player A previously terminated a match by choice (i.e., concealed an idea), either as player A or player B, and = 0 otherwise. Both *Other Terminated* and *Own Terminated* are = 0 in the very first match. *Subject Dummies* indicates whether or not subject dummies are included to control for subject-specific effects. For models (IV) through (IX), a reduced sample with player A who exhibit varying decisions across matches is considered, implying that 43.9% of observations are dropped (69.7% of observations in treatment High, 25.8% in Low, 55.3% in Low-High, and 31.8% in High-Low). The number of observations is the number of Round 1 decisions by player A. Reported marginal effects in column ME are average marginal effects; reported ME for dummy variables *Other Terminated* and *Own Terminated* are for a discrete change from 0 to 1. We report standard errors in parentheses.

	Dependent variable = 1 if player A shares in Round 1 and = 0 otherwise								
	(I) ME	(II) ME	(III) ME	(IV) ME	(V) ME	(VI) ME	(VII) ME	(VIII) ME	(IX) ME
Cross p_B	0.0035 ***	0.0036 ***	0.0036 ***	-0.0002	0.0015	0.0011	-0.0088	-0.0054	-0.0077
	(0.0008)	(0.0007)	(0.0007)	(0.0013)	(0.0013)	(0.0013)	(0.0056)	(0.0058)	(0.0056)
Expect. $\tilde{\sigma}_B$	0.0055 ***	0.0045 ***	0.0045 ***	0.0063 ***	0.0058 ***	0.0058 ***	0.0062 ***	0.0065 ***	0.0060 ***
	(0.0004)	(0.0004)	(0.0004)	(0.0006)	(0.0006)	(0.0006)	(0.0008)	(0.0008)	(0.0008)
Own p_A	0.0013	0.0014 *	0.0014 *	-0.0015	-0.0008	-0.0010	-0.0008	0.0013	-0.0001
	(0.0008)	(0.0007)	(0.0007)	(0.0012)	(0.0012)	(0.0012)	(0.0057)	(0.0060)	(0.0059)
Other Terminated	-0.0837 **		0.0409	-0.2005 ***		-0.0810	-0.2875 ***		-0.2142 ***
	(0.0361)		(0.0371)	(0.0494)		(0.0571)	(0.0573)		(0.0691)
Own Terminated		-0.2628 ***	-0.2782 ***		-0.2738 ***	-0.2332 ***		-0.2144 ***	-0.1092 *
		(0.0271)	(0.0305)		(0.0449)	(0.0537)		(0.0457)	(0.0560)
Subject Dummies	No	No	No	No	No	No	Yes	Yes	Yes
Observations	578	578	578	324	324	324	324	324	324
pseudo R^2	0.2377	0.3386	0.3404	0.2012	0.2351	0.2395	0.3865	0.3733	0.3949
Log-likelihood	-261.36	-226.77	-226.17	-176.55	-169.06	-168.07	-135.60	-138.51	-133.73

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

et al., 2006:574). This means that, when player A observes herself concealing an idea, the general incentives of sharing and concealing become more salient, resulting in less optimistic expectations about player B's intentions in a given match.

We argue that past experience enters a player's decision to share information through her expectations of the other player's intentions. Controlling for player i fixed effects in models (IV) and (V) of Table 6, we indeed see a negative effect of *Other Terminated* on player i's beliefs $\tilde{\sigma}_j$. Moreover, the positive interaction effect in model (V) suggests that the two types of past experience (self-inflicted and by the rival) are not cumulative. The effect of *Other Terminated* is stronger when player i herself has not terminated an earlier match. Similarly, player i adjusts her expectations of player j's behavior downward in response to *Own Terminated* only if she has not already seen a rival terminate a match.

The negative effect of *Other Terminated* on a player's expectations about her rival is consistent with updating beliefs $\bar{\gamma}_{m-1}$ about the rival's type—as formally introduced above. A negative past experience makes the player more pessimistic about the pool of possible future rivals. In the Online Appendix, we further present results for the effect of past experience separately for early and late matches. We predict that the negative effect of *Other Terminated* on a player's beliefs is weaker as the number of completed matches increases. Indeed, our estimation results for the late-match sample suggest such a null effect.

5 Concluding Remarks

Competitors frequently share information with others in the hopes of receiving new information in return. This type of (often informal) information exchange has been documented in academia, finance, and numerous other industries. A number of studies have argued that individuals and firms are more likely to share information when the costs of sharing ("current loss of competitiveness") are offset by its future benefits ("potential future reciprocity") (Häussler et al., 2014).

We experimentally test whether a *feedback loop* that captures this tradeoff can shed light on the role of two central components of an individual's or a firm's expectations of reciprocity: a recipient's *ability* to share new information and a recipient's *intention* to share new information. We use an indeterminate-horizon centipede game to establish a feedback loop in the laboratory and find that an individual's expectations about a recipient's intent to reciprocate are a more important determinant for the exchange of information than the recipient's ability to share. We further show that prior negative experience (in episodes of information exchange) lingers and has a negative effect on individuals' and firms' incentives to initiate and maintain such exchange in the future.

The basic tradeoff that individuals face when deciding to share private information applies to both competitors across organizations and team members within organizations. Thus, our results inform the literature on collective invention (Allen, 1983; Powell and Giannella, 2010) and innovation networks (Schilling and Phelps, 2007; von Hippel, 2007), as well as suggest implications for the design of formal and informal platforms of knowledge exchange. We find that reducing the strategic uncertainty about the behavior of linked members in one's network is more important for the diffusion of knowledge than reducing the environmental uncertainty about other members' productive potential.[32] This means that the presence of reliable members is more conducive to a sustainable information exchange than the presence of members with high productive potential.

The implications of our results are consistent with observed practice in many innovation-focused organizations. For instance, the Internet Engineering Task Force (IETF), a standard-development body that develops, maintains, and promotes technology standards and protocols that are the foundation of the internet, conducts a large part of its development work remotely via email listservs. But, importantly, it also organizes three annual meetings where contributors meet in person. The atmosphere has been described as fun and more social than at other standards bodies or computer industry conferences (Hoffman, 2018). Because the IETF has no formal membership requirements, these face-to-face meetings are a way for the organization to facilitate the building of stronger ties among its contributors and thus foster better development.

[32] A player with strong links is more likely to expect another node to reciprocate with new information (implying higher expectations of the *intention* to share). Conversely, a node with more links is more likely to face a node that is able to generate new information to share (implying higher expectations of the *ability* to share).

References

AGARWAL, V. AND N. Y. NAIK (2000): "Multi-Period Performance Persistence Analysis of Hedge Funds," *Journal of Financial and Quantitative Analysis*, 35, 327–342.

ALLEN, R. C. (1983): "Collective Invention," *Journal of Economic Behavior & Organization*, 4, 1–24.

ANDREI, D. AND J. CUJEAN (2017): "Information Percolation, Momentum, and Reversal," *Journal of Financial Economics*, 123, 617–645.

APPLEYARD, M. M. (1996): "How Does Knowledge Flow? Interfirm Patterns in the Semiconductor Industry," *Strategic Management Journal*, 17, 137–154.

ARROW, K. J. (1962): "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity*, ed. by R. R. Nelson, Cambridge, Mass.: National Bureau of Economic Research, 609–625.

AUGENBLICK, N. AND A. BODOH-CREED (2018): "To Reveal or Not to Reveal: Privacy Preferences and Economic Frictions," *Games and Economic Behavior*, 110, 318–329.

BERG, J., J. DICKHAUT, AND K. MCCABE (1995): "Trust, Reciprocity, and Social History," *Games and Economic Behavior*, 10, 122–142.

BINMORE, K. (1987): "Modeling Rational Players: Part 1," *Economics and Philosophy*, 3, 179–214.

BORNSTEIN, G., T. KUGLER, AND A. ZIEGELMEYER (2004): "Individual and Group Decisions in the Centipede Game: Are Groups More "Rational" Players?" *Journal of Experimental Social Psychology*, 40, 599–605.

BOTELHO, T. L. (2018): "Here's an Opportunity: Knowledge Sharing among Competitors as a Response to Buy-in Uncertainty," *Organization Science*, 29, 1033–1055.

BOUTY, I. (2000): "Interpersonal and Interaction Influences on Informal Resource Exchanges between R&D Researchers across Organizational Boundaries ," *Academy of Management Journal*, 43, 50–65.

BROWN, J. R., Z. IVKOVIĆ, P. A. SMITH, AND S. WEISBENNER (2008): "Neighbors Matter: Causal Community Effects and Stock Market Participation," *Journal of Finance*, 63, 1509–1531.

COHEN, L., A. FRAZZINI, AND C. MALLOY (2008): "The Small World of Investing: Board Connections and Mutual Fund Returns," *Journal of Political Economy*, 116, 951–979.

COLLA, P. AND A. MELE (2010): "Information Linkages and Correlated Trading," *Review of Financial Studies*, 23, 203–246.

COOPER, R. W., D. V. DEJONG, R. FORSYTHE, AND T. W. ROSS (1990): "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, 80, 218–233.

CRAWFORD, S. S., W. R. GRAY, AND A. E. KERN (2017): "Why Do Fund Managers Identify and Share Profitable Ideas?" *Journal of Financial and Quantitative Analysis*, 52, 1903–1926.

Duffie, D. and G. Manso (2007): "Information Percolation in Large Markets," *American Economic Review*, 97, 203–209.

Dziuda, W. and R. Gradwohl (2015): "Achieving Coordination under Privacy Concerns," *American Economic Journal: Microeconomics*, 7, 142–173.

Fehr, E., G. Kirchsteiger, and A. Riedl (1993): "Does Fairness Prevent Market Clearing? An Experimental Investigation," *Quarterly Journal of Economics*, 108, 437–459.

Fehr, E. and K. M. Schmidt (1999): "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 114, 817–868.

——— (2006): "The Economics of Fairness, Reciprocity and Altruism: Experimental Evidence and New Theories," in *Handbook of the Economics of Giving, Altruism and Reciprocity*, ed. by S.-C. Kolm and J. M. Ythier, Amsterdam, The Netherlands: North-Holland, vol. 1, chap. 8, 615–691.

Feng, L. and M. S. Seasholes (2004): "Correlated Trading and Location," *Journal of Finance*, 59, 2117–2144.

Fischbacher, U. (2007): "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments," *Experimental Economics*, 10, 171–178.

Gächter, S., G. von Krogh, and S. Haeflinger (2010): "Initiating Private-Collective Innovation: The Fragility of Knowledge Sharing," *Research Policy*, 39, 893–906.

Ganglmair, B. and E. Tarantino (2014): "Conversation with Secrets," *RAND Journal of Economics*, 45, 273–302.

Gerber, A. and P. C. Wichardt (2010): "Iterated Reasoning and Welfare-Enhancing Instruments in the Centipede Game," *Journal of Economic Behavior & Organization*, 74, 123–136.

Goetzmann, W. N., J. E. Ingersoll, and S. A. Ross (2003): "High-Water Marks and Hedge Fund Management Contracts," *Journal of Finance*, 58, 1685–1717.

Grinblatt, M. and M. Keloharju (2001): "How Distance, Language, and Culture Influence Stockholdings and Trades," *Journal of Finance*, 56, 1053–1073.

Guttman, I., I. Kremer, and A. Skrzypacz (2014): "Not Only What but Also When: A Theory of Dynamic Voluntary Disclosure," *American Economic Review*, 104, 2400–2420.

Harsanyi, J. and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*, Cambridge, Mass.: MIT Press.

Hau, H. (2001): "Location Matters: An Examination of Trading Profits," *Journal of Finance*, 56, 1959–1983.

Häussler, C. (2011): "Information-Sharing in Academia and the Industry: A Comparative Study," *Research Policy*, 40, 105–122.

Häussler, C., L. Jiang, J. Thursby, and M. C. Thursby (2014): "Specific and General Information Sharing Among Academic Scientists," *Research Policy*, 43, 465–475.

Hellmann, T. F. and E. C. Perotti (2011): "The Circulation of Ideas in Firms and Markets," *Management Science*, 57, 1813–1826.

Hoffman, P. (2018): "The Tao of IETF: A Novice's Guide to the Internet Engineering Task Force," Retrieved from https://www.ietf.org/tao.html (June 14, 2019).

Holt, C. A. and S. K. Laury (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92, 1644–1655.

Hong, H., J. D. Kubik, and J. C. Stein (2004): "Social Interaction and Stock-Market Participation," *Journal of Finance*, 59, 137–163.

Ingram, P. and P. W. Roberts (2000): "Friendships among Competitors in the Sydney Hotel Industry," *American Journal of Sociology*, 106, 387–423.

Ivković, Z. and S. Weisbenner (2005): "Local Does as Local Is: Information Content of the Geography of Individual Investors' Common Stock Investments," *Journal of Finance*, 60, 267–306.

Kawagoe, T. and H. Takizawa (2012): "Level-k Analysis of Experimental Centipede Games," *Journal of Economic Behavior & Organization*, 82, 548–566.

Levitt, S. D., J. A. List, and S. E. Sadoff (2011): "Checkmate: Exploring Backward Induction Among Chess Players," *American Economic Review*, 101, 975–990.

Lévy-Garboua, L., C. Meidinger, and B. Rapoport (2006): "The Formation of Social Preferences: Some Lessons from Psychology and Biology," in *Handbook of the Economics of Giving, Altruism and Reciprocity*, ed. by S.-C. Kolm and J. M. Ythier, Amsterdam, The Netherlands: North-Holland, vol. 1, chap. 7, 545–613.

Manela, A. (2014): "The Value of Diffusing Information," *Journal of Financial Economics*, 111, 181–199.

McCrary, S. A. (2002): *How to Create and Manage a Hedge Fund: A Professional's Guide*, Hoboken, N.J.: John Wiley & Sons.

McKelvey, R. D. and T. R. Palfrey (1992): "An Experimental Study of the Centipede Game," *Econometrica*, 60, 803–836.

Nagel, R. and F. F. Tang (1998): "Experimental Results on the Centipede Game in Normal Form: An Investigation on Learning," *Journal of Mathematical Psychology*, 42, 356–384.

Ortmann, A., J. Fitzgerald, and C. Boeing (2000): "Trust, Reciprocity, and Social History: A Re-examination," *Experimental Economics*, 3, 81–100.

Pool, V. K., N. Stoffman, and S. E. Yonker (2015): "The People in Your Neighborhood: Social Interactions and Mutual Fund Portfolios," *Journal of Finance*, 70, 2679–2732.

Powell, W. W. and E. Giannella (2010): "Collective Invention and Inventor Networks," in *Handbook of the Economics of Innovation*, ed. by B. H. Hall and N. Rosenberg, Amsterdam, The Netherlands: North-Holland, vol. 1, chap. 13, 575–605.

Rantala, V. (forthcoming): "How Do Investment Ideas Spread through Social Interaction? Evidence from a Ponzi Scheme," *Journal of Finance*.

Rosenthal, R. W. (1981): "Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox," *Journal of Economic Theory*, 25, 92–100.

Schilling, M. A. and C. C. Phelps (2007): "Interfirm Collaboration Networks: The Impact of Large-Scale Network Structure on Firm Innovation," *Management Science*, 53, 1113–1126.

Schrader, S. (1991): "Informal Technology Transfer Between Firms: Cooperation Through Information Trading," *Research Policy*, 20, 153–170.

Shiller, R. J. (2000): *Irrational Exuberance*, Princeton, N.J.: Princeton University Press.

Shiller, R. J. and J. Pound (1989): "Survey Evidence on Diffusion of Interest and Information Among Investors," *Journal of Economic Behavior & Organization*, 12, 47–66.

Shive, S. (2010): "An Epidemic Model of Investor Behavior," *Journal of Financial and Quantitative Analysis*, 45, 169–198.

Stein, J. C. (2008): "Conversations among Competitors," *American Economic Review*, 98, 2150–2162.

Trautmann, S. T. and G. van de Kuilen (2015): "Belief Elicitation: A Horse Race among Truth Serums," *Economic Journal*, 125, 2116–2135.

von Hippel, E. (1987): "Cooperation Between Rivals: Informal Know-How Trading," *Research Policy*, 16, 291–302.

——— (2007): "Horizontal Innovation Networks–by and for Users," *Industrial and Corporate Change*, 16, 293–315.

von Hippel, E. and G. von Krogh (2006): "Free Revealing and the Private-Collective Model for Innovation Incentives," *R&D Management*, 36, 295–306.

Zvilichovsky, D., Y. Inbar, and O. Barzilay (2015): "Playing Both Sides of the Market: Success and Reciprocity on Crowdfunding Platforms," Unpublished manuscript, available at https://ssrn.com/abstract=2304101.

Appendix

A Model: Extended Presentation with Formal Proofs

We first summarize the formal results in Propositions A.1 through A.4 and then provide our formal predictions of the model.

A.1 Pure-Strategy and Mixed-Strategy Equilibria

We can now summarize the pure-strategy equilibria of the game of word-of-mouth communication as follows.

Proposition A.1 (Equilibria in Pure Strategies)**.** *Suppose a player i, believing its rival player j will share an idea with certainty so that $\tilde{\sigma}_j = 1$, has non-negative net benefits from sharing, $\tilde{\phi}_i(1) \geq 0$. The game of word-of-mouth communication has two pure-strategy Nash equilibria: (1) both players never share an idea and (2) both players always share a newly generated idea.*

Proof. We can rewrite the sharing condition (4) (for a given $\tilde{\sigma}_j$) as follows:

$$\tilde{\sigma}_j \geq \frac{\theta}{(1-\theta+\beta p_i)\,\beta p_j}. \tag{A.1}$$

This condition defines player i's best response function, $s_i : [0,1] \to \{\text{share, conceal}\}$. If player j is expected to share with sufficiently high probability, that means, if $\tilde{\sigma}_j$ is sufficiently high, then player i will share. Conversely, if player i expects player j to share a newly generated idea with low probability, then player i will in return choose to conceal her idea and end the conversation:

$$s_i(\tilde{\sigma}_j) = \begin{cases} \text{share} & \text{if } \tilde{\sigma}_j \geq \dfrac{\theta}{(1-\theta+\beta p_i)\,\beta p_j} \\ \text{conceal} & \text{if } \tilde{\sigma}_j < \dfrac{\theta}{(1-\theta+\beta p_i)\,\beta p_j}. \end{cases} \tag{A.2}$$

To show the claims in the Proposition, first note that, in equilibrium, $\tilde{\sigma}_j = \sigma_j$. Recall that we assume time-invariant strategies σ_i for $i = A, B$.

1. Suppose player j always conceals and $\sigma_j = 0$. Then $\tilde{\phi}_i(0) = -\mu\theta < 0$ and player i's sharing condition (4) is violated in all t. As a result, $\sigma_i = 0$. For $\sigma_i = 0$, player j's sharing condition is violated in all t because $\tilde{\phi}_j(0) = -\mu\theta < 0$ so that $\sigma_j = 0$, inducing player i to conceal in all t.

2. This proof immediately follows from Stein (2008). In order for a player i to share, her necessary condition $\tilde{\phi}_i(\tilde{\sigma}_j) \geq 0$ must be satisfied, given player j's strategy σ_j (and player i's beliefs thereof). We first show that if the condition is satisfied for $\sigma_i = 1$ for both $i = A, B$, then both players always share a newly generated idea. We then show that, if at least one of them is violated for $\sigma_i = 1$, neither player i nor player j will ever share a newly generated idea.

- First, observe that if both players always share and $\sigma_i = \sigma_j = 1$, then $\phi_i := \tilde{\phi}_i(1) \geq 0$ for $i = A, B$. If $\phi_i \geq 0$ and player i anticipates (in equilibrium) that player j continues in all $t' > t$ so that $\sigma_j = 1$, then player i continues in any t because her necessary condition $\phi_i \geq 0$ holds. Then $\sigma_i = 1$. If $\phi_j \geq 0$ and player j anticipates (as player i's best response to σ_j) that player i continues in all $t' > t$ so that $\sigma_i = 1$, then player j continues in any t because her necessary condition $\phi_j \geq 0$ holds. Then $\sigma_j = 1$.
- Now, suppose that $\phi_j \geq 0$ but $\phi_i < 0$. This implies that $\tilde{\phi}_i(1) < 0$, and $\tilde{\phi}_i(\sigma_j) < 0$ for all σ_j because $\tilde{\phi}_i(\sigma_j)$ increases in σ_j (see the proof of Proposition A.2 below). This means that for any strategy σ_j, player i conceals an idea in t. Anticipating this, player j expects in $t-1$ payoffs of $\text{share}_{j,t-1} = \mu(1-\theta)(1-\beta^{t-1})$ when she shares and $\text{conceal}_{t-1} = \mu\left[(1-\beta^{t-1}) - \theta(1-\beta^{t-2})\right]$ when she conceals. She decides to conceal because $\text{share}_{j,t-1} < \text{conceal}_{t-1}$ as $1-\beta^{t-1} > 1-\beta^{t-2}$. Because player i conceals in any t, player j will respond by concealing in any $t-1$. The game therefore unravels and player A conceals in $t = 1$. The analogous argument applies to the case of $\phi_i \geq 0$ but $\phi_j < 0$. Q.E.D.

In Proposition A.1, we characterize two pure-strategy equilibria. The game also has a mixed-strategy equilibrium. We characterize the equilibrium in the following proposition:

Proposition A.2 (Equilibrium in Mixed Strategies)**.** *Let $\tilde{\phi}_i(1) \geq 0$ for $i = A, B$. The communication game has a mixed-strategy equilibrium in which player $i = A, B$, $i \neq j$, shares newly arrived ideas with probability*

$$\sigma_i^* = \frac{\theta}{(1-\theta+\beta p_j)\,\beta p_i}. \tag{A.3}$$

Proof. In equilibrium, a player's expectations about the rival's strategy are correct, $\tilde{\sigma}_j = \sigma_j$. Moreover, in a mixed-strategy equilibrium, player i chooses a mixed strategy if she is indifferent between *share* and *conceal*. By the expression in (A.2), player i is indifferent if $\tilde{\sigma}_j = \sigma_j = \frac{\theta}{(1-\theta+\beta p_i)\beta p_j}$, and therefore indifferent between the pure actions and any mixture $\sigma_i \in [0,1]$. If $\sigma_i = \frac{\theta}{(1-\theta+\beta p_j)\beta p_i}$, then player j is indifferent and willing to play a strategy σ_j as above. Q.E.D.

This mixed-strategy equilibrium is payoff-dominated by the sharing equilibrium in Proposition A.1, and it payoff-dominates the nonsharing equilibrium in that proposition. The mixed strategies are time-invariant. We calculate the expected duration of our game of word-of-mouth communication and summarize in the following proposition.

Proposition A.3. *The expected duration of word-of-mouth communication is*

$$1 + \frac{\sigma_A p_B}{1 - \sigma_A \sigma_B p_A p_B}.$$

It is finite if σ_i and p_i such that $\sigma_A \sigma_B p_A p_B < 1$. The effect of p_B on this expected duration is stronger than the effect of p_A if, and only if, σ_i and p_B such that $\sigma_A \sigma_B p_B < 1$.

Proof. To determine the expected duration of communication, we determine the probabilities δ_t that the game ends in a stage t.

- The game ends in Round 1 when (i) player A conceals or (ii) when player A shares and player B fails. The probability of (i) or (ii) is

$$\delta_1 = 1 - \sigma_A + \sigma_A (1 - p_B) = 1 - \sigma_A p_B.$$

- The game ends in Round 2 when (i) player A shares, player B is successful, and player B conceals; or (ii) player A shares, player B is successful, player B shares, and player A fails. The probability of (i) or (ii) is

$$\begin{aligned} \delta_2 &= \sigma_A p_B (1 - \sigma_B) + \sigma_A p_B \sigma_B (1 - p_A) \\ &= \sigma_A p_B (1 - \sigma_B p_A). \end{aligned}$$

- The game ends in Round 3 when (i) player A shares, player B is successful, player B shares, player A is successful, and player A conceals; or (ii) player A shares, player B is successful, player B shares, player A is successful, player A shares, and player B fails. The probability of (i) or (ii) is

$$\begin{aligned} \delta_3 &= \sigma_A p_B \sigma_B p_A (1 - \sigma_A) + \sigma_A p_B \sigma_B p_A \sigma_A (1 - p_B) \\ &= \sigma_A p_B \sigma_B p_A (1 - \sigma_A p_B). \end{aligned}$$

- The probability that the game ends in Round 4 is $\delta_4 = (\sigma_A p_B)^2 \sigma_B p_A (1 - \sigma_B p_A)$; the probability that the game ends in Round 5 is $\delta_5 = (\sigma_A p_B)^2 (\sigma_B p_A)^2 (1 - \sigma_A p_B)$; the probability that the game ends in Round 6 is $\delta_6 = (\sigma_A p_B)^3 (\sigma_B p_A)^2 (1 - \sigma_B p_A)$; the probability that the game ends in Round 7 is $\delta_7 = (\sigma_A p_B)^3 (\sigma_B p_A)^3 (1 - \sigma_A p_B)$; and so forth.

The expected duration of word-of-mouth communication (i.e., the expected round in which it ends) is

$$\begin{aligned} D &= \sum_{q=0}^{\infty} \delta_{q+1} (q+1) \\ &= \sum_{q=0}^{\infty} (\sigma_A p_B)^q (\sigma_B p_A)^q [(1 - \sigma_A p_B)(1+q) + \sigma_A p_B (1 - \sigma_B p_A)(2+q)] \\ &= 1 + \frac{\sigma_A p_B}{1 - \sigma_A \sigma_B p_A p_B}. \end{aligned} \tag{A.4}$$

The derivative of the last expression, D, with respect to p_A is

$$\frac{\partial D}{\partial p_A} = \frac{p_B^2 \sigma_A^2 \sigma_B}{(1 - \sigma_A \sigma_B p_A p_B)^2} > 0.$$

The derivative of D with respect to p_B is

$$\frac{\partial D}{\partial p_B} = \frac{\sigma_A}{(1 - \sigma_A \sigma_B p_A p_B)^2} > 0.$$

At last,

$$\frac{\partial D}{\partial p_B} > \frac{\partial D}{\partial p_A} \iff \sigma_A \sigma_B p_B^2 < 1, \tag{A.5}$$

implying that the effect of player B's success probability is stronger than player A's success probability if and only if σ_i and p_i such that $\sigma_A \sigma_B p_B < 1$ Q.E.D.

Our predictions of the model discussed in the main text use payoff-dominance as criterion for equilibrium selection when both a sharing and nonsharing equilibrium exist (for $\tilde{\phi}_i$). The following proposition summarizes the comparative statics of $\tilde{\phi}_i \geq 0$ with respect to our parameters of interest.

Proposition A.4 (Comparative Statics)**.** *Suppose a sharing equilibrium with $\tilde{\phi}_i(1) \geq 0$ for $i = A, B$ exists. Player i's net benefits from sharing $\tilde{\phi}_i$ have the following properties: (1) $\tilde{\phi}_i$ is increasing in p_j, $\tilde{\sigma}_j$, and $\pi_j = p_j \tilde{\sigma}_j$, (2) $\tilde{\phi}_i$ is increasing in p_i, and (3) the marginal effect of p_j on $\tilde{\phi}_i$ is larger than the effect of p_j for sufficiently large values of p_i.*

Proof. The first derivatives of $\tilde{\phi}_i$ with respect to π_j, p_j, and $\tilde{\sigma}_j$ are

$$\begin{aligned} \frac{\partial \tilde{\phi}_i}{\partial p_j} &= \frac{\beta (1 + \beta p_i) \tilde{\sigma}_j}{(1 + \beta \tilde{\sigma}_j p_j)^2} > 0, \\ \frac{\partial \tilde{\phi}_i}{\partial \tilde{\sigma}_j} &= \frac{\beta (1 + \beta p_i) p_j}{(1 + \beta \tilde{\sigma}_j p_j)^2} > 0, \quad \text{and} \\ \frac{\partial \tilde{\phi}_i}{\partial p_i} &= \frac{\beta^2 \tilde{\sigma}_j p_j}{1 + \beta \sigma_j p_j} > 0. \end{aligned}$$

From the cross-probability effect of p_j and the own-probability effect of p_i, we can see that

$$\frac{\partial \tilde{\phi}_i}{\partial p_j} > \frac{\partial \tilde{\phi}_i}{\partial p_i} \iff \frac{1 + \beta p_i}{1 + \beta p_j \tilde{\sigma}_j} > \beta p_j. \tag{A.6}$$

This means that the effect of player j's ability p_j is stronger than the effect of player i's own ability p_i if p_i is not too low. Q.E.D.

A.2 Predictions from the Model

For our first prediction, we expect the parties to play the sharing strategy (and thus coordinate on the payoff-dominant sharing equilibrium) more often when the net benefits of sharing $\tilde{\phi}_i$ are higher.

Prediction 1. *Players are expected to share and coordinate on the sharing equilibrium more often when the net benefits from sharing $\tilde{\phi}_i$ are higher.* [Hypothesis 1]

By our payoff-dominance argument, anything that relaxes the condition in equation (4) increases the net benefits from sharing and will induce players to choose the sharing strategy more often.

Prediction 2. *Player i is expected to share a new idea more often when she expects to receive information in return with higher probability.* [Hypotheses 2 and 3]

Unless $p_j = \tilde{\sigma}_j$, we expect the two components of reciprocity (ability and expected intentions) to have different effects on player i's behavior. More specifically, if $p_j < \tilde{\sigma}_j$, the effect of ability p_j is stronger than that of expected intentions $\tilde{\sigma}_j$. If, instead, $p_j > \tilde{\sigma}_j$, this pattern is reversed.[33] Moreover, the larger the difference between p_j and $\tilde{\sigma}_j$ (or the larger their ratio $p_j/\tilde{\sigma}_j$), the bigger this differential effect—where the difference of marginal effects of p_j and $\tilde{\sigma}_j$ is positive for $p_j/\tilde{\sigma}_j < 1$ and negative if $p_j/\tilde{\sigma}_j > 1$. The relative impact of environmental uncertainty (captured by p_j) and strategic uncertainty (captured by $\tilde{\sigma}_j$) on player i's behavior is thus determined by how much they each contribute to the probability of reciprocity. This gives rise to our third prediction.

Prediction 3. *The difference of the marginal effects of p_j and $\tilde{\sigma}_j$ on player i's willingness to share is negative if $p_j/\tilde{\sigma}_j < 1$, positive if $p_j/\tilde{\sigma}_j > 1$, and zero otherwise.* [Hypothesis 4]

For our next prediction, recall that, in the symmetric baseline model, a player's and a rival's abilities have the same effect on player i's decision to share (because $p_i = p_j = p$). We can disentangle one from the other. We predict that a player's own ability to send information in $t + 2$ has a positive effect on her decision to send information in t.

Prediction 4. *Holding $\tilde{\sigma}_j$ constant, player i's willingness to share increases with her own ability, p_i.* [Hypothesis 5]

Holding expectations $\tilde{\sigma}_j$ constant, a player's own ability p_i is expected to be weaker than the effect of the other player's cross ability. This is because the cross-ability effect is immediate (in $t + 1$), whereas the own-ability effect is delayed (in $t + 2$).

Prediction 5. *Holding $\tilde{\sigma}_j$ constant, player i's own ability p_i has a weaker effect on her decision to share than player j's cross ability p_j.* [Hypothesis 6]

The effect of own ability p_i on player i's decision to share is two-fold. First, a player's own ability increases the value of ongoing information exchange because, with a higher p_i, the exchange is expected to last longer and therefore more ideas are generated and shared. We refer to this effect as direct effect. Second, an indirect effect of own ability on the player's decision acts through her expectations of the other player's intentions. In Prediction 2, we predict that p_j has a positive effect on player i's decision. We therefore expect a positive effect of p_i on player j's decision. Player i, forming beliefs about player j's intentions to share, will anticipate this effect. As a consequence, a player's own ability is expected to have a positive effect on her expectations of player j's intentions, and p_i has thus an *indirect* or second-order effect on her decision to share.

[33] We can see this from comparative statics. The first derivative of $\tilde{\phi}_i(\tilde{\sigma}_j)$ with respect to p_j is $\frac{d\tilde{\phi}_i}{dp_j} = \frac{(1+\beta p_i)\beta\tilde{\sigma}_j}{(1+\beta\pi_j)^2}$. The first derivative of $\tilde{\phi}_i(\tilde{\sigma}_j)$ with respect to $\tilde{\sigma}_j$ is $\frac{d\tilde{\phi}_i}{d\tilde{\sigma}_j} = \frac{(1+\beta p_i)\beta p_j}{(1+\beta\pi_j)^2}$. We have $\frac{d\tilde{\phi}_i}{dp_j} > \frac{d\tilde{\phi}_i}{d\tilde{\sigma}_j}$ if, and only if, $p_j < \tilde{\sigma}_j$.

Prediction 6. *Player i's own ability, p_i, has a stronger effect on her expectations about player j's intentions than player j's cross ability, p_j.* [Hypothesis 7]

These predictions address our three research questions as laid out in the introduction. We translate these predictions into testable hypotheses in Section 3 and present answers to our questions in Section 4.

B Robustness Results

Our results for the expectations of reciprocity are robust to a set variables capturing trust, fairness, and personal connections, all of which have been associated with increased cooperative or pro-social behavior. We report these results in Table A.1. We provide detailed descriptions and summary statistics for these control variables in Table A.2.

B.1 Personal Connections

Indicators of personal connections or social bonds (i.e., the number of people a participant recognizes in the experimental session, *Acquaintances*, and the number of people in the session a participant considers friends, *Friends*[34]) do not affect our results for the probability of reciprocity (p_B and $\tilde{\sigma}_B$) or a player's own-success probability. Moreover, only *Friends* exhibits a statistically significant effect on player A's sharing in Round 1.

A small number of papers have presented results that suggest that social interactions and peer effects influence stock market participation (Hong et al., 2004) or provide a mechanism through which asset prices incorporate private information (Cohen et al., 2008). To understand how personal connections or social bonds affect word-of-mouth communication, we need to draw a distinction between the effect at the *extensive margin* and at the *intensive margin*. The former describes how players choose to form connections or a network with which to share private information (selection). The latter captures the effect on the willingness to share when a connection or network has already been formed. We find that, given an exchange network (taking the extensive margin as given), the presence of personal connections or social bonds plays little to no role in the player's decision to share an idea. Our results are complementary to Crawford et al. (2017) who also take a social network as given and observe word-of-mouth communication at the intensive margin.

[34]Recall that, by the design of the experiment, subjects did not know with whom they had been grouped. The answers to the above questions, therefore, apply to the session (two groups) rather than to the subject's group.

Table A.1: Robustness Results for the Effects of Ability and Intentions

We report probit results of a set of sensitivity analyses for all four treatments. The dependent variable is a dummy variable = 1 if player A shares in Round 1, and = 0 otherwise. Player A's expectations of reciprocity are captured by *Cross* p_B (player B's cross-success probability) and *Expect.* $\tilde{\sigma}_B$ (player A's expectations that player B will share in Round 2). *Own success:* p_A is player A's own-success probability. Further co-variates are the *Match* number; the number of a participant's *Acquaintances* in the experimental session; the number of a participant's *Friends* in the experimental session; a participant's perception of general *Fairness* and *Trustworthiness* of people (ranging from 1 to 10 with higher numbers indicating more fairness or trustworthiness); and a *Risk Aversion* measure by the *Holt-Laury* risk preference task (ranging from 1 to 10 with higher numbers reflecting higher degrees of risk aversion). See Table A.2 for more detailed definitions and descriptive statistics of these co-variates. We reproduce model (V) from Table 4 with the main results in the first column. The number of observations is the number of Round 1 decisions by player A. Reported marginal effects (ME) are average marginal effects. We report standard errors in parentheses.

	Dependent variable = 1 if player A shares in Round 1 and = 0 otherwise								
	Table 4:(V) ME	(I) ME	(II) ME	(III) ME	(IV) ME	(V) ME	(VI) ME	(VII) ME	(VIII) ME
Cross p_B	0.0035 ***	0.0032 ***	0.0031 ***	0.0032 ***	0.0031 ***		0.0030 ***	0.0029 ***	0.0030 ***
	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0008)		(0.0009)	(0.0009)	(0.0009)
Expect. $\tilde{\sigma}_B$	0.0056 ***	0.0054 ***	0.0053 ***	0.0053 ***	0.0053 ***		0.0052 ***	0.0052 ***	0.0052 ***
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0005)		(0.0005)	(0.0005)	(0.0005)
Own p_A	0.0014 *	0.0011	0.0011	0.0011	0.0012		0.0018 **	0.0011	0.0018 **
	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0009)		(0.0009)	(0.0009)	(0.0009)
Match		-0.0155 ***	-0.0162 ***	-0.0159 ***	-0.0164 ***	-0.0285 ***	-0.0138 ***	-0.0135 ***	-0.0138 ***
		(0.0046)	(0.0046)	(0.0046)	(0.0046)	(0.0052)	(0.0050)	(0.0050)	(0.0050)
Acquaintances			-0.0097		-0.0080	-0.0062	-0.0150		-0.0150
			(0.0072)		(0.0075)	(0.0085)	(0.0104)		(0.0104)
Friends			0.0169 *		0.0152	0.0179	0.0415 ***		0.0415 ***
			(0.0098)		(0.0100)	(0.0116)	(0.0144)		(0.0145)
Fairness				-0.0003	-0.0005	0.0120	0.0015		0.0015
				(0.0078)	(0.0077)	(0.0087)	(0.0080)		(0.0080)
Trustworthiness				0.0075	0.0056	0.0109	-0.0027		-0.0027
				(0.0085)	(0.0087)	(0.0100)	(0.0091)		(0.0091)
Risk Aversion								0.0114	0.0005
								(0.0114)	(0.0119)
Observations	578	578	578	578	578	578	481	481	481
pseudo R^2	0.2299	0.2456	0.2499	0.2473	0.2507	0.0534	0.2774	0.2571	0.2774
Log-likelihood	-265.05	-258.66	-257.20	-258.09	-256.92	-324.58	-202.10	-207.77	-202.10

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.2 Fairness and Trustworthiness

The experimental literature in economics has shown that considerations of fairness of others and trust toward others play an important role in how people make decisions.[35] In order to see the effect of fairness and trust on a player's decision to share a new idea, we control for two variables obtained in an exit survey. First, we survey the participants' perceptions of other people's fairness, *Fairness*; second, we ask for participants' perceptions of other people's trustworthiness, *Trustworthiness*. Again, our main results are robust to the inclusion of these indicators. Moreover, subjects' views of fairness and trustworthiness do not exhibit statistically significant effects on player A's sharing in Round 1. We therefore do not find evidence for an effect of general perceptions of fairness and trustworthiness of others on a player's decision to share private information.

B.3 Risk Aversion

We further find that risk aversion does not drive our main results because the marginal effect of *Risk Aversion* on player A's sharing behavior is not statistically significantly different from zero. We derive our risk-aversion measure from the Holt and Laury (2002) risk-preference tasks; our numbers are consistent with those in Holt and Laury (2002).[36] We take a conservative approach, and for our analyses—utilizing the Holt-Laury risk-preference measure in models (VI) through (VIII)—we use only observations from matches with subjects making consistent choices.

[35]For fairness, see, for instance, Fehr et al. (1993), Fehr and Schmidt (1999), or Fehr and Schmidt (2006). In the context of information exchange, Gächter et al. (2010) argue that knowledge sharing in private-collective innovation (i.e., privately funded public goods innovation) is affected by fairness. For trust, see, for instance, Berg et al. (1995) or Ortmann et al. (2000).

[36]Most subjects are risk averse and made choices between 5, 6, and 7 in the risk-aversion elicitation task. This implies risk-aversion coefficients of 0.15 and 0.97 in terms of a CARA expected utility framework. About 22% of the subjects exhibit inconsistent choices (selecting back and forth between lottery A and lottery B as the probability of the higher payoff increased).

Table A.2: Definitions and Summary Statistics of Explanatory Variables

Definitions for round-level data	
Own success p_i	Player i's success probability (i.e., the probability of generating a new idea conditional on player j having shared an idea in the previous round). Subjects know their own and their rival's success probabilities.
Cross success p_j	Player j's success probability (i.e., the probability of generating a new idea conditional on player i having shared an idea in the previous round). Subjects know their own and their rival's success probabilities.
Expected intentions $\tilde{\sigma}_j$	Player i's expectations that player j will share a newly generated idea in the next round.
Round	Round number of a given match.
Other Terminated	Dummy variable = 1 if player i has previously had a rival who terminated the match by choice either as player A (in odd rounds) or player B (in even rounds). By definition, *Other Terminated* = 0 for the first match.
Own Terminated	Dummy variable = 1 if player i has previously terminated a match by choice either as a player A (in odd rounds) or player B (in even rounds). By definition, *Own Terminated* = 0 for the first match.
Definitions for subject-level data	
Acquaintances	Number of people each participant recognized in the experimental session (Survey question: "How many people in this session do you recognize?").
Friends	Number of a participant's friends that are participating in the same session (Survey question: "How many would you consider friends?").
Fairness	Participant's perception of other people's fairness with higher values indicating more fairness (Survey question: "Do you think that most people would try to take advantage of you if they got a chance, or would they try to be fair?" This question is adapted from the World Values Survey. The questionnaire can be found at `http://www.worldvaluessurvey.org/WVSDocumentationWV6.jsp`.).
Trustworthiness	Participant's perception of other people's trustworthiness with higher values indicating higher levels of trust (Survey question: "Generally speaking, would you say that most people can be trusted, or that you need to be very careful in dealing with people?" This question is adapted from the World Values Survey. The questionnaire can be found at `http://www.worldvaluessurvey.org/WVSDocumentationWV6.jsp`.).
Risk Aversion	Risk aversion category by the Holt and Laury (2002) risk preference task, ranging from 1 to 10 with higher numbers reflecting higher degrees of risk aversion. Risk-aversion results are consistent with the results from Holt and Laury (2002) in that most subjects are risk averse and choose between 5 (21.3%), 6 (14.9%), and 7 (30.3%) in the risk-aversion elicitation task. This implies risk-aversion coefficients of 0.15 and 0.97 in terms of a CARA expected utility framework. Subjects that exhibit inconsistent behavior (selecting back and forth between lottery A and lottery B as the probability of the higher payoff increased), are dropped from the sample when Holt-Laury is used as independent variable.

Summary Statistics

	N	Mean	Std.Dev.	Min	Max
Own success p_i (for Round 1)	578	68.27	19.94	50	90
Cross success p_j (for Round 1)	578	68.27	19.94	50	90
Expected intentions $\tilde{\sigma}_j$ (for Round 1)	578	60.70	30.10	0	100
Round	1574	3.67	3.77	1	22
Other Terminated	1574	0.72	0.45	0	1
Own Terminated	1574	0.53	0.50	0	1
Acquaintances	100	2.92	2.36	0	12
Friends	100	1.81	2.64	0	12
Fairness	100	4.85	2.36	1	10
Trustworthiness	100	5.45	2.61	1	10
Risk Aversion	82	6.95	1.51	3	10

C Means Tests Results for Hypotheses 2, 3, 5, and 6

Table A.3: Average Treatment Effects (Hypotheses 2, 5, and 6)

In the top portion of the table, we report the average level of sharing in Round 1 by player A for treatments HIGH, LOW, LOW-HIGH, and HIGH-LOW. In the bottom portion of the table, we report the results of one-tailed unpaired two-sample t-tests of the pair-wise difference of the mean of sharing (in Round 1 by player A) for Hypotheses 2, 5, and 6. We provide results for the full sample, as well as by three groups of player A's expectations $\tilde{\sigma}_B$ about B's sharing in Round 2: "Low" for $\tilde{\sigma}_B \in [0\%, 33\%]$, "Medium" for $\tilde{\sigma}_B \in (33\%, 66\%]$, and "High" for $\tilde{\sigma}_B \in (66\%, 100\%]$. The prediction is a positive average treatment effect on sharing (e.g., Sharing (Round 1) in HIGH > Sharing (Round 1) in HIGH-LOW). We report the average treatment effects with standard errors in parentheses.

Treatment		Sharing in Round 1 (Player A) Mean (s.e.)	N
HIGH	($p_A = 90\%$, $p_B = 90\%$)	0.8939 (0.026)	132
LOW	($p_A = 50\%$, $p_B = 50\%$)	0.5934 (0.036)	182
LOW-HIGH	($p_A = 50\%$, $p_B = 90\%$)	0.8106 (0.034)	132
HIGH-LOW	($p_A = 90\%$, $p_B = 50\%$)	0.6287 (0.042)	132

Differences: Unpaired two-sample t-test

Prediction	Average treatment effect on sharing (s.e.)			
		by player A's expectations subgroup		
	Full	Low	Medium	High
Hypothesis 2: Positive effect of cross success probability				
HIGH > HIGH-LOW	0.2651***	-0.0051	0.0631	0.1849***
	(0.050)	(0.145)	(0.087)	(0.051)
LOW-HIGH > LOW	0.2171***	0.3882***	0.1247*	0.1012*
	(0.051)	(0.123)	(0.082)	(0.061)
Hypothesis 5: Positive effect of own-success probability				
HIGH > LOW-HIGH	0.0833*	-0.3882**	0.1209	0.0520*
	(0.043)	(0.190)	(0.094)	(0.036)
HIGH-LOW > LOW	0.0353	0.0051	0.1825**	-0.0316
	(0.055)	(0.088)	(0.079)	(0.082)
Hypothesis 6: Effect of cross-success probability is stronger than of own-success probability				
LOW-HIGH > HIGH-LOW	0.1818***	0.3831***	-0.0578	0.1328**
	(0.054)	(0.128)	(0.077)	(0.070)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.4: Average Treatment Effects (Hypothesis 3)

We report the results of one-tailed unpaired two-sample t-tests of the pair-wise difference of mean sharing (in Round 1 by player A) between different belief groups ("Low", "Medium", and "High") for treatments High, Low, Low-High, and High-Low. The three groups of player A's beliefs $\tilde{\sigma}_B$ about B's sharing in Round 2 are "Low" for $\tilde{\sigma}_B \in [0\%, 33\%]$, "Medium" for $\tilde{\sigma}_B \in (33\%, 66\%]$, and "High" for $\tilde{\sigma}_B \in (66\%, 100\%]$. The prediction is a positive average treatment effect on sharing between belief groups (e.g., mean of sharing in "Medium" > mean of sharing in "Low"). We report the average treatment effects (ATE) with standard errors in parentheses.

Treatment	Belief group	Sharing in Round 1		Prediction	Comparison across expectation groups	
		Mean (s.e.)	N		ATE	(s.e.)
High	Low	0.2000 (0.133)	10			
	Medium	0.8846 (0.063)	26	Medium > Low	0.6846^{***}	(0.131)
	High	0.9687 (0.017)	96	High > Medium	0.0841^{**}	(0.047)
Low	Low	0.2000 (0.060)	45			
	Medium	0.6388 (0.057)	72	Medium > Low	0.4388^{***}	(0.086)
	High	0.8153 (0.048)	65	High > Medium	0.1764^{**}	(0.075)
Low-High	Low	0.5882 (0.123)	17			
	Medium	0.7636 (0.057)	55	Medium > Low	0.1754^{*}	(0.124)
	High	0.9166 (0.035)	60	High > Medium	0.1530^{**}	(0.066)
High-Low	Low	0.2051 (0.065)	39			
	Medium	0.8214 (0.051)	56	Medium > Low	0.6163^{***}	(0.082)
	High	0.7837 (0.068)	37	High > Medium	-0.0376	(0.084)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Download ZEW Discussion Papers from our ftp server:

http://ftp.zew.de/pub/zew-docs/dp/

or see:

https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html
https://ideas.repec.org/s/zbw/zewdip.html

//

IMPRINT

ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim

ZEW – Leibniz Centre for European Economic Research

L 7,1 · 68161 Mannheim · Germany
Phone +49 621 1235-01
info@zew.de · zew.de