Discussion Paper Series – CRC TR 224

Discussion Paper No. 040
Project B 02

Expectations of Reciprocity and Feedback when Competitors Share Information: Experimental Evidence

Bernhard Ganglmair*
Alex Holcomb**
Noah Myung***

September 2018

*University of Mannheim, Department of Economics, L7, 3{5, 68161 Mannheim, Germany; E-mail: ganglmair@uni-mannheim.de
**The University of Texas at El Paso, College of Business Administration, 500 West University Avenue, El Paso, Texas 79968, USA; E-mail: ajholcomb@utep.edu
***University of Virginia, Frank Batten School of Leadership and Public Policy, Charlottesville, VA 22904; E-mail: noah.myung@virginia.edu

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Expectations of Reciprocity and Feedback when Competitors Share Information: Experimental Evidence*

Bernhard Ganglmair† Alex Holcomb‡ Noah Myung§

August 22, 2018

Abstract

Informal know-how trading and exchange of information among competitors has been well-documented for a variety of industries, including in science and R&D, and an individual's expectations of reciprocity is understood to be a key determinant of such flow of information. We establish a *feedback loop* (as a representation of information trading) in the laboratory and show that an individual's expectations of the recipient's intentions to reciprocate matter more than a recipient's ability to do so. This implies that reducing strategic uncertainty about competitors' behavior has a bigger effect on the flow of information than reducing environmental uncertainty (about their ability to generate new information). We also show that the formation of beliefs about a recipient's intentions to reciprocate are heavily influenced by past experience, where prior experience lingers and can have negative effects on the sustainability of productive and fruitful information exchange.

Keywords: knowledge diffusion; information sharing; reciprocity; collective innovation; R&D; 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, Joel van der Weele, and conference and seminar participants at the City University London, George Mason University, University of Maryland, University of Texas at Dallas, University of Texas at El Paso, Society of Experimental Finance (2014, Zurich; 2016, Tucson), ESA 2014, the Dublin Colloquium of Behavioral and Experimental Finance 2014, the Virginia Experimental Social Science Conference 2016, and the Research in Behavioral Finance Conference 2016 (Amsterdam) for comments and suggestions. Financial support from the Center and Laboratory for Behavioral Operations and Economics (CLBOE) at the University of Texas at Dallas and the CRC TR 224 - EPoS is gratefully acknowledged. Internal Review Board approval has been obtained at the University of Texas at Dallas (MR 14-127).

†University of Mannheim and ZEW. Address: University of Mannheim, Department of Economics, L7, 3–5, 68161 Mannheim, Germany; E-mail: ganglmair@uni-mannheim.de; Twitter: @ganglmair.

‡The University of Texas at El Paso, College of Business Administration, 500 West University Avenue, El Paso, Texas 79968, USA; E-mail: ajholcomb@utep.edu.

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

1 Introduction

The existence of markets for ideas and the functioning of the resultant information exchange has long been of interest to economists (e.g., Arrow, 1962). Of particular interest are the often puzzling cases where competitors share information that is of strategic value to them. Numerous empirical studies provide evidence of such knowledge exchange in a variety of competitive settings, including the steel minimills industry (von Hippel, 1987; Schrader, 1991), the semi-conductor industry (Appleyard, 1996), academic research (Bouty, 2000; Häussler, 2011; Häussler et al., 2014), French *Haute cuisine* (Fauchart and von Hippel, 2008), or financial investment (Crawford et al., 2017; Botelho, forthcoming). One common theme in these studies is an individual's expectation of reciprocity as 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 wherein two parties are more likely to share information with each other if they expect the other party to reciprocate in kind.[1] In this paper, we investigate how expectations of reciprocity affect the trading of ideas (via a feedback loop) in a controlled laboratory environment. More specifically, we argue that a player must have both the *ability* and the *intention* to reciprocate, and we ask how a player's expectations of ability and intentions differentially affect the incentive to share information.

Our framework (as 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 is via an escalating feedback loop (with potential payoffs increasing in the number of interactions) that breaks if a player conceals new information or has no new information to share. The

[1] This information or feedback 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 use the term *feedback* for the latter kind.

players' payoffs accrue from a *monopolistic segment* of the market (increasing in the absolute amount of information they hold) and a *competitive segment* (increasing in the relative amount of information). Players face a simple trade-off: concealing new information gives them an immediate payoff advantage (through an information advantage that translates into higher payoffs in the competitive segment), but it also breaks the feedback loop, and they forego the chance to generate (and share) more information (increasing payoffs in the monopolistic segment). Stein (2008) shows that if the probability that a player generates new information—and is therefore *able* to share information—is sufficiently high, a *sharing equilibrium*, in which both players always share, exists.

We introduce two simple extensions to the baseline model to investigate the role of expectations of reciprocity in this setting. First, we introduce a player's expectations of a rival player's intentions because, in order to receive feedback, the other player 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 this adaptation, we answer the three 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 decision to initiate and maintain the information exchange? Third, what elements of players' experience 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.[2] We choose this finance-related frame because context is important, and a financial context would be more familiar to our experimental subjects (most of whom are business school students).[3] Because we do not consider trading *per se* but rather focus on the ex-

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

[3]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 funds incur to ensure informational security

change of ideas between competitors, the ways that expectations of reciprocity incentivize information exchange are similar to those highlighted in the existing empirical literature. These effects are the consequence of an analogous trade-off, where a player juxtaposes short-term gains from concealing ideas and long-term gains from sharing (e.g., Mukherjee and Stern, 2009). The game proceeds so that in each round, a fund manager has the chance to generate a new idea to share. The probability of doing so reflects a manager's ability. If a manager generates a new idea, she must decide to share the idea with her rival or conceal it. If she conceals the idea, then the feedback loop breaks and the game ends. If she shares the idea, then her rival is given the chance to generate a new idea according to her ability. If the rival generates a new idea, she can either share this idea (and thus provide "feedback") or conceal it—*et cetera*. If either of the players fails to generate a new idea, then no new information can be shared and the game ends. The game continues an indeterminate number of rounds and ends if a player either decides to conceal an idea (the game is terminated *by choice*) or fails to generate a new idea (the game is terminated *by chance*).

In the theoretical model, if the benefits from sharing information (i.e., the long-term payoffs from an escalating 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 that with higher 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

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; Gerritzen et al., 2016), 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. Our results contribute to this debate.

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 to continue providing feedback and her intention to reciprocate. This means that both players must have valuable information to share, and that they must also intend to share it as reciprocal feedback. We further find that a subject's expectations of the recipient's intentions to share, rather than the 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. We also document how negative past experience (either by a rival or self-inflicted) results in subjects that are less inclined to share information.

Our paper contributes to the literature in a number of ways. First, 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 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. They show that, in equilibrium, the secret holder voluntarily reveals her secret to salvage the cooperative (and profitable) feedback loop. The context of their article is one where product-market competitors cooperate to establish standards in production and design that mutually benefit the entire industry.

Numerous empirical studies, covering a variety of industries, have highlighted the role of expected reciprocity and feedback as a driver of individuals' incentives to share information.[4] von Hippel (1987) and Schrader (1991) report empirical evidence of know-how sharing (or trading) of competing firms in the steel minimill industry. Bouty (2000), Häussler (2011), and Häussler et al. (2014) present results for knowledge sharing in academic

[4] Botelho (forthcoming) provides a detailed review of the literature.

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. Fauchart and von Hippel (2008) report that French chefs who share new recipes have higher expectations of receiving information in the future.

In finance, Crawford et al. (2017) and Botelho (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 (in Stein (2008)). Beyond these two studies, the relevant literature is limited by the data and merely suggestive of managers collaborating in the manner we describe (e.g., Hong et al., 2004; Cohen et al., 2008; Pool et al., 2015).[5]

Last, the structure of our model resembles that of a centipede game (Rosenthal, 1981; Binmore, 1987), with the exception that our game is of indeterminate horizon. In the finite-horizon centipede game, the unique subgame perfect equilibrium is for the first player to defect in the first round and end the 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.

[5]Hong et al. (2004) argue that stock-market participation is strongly influenced by social interaction; Cohen et al. (2008) extend this basic idea and claim that social networks may be important mechanisms through which asset prices incorporate private information; and Pool et al. (2015) show that managers who live in the same neighborhood have significantly higher overlap in their portfolio holdings than managers who live in the same city (but are not neighbors) because they have a greater chance of being socially connected.

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.[6]

The paper is structured as follows. In Section 2, we introduce our adaptation of the work by Stein (2008) and list the theoretical predictions of this model. In Section 3, we discuss the experimental design and develop our empirical hypotheses. In Section 4, we present our main results. We conclude in Section 5 with a discussion of implications for the design of formal and informal platforms of knowledge within and across organizations.

2 Theoretical Framework

We consider an asymmetric version of the model of word-of-mouth communication in Stein (2008). In this model, players exchange private information that yields a temporal payoff advantage if concealed. If players share the information, both players' payoffs may increase.[7] 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.

2.1 A Model of Word-of-Mouth Communication (Stein, 2008)

The story of the model as presented in Stein (2008) is as follows: Two players take turns in generating and sharing new ideas (e.g., for investment opportunities). 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

[6] 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.

[7] 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.

Figure 1: Timeline of Word-of-Mouth Communication

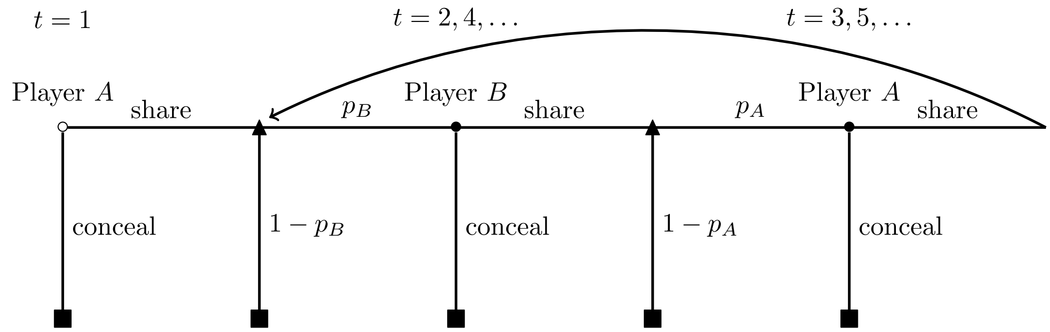

to conceal the idea, the game ends. If she shares, the game continues. Upon observing A's idea, in $t = 2$, player B generates a new idea with probability p. If she is successful, she must decide whether to share this idea with A or keep it to herself. If she conceals the idea, the game ends; if she shares, the game proceeds to $t = 3$. The game proceeds in this fashion in all t. See the graphical depiction of this game in Figure 1.

More ideas increase a player's potential payoffs. The number of ideas a player i can access, say n_i, is the number of ideas she herself has generated plus the number of ideas the other player has shared with her. Stein uses a simple model in which ideas reduce a player's production costs β^{n_i} with $\beta < 1$. See the graphical illustration of the payoff structure in Figure 2. 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 $\theta\mu\,(\beta^{n_j} - \beta^{n_i})$. If player i has fewer ideas than player j, then player i's payoffs in the competitive segment are equal to zero.[8] The realized

[8] 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

Figure 2: Competition and Payoffs

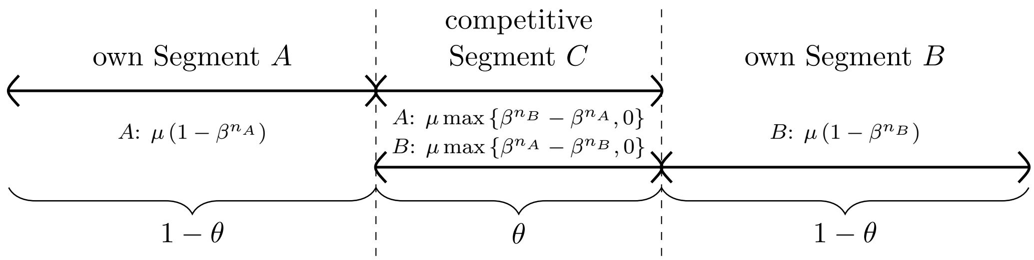

payoffs for player i, given n_i and n_j, are

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

The payoffs in the model give rise to a simple tradeoff for a player who was successful in generating a new idea in t. Concealing this 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 (and no further cost reductions realized). We will later 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 this idea, on the other hand, gives the other player j, in $t+1$, a chance to generate a new idea (with probability p) and then share this idea back with player i, which in turn, will give i the chance to continue the game in $t+2$; and so forth.

The tradeoff player i faces in t is one of short-term gains from concealing versus potential future gains from sharing. Suppose both players always share a new idea if they

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).

generate one. The game thus ends only if a player fails to generate a new idea (we will later refer to this type of termination as *termination by chance*). In a given t, player i's expected benefits from sharing can be written as

$$\text{share}_t = \mu\,(1-\theta)\sum_{k=0}^{\infty}(1-p)\,p^k\left(1-\beta^{t+k}\right). \tag{3}$$

Stein argues that, if the long-run benefits from always sharing outweigh the short-run benefits from concealing in all t, then both players will always share. In this case, the expected payoffs share_t are expected payoffs in equilibrium. Stein then shows that if the net benefits from sharing, ϕ_i, are positive,

$$\phi_i \equiv \mu\,[p\beta-\theta] \geq 0, \tag{4}$$

then a sharing equilibrium exists in which both players always share newly generated ideas.

2.2 Asymmetric Abilities and Players' Intentions

It is this last property of the baseline model that allows Stein (2008) to conclude that competitors are more likely to share valuable information when they are more optimistic about receiving feedback. Notice, however, that this version of the model is symmetric in the sense that an increase in p implies an increase in both players' abilities. This means, an increase in p results in an increase in a player's probability of receiving feedback, as well as that player's probability of being able to give feedback in future rounds of the conversation.

In order to draw more nuanced conclusions about the effect of receiving feedback, we introduce two simple extensions to the baseline model. First, we assume asymmetric abilities. Player A's probability of generating a new idea is denoted by p_A, whereas player B's is denoted by p_B. These probabilities are common knowledge. Second, we introduce a player's expectations of the other player's intentions. To receive feedback, it is not enough for the other player to be able to generate a new idea. The other player must also intend to share a new idea. To this end, suppose that a player j shares a newly generated idea with probability σ_j.[9] We denote player i's expectations of player j's intentions by $\tilde{\sigma}_j \equiv E(\sigma_j)$. Combining the two components of feedback, a player i's expectations of reciprocity in $t+1$ upon sharing an idea in t are

$$\pi_j \equiv p_j \tilde{\sigma}_j. \tag{5}$$

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

$$\text{share}_{it} = \mu\,(1-\theta) \sum_{k=0}^{\infty} p_i^k \pi_j^k \left[(1-\pi_j)\left(1-\beta^{t+2k}\right) + \pi_j \left(1-\beta^{t+1+2k}\right)\right]. \tag{6}$$

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

$$\tilde{\phi}_i(\tilde{\sigma}_j) \equiv \mu \left[\frac{1+\beta p_i}{1+\beta \pi_j} \beta \pi_j - \theta\right] \geq 0. \tag{7}$$

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.[10] The net benefits from sharing are increasing in $\pi_j = p_j \tilde{\sigma}_j$ and p_i.

Condition (7) is necessary for a sharing equilibrium, as is condition (4) in the baseline model. To be more specific, 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:[11] First, in a sharing equilibrium, all players always share newly generated ideas (so that in equilibrium $\sigma_j = \tilde{\sigma}_j = 1$ for $j = A, B$). The game

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

[10] Condition (7) reduces to the condition (4) (as in the baseline model in Stein (2008)) when $\tilde{\sigma}_j = 1$ and $p_i = p_j = \pi_j = p$.

[11] 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).

continues until it is terminated by chance. Second, in a non-sharing equilibrium, neither player ever shares a new idea. The game is terminated by choice (by player A) in $t = 1$.

2.3 Predictions from the Model

With condition (7) satisfied, the game has multiple equilibria. For our theoretical predictions, we use payoff dominance as selection criterion (Harsanyi and Selten, 1988; Cooper et al., 1990). The sharing equilibrium payoff-dominates the non-sharing equilibrium, which follows from the sharing condition in equation (7). We therefore 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 $\tilde{\phi}_i(\tilde{\sigma}_j)$ 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.*

For the theoretical predictions that follow, we make use of this payoff-dominance argument. Anything that relaxes the condition in equation (7), by increasing the value of $\tilde{\phi}_i(\tilde{\sigma}_j)$, also increases the net benefits from sharing and will induce players to choose the sharing strategy more often. With our simple extension of the model, we can separate the effects of receiving feedback from a player's own ability to give feedback. Increasing a player i's expectations of reciprocity π_j increases player i's net benefits of sharing. We then expect the player 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 feedback with higher probability.*

Both components of feedback (ability and expected intentions) are expected to have the same effect on player i's behavior. We refer to ability and intentions as perfect substitutes. This prediction also implies that environmental uncertainty (captured by p_j) and strategic uncertainty (captured by $\tilde{\sigma}_j$) exhibit the same effects on player i's behavior.

Prediction 3. *A player j's ability to send feedback, p_j, and her (expected) intentions, $\tilde{\sigma}_j$, are expected to have the same effect on player i's decision to share.*

In the baseline model, a player's and recipient's abilities to send feedback have, by design, 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 give feedback in $t + 2$ has a positive effect on her decision to give feedback in t. This comes through a direct and indirect effect. The direct effect is the positive effect of p_i on $\tilde{\phi}_i(\tilde{\sigma}_j)$, holding everything else—more specifically, the player's expectations $\tilde{\sigma}_j$—constant. A player's own ability increases the value of ongoing information exchange because, with a higher p_i, the conversation is expected to last longer and therefore more ideas are generated and exchanged.

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

This direct effect of 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 to send feedback, p_i, has a weaker effect on her decision to share than player j's cross ability, p_j.*

The indirect effect of own ability on the player's decision channels through her expectations of the other player's intentions. We predict that p_j has a positive effect on player i's decision (via Prediction 2). 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.

Prediction 6. *Player i's own ability to send feedback, p_i, has a stronger effect on her expectations about player j's intentions than player j's cross ability, p_j.*

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.

3 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 (High, Low, High-Low, and Low-High). Each subject participated in only one treatment. Each session lasted anywhere from 80 to 120 minutes, depending on the treatment. Payments ranged from \$10 to \$30, averaging \$19.30.[12] 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.[13] The main experiment was programmed and executed via zTree (Fischbacher, 2007). The outcomes of the lottery in the Holt-Laury risk preference task

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

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

and the respective payoffs were revealed after the computerized experiment at the end of the session. Subjects were provided with detailed printed instructions for both the Holt-Laury task and the computerized experiment, and a short quiz was conducted after the instructions had been read out by the experimenter.

In the computerized experiment, at the beginning of each period, subjects were randomly matched into pairs without replacement. After the matches were determined, the subjects in each match were randomly assigned the roles of player A and player B. 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 ideas for investment opportunities and compete over capital of new investors.[14] 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 j's intentions, $\tilde{\sigma}_j$.[15] 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).[16] 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.[17] If player i has shared an idea in a round t, then player j sees the decision screen in round $t+1$ (provided she has generated a new idea with probability p_j). If instead, player i decides to conceal the idea, then the match is terminated. After all matches had been terminated, the subjects observed their payoffs from the current match and their

[14]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."

[15]For all odd rounds, we obtain player A's expectations $\tilde{\sigma}_B$; for even rounds, we obtain player B's expectations $\tilde{\sigma}_A$.

[16]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?"

[17]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.

accumulated payoffs from all previous matches.[18] This concluded a match. The subjects were then rematched within their respective group, and a new game was played.

We would like to emphasize the following two design considerations. First, this experiment consists of a one-shot game of *indeterminate* horizon (i.e., played multiple times by each subject with a randomly matched partner). The game ends if one of the players either fails to generate a new idea or conceals an idea. We do not force a match to end prematurely.[19] Second, we do not incentivize the reporting of expectations $\tilde{\sigma}_j$, but we do incentivize the formation of these beliefs.[20] Because 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 $\tilde{\sigma}_j$. This means that, for player i, the formation of these expectations is fully incentivized within the game itself. While our approach of surveying subjects' expectations about their match partners' next-round behavior does not provide incentives for truthful reporting of these expectations, we are confident that, on average, expectations are reported truthfully, albeit with more noise.[21]

The incentive effects of expectations of reciprocity on sharing new ideas that we establish in our finance-related frame are the same as those highlighted in empirical literature discussed in the introduction (e.g., von Hippel, 1987; Schrader, 1991; Fauchart and von Hippel, 2008; Häussler et al., 2014), and our results carry over without limitations. In our narrative, fund managers compete across potential investors for funds, and each man-

[18] From the current match's payoffs, player i is able to infer whether the match has been terminated by chance (player j failed to generate a new idea) or by choice (player j decided to conceal a new idea). We make this point explicit in the printed instructions.

[19] We acknowledge that, in practice, reputational concerns vis-à-vis a player's community as a whole may play a role in her decision to share. Sharing information in one instance may give a player a reputation of being cooperative, inducing other players to update beliefs about this player. Because of the random matching design of the experiment, we do not expect such reputational concerns to play a role in our empirical results.

[20] A similar approach is chosen, for instance, in Cohn et al. (2015). We do not provide incentives in eliciting subjective expectations for practical reasons: incentivizing player i through, for instance, a scoring rule with higher payments for lower linear, logarithmic, or squared difference between the stated expectations and player j's actual decision is not practical when the game is terminated by chance before it is player j's turn to share or conceal (so that no decision by player j is observed).

[21] Trautmann and van de Kuilen (2015) find that introspective beliefs are no less accurate than incentive-elicited beliefs. Their study supports the notion that introspection is a valid method to measure subjective beliefs. See Palfrey and Wang (2009) for a broad set of related results.

ager exerts monopolistic control over a fraction of the market[22] and competes with their rivals over the remaining portion.[23] A manager's payoff increases in the absolute number of investment ideas that she possesses at the end of the game. This part of a manager's payoffs accrues from the monopolistic segment of the market. It is akin to the portion of a fund manager's compensation that stems from good performance with the investor capital already under her control. In addition, if a manager holds more ideas than her rival, she captures the competitive segment of the market as well. The fund manager with the best relative investment performance captures all the remaining uncommitted investor capital. This combination of the absolute and relative number of ideas introduces the same trade-off as discussed in Section 2: on the one hand, sharing ideas increases the potential number of ideas (via a sustained feedback loop in a sharing equilibrium) and thus increases payoffs from the monopolistic side of the market; on the other hand, concealing an idea allows a player to capture the (immediate) competitive side of the market.

3.2 Calibration

We implement the game depicted in Figures 1 and 2 with the realized payoffs in (1). We set $\mu = 400$, $\beta = 3/4$, and $\theta = 3/8$ but vary the success probabilities p_A and p_B, assigning values $p_i \in \{50\%, 90\%\}$.[24] We summarize the calibrations for the four treatments of the experiment in Table 1. For treatment High, we assume symmetric success probabilities $p_A = p_B = p = 90\%$; for treatment Low, the symmetric success probabilities are $p_A = p_B = p = 50\%$. For treatments Low-High ($p_A = 50\%$ and $p_B = 90\%$) and High-Low ($p_A = 90\%$ and

[22] Initial investors in a fund often agree to a significant lock-in period in which they cannot withdraw any invested money from the fund (Agarwal and Naik, 2000). Alternatively, funds may form what are called "side pockets" which are frozen by managers so that redemptions do not force the inefficient early liquidation of assets (McCrary, 2002:192). In the context of our model, these initial investors are fully captured by the fund manager and represent the monopolist portion of the capital market.

[23] The majority of hedge funds are not closed to new investors for a significant amount of time after creation. This means that the fund manager will continue to compete to raise additional capital either from the fund's existing investors or from new investors (Goetzmann et al., 2003). Available (or not yet committed) capital represents the competitive portion of the capital market.

[24] 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 1: Calibration and Treatments (with $\mu = 400$, $\beta = 3/4$, and $\theta = 3/8$)

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

$p_B = 50\%$), we assume asymmetric success probabilities. In all four treatments, the sharing condition in equation (7) 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 alongside the nonsharing equilibrium.

3.3 Hypothesis Development

We derive our hypotheses from the predictions in Section 2. For Hypotheses 1 through 6, we focus on the first decision by player A in each match. This way, the player's beliefs are not affected by an earlier decision of the same rival (within the same match).

Prediction 1 states that player A shares more often when her net benefits $\tilde{\phi}_A$ are higher. More specifically, Prediction 2 states that player A is predicted to share more often when her expectations for π_B are higher. Because $\pi_B = p_B\tilde{\sigma}_B$, both p_B and $\tilde{\sigma}_B$ have a positive effect on player A's decision to share, ceteris paribus. We capture the two predictions in the following hypotheses:

Hypothesis 1. *Player A shares more often in $t = 1$ when $\tilde{\phi}_A$ is high; that means, more often in treatment* HIGH *relative to* LOW.

Hypothesis 2. *Player A shares more often in $t = 1$ when p_B is high; that means, more often in treatment* HIGH *relative to* HIGH-LOW *and* LOW-HIGH *relative to* LOW.

Hypothesis 3. *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.*

Prediction 3 states that, ceteris paribus, p_B and $\tilde{\sigma}_B$ have the same effect on player A's decision to share in $t = 1$.

Hypothesis 4. *Player A's decision in $t = 1$ is the same for $(p_B, \tilde{\sigma}_B)$ such that $\pi_B = p_B \tilde{\sigma}_B$ is constant.*

Hypothesis 2 (and Prediction 2) relates to the effect of a rival player's success probability on one's own willingness to share. A similar positive effect on sharing stems from a player's own success probability (Prediction 4). This effect is two-pronged. A player i's own-success probability has a positive direct effect and a positive indirect effect (through expectations $\tilde{\sigma}_j$) on her likelihood to share. The following hypothesis reflects this joint effect.

Hypothesis 5. *Player A shares more often in $t = 1$ when p_A is high; that means, more often in treatment* High *relative to* Low-High *and* High-Low *relative to* Low.

Prediction 5 states that, when holding $\tilde{\sigma}_B$ constant, player A's own probability of success p_A has a weaker effect on her decision to share than her rival's probability p_B. This effect of p_A is the direct effect as discussed above, whereas the indirect effect is zero because $\tilde{\sigma}_B$ is held constant.

Hypothesis 6. *Holding expectations $\tilde{\sigma}_B$ constant, the 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.

Finally, Prediction 6 states that a player's own ability is expected to have a positive effect on her expectations of player j's intentions, and this effect is expected to be stronger than the effect of p_j. This captures the indirect (or second-order) effect of p_i on player i's decision to share.

Hypothesis 7. *The effect of p_i on player i's beliefs in t about player j's willingness to share an idea in $t + 1$ is stronger than the effect of p_j.*

For this last hypothesis, we use data from all rounds of a given match.

4 Experimental Results

4.1 Sharing Equilibrium

In this section, we present data and results 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 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 2, we report the total number of subjects and matches for each treatment, the average duration of each match (i.e., the average number of rounds and decisions by the players), and the average earnings (per match) for each subject. We explain how the matches in each treatment were terminated (by either chance or choice) and how players expected their rivals to behave. We also report the (theoretical) net benefits of a player continuing in any given period t, $\tilde{\sigma}_A$ and $\tilde{\sigma}_B$.

Matches in treatments with a higher average success probability exhibit a longer duration. First-order (direct) and higher-order (indirect) effects are at play here. For the direct or first-order 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 (or higher-order effect), higher success probabilities are likely to increase the values of $\tilde{\sigma}_i$. We see this effect on expectations when we compare expectations $\tilde{\sigma}_B$ by player A in treatment High-Low (50.9%) to expectations in treatment High (82.4%). Likewise, we see the same pattern when comparing player B's expectations $\tilde{\sigma}_A$ in treatment High-Low (59.8%) to the expectations in treatment High (78.7%). As we will formally test below, with more optimistic expectations about a rival-partner's intentions, a player is less likely to terminate a match by *choice*, increasing the duration of a given match.

Longer matches are indicative of the players choosing sharing-equilibrium strategies. However, because of the direct effect of higher success probabilities on duration, a simple

Table 2: 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 1. All treatments were conducted in one session with two groups of equal size s_g. For the calibration of the treatments, see Table 1. 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 percentage of matches terminated by chance (when either player A or player B has failed to generate a new idea), by player A in an odd round, or by player B in an even round. Because a match is terminated by either chance or choice, these percentages sum up to 100% (with rounding errors). We also provide player A's (theoretical) net benefits from sharing, ϕ_A, in Round 1 when $\tilde{\sigma}_B = 1$.

	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
Percentage of matches terminated by chance …				
… b/c player A has failed	26.5%	37.4%	8.3%	47.0%
… b/c player B has failed	27.3%	10.4%	45.5%	4.5%
Percentage of matches terminated by choice …				
… by player A	21.2%	40.6%	25.0%	38.6%
… by player A in Round 1	10.6%	40.6%	18.9%	37.1%
… by player B	25.0%	11.5%	21.2%	9.8%
Expected intentions $\tilde{\sigma}_j$ …				
… 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 ϕ_A in Round 1	120.00	0.00	71.64	32.73
Theoretical expected net benefits ϕ_B in Round 2	120.00	0.00	32.73	71.64

comparison of the duration of matches across treatments is misleading. Instead, we consider the players' termination choices to summarize our first result (Hypothesis 1) as follows:

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

In matches that are terminated by choice, the terminating player chooses a strategy other than a sharing-equilibrium strategy. The data in Table 2 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 matches being terminated by choice in treatment HIGH

(46.2%) than in treatment Low (52.1%).[25] In 10.6% (treatment High) to 40.6% (treatment Low) of all matches, player A terminates in Round 1, and the equilibrium outcome is that of the payoff-dominated non-sharing equilibrium (in which neither player shares).[26]

Despite its preliminary nature, Result 1 represents our baseline result. We have calibrated our model such that, theoretically, both sharing and non-sharing are Nash equilibrium strategies. In fact, 53.8% of all matches in the High treatment and 47.9% in the Low treatment play the sharing equilibrium.[27] We have argued before that the sharing equilibrium is payoff-dominant and that higher (theoretical) net benefits from sharing should increase the effect of payoff dominance for equilibrium selection. In other words, as the net benefits from sharing in any given round increase, a player is more likely to choose a sharing strategy (or is less likely to make a mistake as the payoff consequences are more severe).

4.2 Reciprocity and Feedback as a Combination of Ability and Intentions

In this section, we present results addressing the second question: how do the players' expectations of reciprocity drive their decision to share information? First, we find support for the simple expectations-of-reciprocity argument (Hypotheses 2 and 3):[28]

Result 2. *Player A is more likely to share an idea when she has a higher expectation of receiving feedback.*

[25] This line of reasoning comes with a caveat. Matches that were terminated by chance are not necessarily the outcomes of a sharing equilibrium. A player's (non-sharing) strategy may prescribe terminating the match after n rounds of sharing. If the match is terminated by chance beforehand, we do not observe the realization of this non-sharing strategy. We therefore treat our conclusions in Result 1 as preliminary and provide more formal tests in the following subsections.

[26] Termination by choice in Round 1 is also the outcome of the finite-horizon centipede game. Previous studies of the centipede game have found that only a small number of games end in the first round as predicted by theory (from 0.7% in McKelvey and Palfrey (1992) to 3.9% in Levitt et al. (2011) with expert chess players a subjects). Studies with linear payoffs report even higher levels of cooperation (Bornstein et al., 2004; Gerber and Wichardt, 2010). In our indeterminate-horizon variant of the centipede game, we expect higher levels of cooperation than in the finite-horizon game but find the opposite.

[27] These figures are upper bounds, as discussed in Footnote 25.

[28] We provide results from simple means tests in the Online Appendix.

We get a preview of this result when comparing player i's actions in a treatment in which player j has low success probability to a treatment with high success probability for player j. Consider player A. Per Prediction 2, player A is more likely to share, and less likely to terminate (by choice) in treatments with high p_B. We observe this by comparing the percentage of matches terminated by choice in treatment HIGH with treatment HIGH-LOW (21.2% < 38.6%), as well as in treatment LOW-HIGH with treatment LOW (25.0% < 40.6%).

These aggregate figures comprise player A's behavior in all rounds, and it is reasonable to assume that the player's experience in earlier rounds of a match has an effect on her later actions. For instance, observing player B sharing to continue the match allows player A to update her beliefs about player B's intentions.[29] To obtain a clearer picture of the treatment effect, 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). In panel (a) of Figure 3, we plot the mean of sharing by player A in Round 1 of each match and report simple means tests. The average treatment effect is as predicted in Hypothesis 2 for all but the last treatment.

Next, we have argued that the probability of receiving feedback, π_B, is composed of a rival-partner's *ability* p_B to share information and her (expected) *intention* $\tilde{\sigma}_B$ to share information. Both components increase the probability of feedback and both components are likely interdependent. For instance, moving from treatment HIGH-LOW to treatment HIGH increases player B's ability while keeping player A's ability constant. At the same time, player A's expectations $\tilde{\sigma}_B$ about player B's intentions also increase. This is the case for player A's expectations across all rounds (50.9% and 82.4%, respectively) and her expectations in Round 1 only (49.1% and 77.3%, respectively). The latter is shown in panel (b) of Figure 3 where we report player A's expectations in Round 1 only.

The results in Table 2 and Figure 3 represent a compound effect. To show the individual effects of ability and intention, we present regression results from probit models in Tables 3 and 4. The dependent variable is a dummy variable equal to 1 if player A shares

[29] We return to this issue further below in Results 8 (updating within a match) and 9 (updating across matches).

Figure 3: Sharing and Expectations in the First Round

This figure plots the average level of sharing in Round 1 by player A (panel (a)) and player A's expectations $\tilde{\sigma}_B$ in Round 1 (panel (b)) for all four treatments. In panel (a), the (theoretical) expected net benefits from sharing, $\tilde{\phi}_i$, as defined in equation (7), for $i = A, B$ and $\tilde{\sigma}_j = 1$ are provided. In panel (b), we provide a box plot for FM A's expectations $\tilde{\sigma}_B$ that player B will share in Round 2. The table below reports 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. The prediction is a positive average treatment effect on sharing for treatments with higher $\tilde{\phi}_A(1)$ relative to lower $\tilde{\phi}_A(1)$ (Prediction 1). We rank treatments by their respective value of $\tilde{\phi}_A(1)$ and predict that mean(High) > mean(Low-High), mean(Low-High) > mean(High-Low), and mean(High-Low) > mean(Low). Median values are shown as horizontal lines and mean values as dots and reported in brackets. We report the average treatment effects with standard errors in parentheses.

Prediction			Avg treatment effect on sharing (s.e.)	
Sharing (Rd 1): High [0.8939]	>	Sharing (Rd 1): Low-High [0.8106]	0.0833**	(0.043)
Sharing (Rd 1): Low-High [0.8106]	>	Sharing (Rd 1): High-Low [0.6287]	0.1818***	(0.054)
Sharing (Rd 1): High-Low [0.6287]	>	Sharing (Rd 1): Low [0.5934]	0.0353	(0.055)

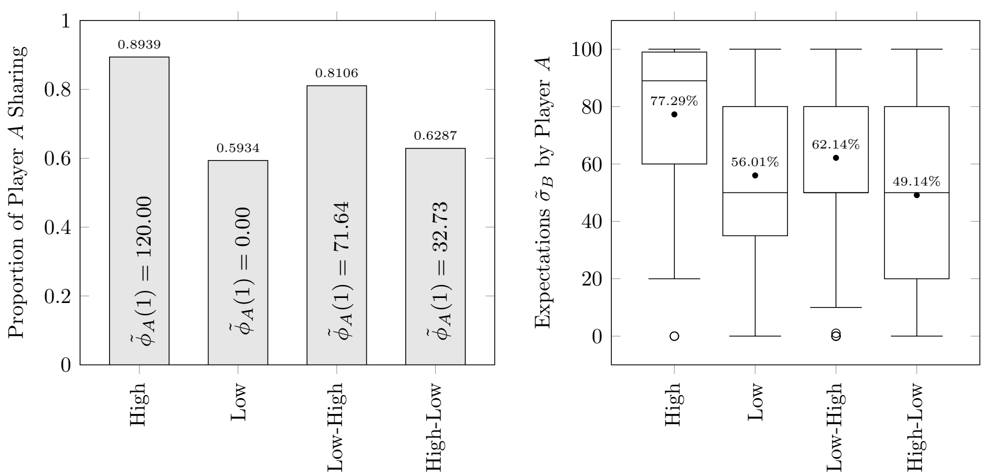

(a) Sharing in Round 1 **(b)** Expectations in Round 1 in Round 1, and equal to 0 if player A does not share the idea in Round 1. As in Figure 3 panel (b), we restrict our data to player A's behavior in Round 1 of each match. We find positive marginal effects of both the cross success probability p_B and the expected intentions $\tilde{\sigma}_B$. For these results, we use minimal specifications. We provide robustness results with extended model specifications in the Appendix.[30]

The marginal effects of the cross-success probability p_B and of player A's expectations $\tilde{\sigma}_B$ (about player B's future actions) are positive and significant ($p < 0.01$), supporting our Hypotheses 2 and 3 and collectively supporting Hypothesis 1. The marginal effects reported in Table 3 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-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. Put differently, a one-standard deviation increase in player A's expectations $\tilde{\sigma}_B$ (that is, an increase of 30.1%) increases the player's probability of sharing an idea in Round 1 by 16.8% to 19.0%.

Result 2 establishes a positive effect of receiving feedback on player A's sharing behavior. This result comports with the collaboration argument in Stein (2008) and Crawford et al. (2017) (using Stein's framework) and confirms the empirical evidence cited earlier: higher expectations of reciprocity increase players' incentives to share.

Feedback depends on the other player's *ability* to give feedback (captured by success probability p_j) and that player's intentions (captured by a player's expectations $\tilde{\sigma}_j$). The following Result 3 summarizes the relative effects of the components of feedback (Hypothesis 4).

Result 3. *The source of feedback is not irrelevant.*

[30]Our results concerning how the probability of feedback (measured by ability and intentions) affects sharing 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 B.1 and provide detailed descriptions and summary statistics for these control variables in Table C.1.

Table 3: Probit Regression 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 feedback 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

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

In contrast to the prediction in Hypothesis 4, we do not find conclusive evidence that the source of feedback is irrelevant. Theoretically, any combination of p_j and $\tilde{\sigma}_j$ induces the same behavior as long as $\pi_j \equiv p_j \tilde{\sigma}_j$ remains constant. This implies that the predicted probabilities of sharing by player A ought to be constant for different values of p_B and $\tilde{\sigma}_B$ such that π_B is constant. The results in Tables 3 and 4 suggest that this substitutability does not prevail in the data. In Table 4, we further show the marginal effects of ability and intentions when evaluated at different combinations of p_B and $\tilde{\sigma}_B$ such that $\pi_B = \bar{\pi}_B$. For all these combinations, we reject the null that the marginal effects of ability and intentions are the same.[31]

It follows from Result 3 (and the negation of Hypothesis 4) that if ability and intention are not perfect substitutes empirically, then subjects must find one more salient than the other, as is evidenced by the marginal effects in Tables 3 and 4 and summarized in Result 4:

[31] A Wald test of simultaneous equality of coefficients for p_j and $\tilde{\sigma}_j$ in the models in Tables 3 and 4 rejects the null of equality with at least $p < 0.10$ in all cases.

Table 4: 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 feedback 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 3 are evaluated at values of p_B and $\tilde{\sigma}_B$, keeping $\pi_B = p_B\tilde{\sigma}_B$ constant at $\bar{\pi}_B = \text{mean}(p_B) \times \text{mean}(\tilde{\sigma}_B)$ (i.e., the sample mean probability of feedback); p_A is at the sample mean. The number of observations is 578; the pseudo R^2 is 0.2299. We report standard errors 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 $p_B = 75\%$, $\tilde{\sigma}_B = \bar{\pi}_B/p_B$	ME evaluated at $p_B = 100\%$, $\tilde{\sigma}_B = \bar{\pi}_B/p_B$
Cross success: p_B	0.0033***	0.0042***	0.0041***
	(0.0010)	(0.0009)	(0.0007)
Expected intentions: $\tilde{\sigma}_B$	0.0054***	0.0068***	0.0066***
	(0.0004)	(0.0008)	(0.0011)
Own success: p_A	0.0014*	0.0017*	0.0017*
	(0.0008)	(0.0010)	(0.0010)
	Test of the difference in coefficients : χ^2		
$p_B - \tilde{\sigma}_B = 0$	5.62**	3.74*	2.92*
	(0.0178)	(0.0533)	(0.0874)

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

Result 4. *Player A's expectations of player B's intentions (captured by $\tilde{\sigma}_B$) have a greater impact on player A's decision to share information than player B's ability (captured by p_B).*

A test of simultaneous equality of the respective coefficients in Tables 3 and 4 shows that intention appears to be more predictive of sharing behavior than ability (both own or rival ability). It implies that strategic uncertainty (i.e., uncertainty about the strategy of the rival partner) is more important for a player's decision to share new information than environmental uncertainty (i.e., uncertainty about the ability of the rival-partner with which a player is matched).

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 player's own ability—summarized in Result 5—and for a rival player's ability—summarized in Result 6.

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

We conclude from models (IV) and (V) in Table 3 that player A is more likely to share an idea in Round 1 when she expects to be successful in generating yet another idea in Round 3.[32] The marginal effect of own-success probability p_A is positive and significant ($p < 0.10$), a result which supports Hypothesis 5. The marginal effects imply that player A is 1.4% to 1.5% more likely to share an idea in Round 1 in response to a 10 percentage point increase in her own-success probability.[33] The marginal effect of p_A in model (IV) is the overall effect, whereas the effect of p_A in model (V), which controls for expected intentions $\tilde{\sigma}_B$, is the direct effect only. A comparison of these two suggests that, if there is an indirect effect of p_A on player A's sharing, which operates through player A's expectations of player B's intentions, then this effect is, at best, small. We do not find strong (preliminary) evidence of higher order beliefs playing a part 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 $\tilde{\sigma}_B$ 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 Table 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 $p_B > p_A$). Our regression results in Table 3 paint an analogous picture. We find that the marginal effects of p_B are greater than those of p_A in

[32] Per the aggregate numbers in Table 2, we see this by comparing treatment High-Low with treatment Low (38% < 40%), as well as treatment High with treatment Low-High (21% < 25%). We provide results from simple means tests in the Online Appendix.

[33] The standard deviation of p_A is 19.9, so that a one-standard deviation increase in her own-success probability increases player A's probability of sharing in Round 1 by roughly 3%.

all specifications. To confirm, we perform a Wald test, which rejects the null that the two effects are the same.[34]

4.3 Players' Beliefs About Competitors' Intentions

In this section, we address the third question: what elements of players' experience affect the formation of their expectations about a rival's intended behavior? In Table 5, we present results detailing the determinants of player i's expectations $\tilde{\sigma}_j$ in a Round t, concerning player j's intentions to share an idea in $t+1$. Unlike in our previous analyses, 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, a left-censoring limit of 0, and a right-censoring limit of 100.

We have shown, in Table 3, a positive effect of own ability (p_A) on player A's willingness to share (Result 5). 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 if 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.*

In Table 2, rational belief formation implies higher expectations $\tilde{\sigma}_B$ in treatment HIGH-LOW than in treatment LOW-HIGH; and higher expectations $\tilde{\sigma}_A$ in treatment LOW-HIGH than in treatment HIGH-LOW.[35] We see the reverse pattern in contradiction to rational belief formation. A similar picture emerges in Table 5. In the specifications without subject

[34]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.

[35]Consider player A's expectations $\tilde{\sigma}_B$. The effect of p_A is $\tilde{\sigma}_B(\text{HIGH}) - \tilde{\sigma}_B(\text{LOW-HIGH})$. The effect of p_B is $\tilde{\sigma}_B(\text{HIGH}) - \tilde{\sigma}_B(\text{HIGH-LOW})$. The effect of p_A is stronger if $\tilde{\sigma}_B(\text{HIGH}) - \tilde{\sigma}_B(\text{LOW-HIGH}) > \tilde{\sigma}_B(\text{HIGH}) - \tilde{\sigma}_B(\text{HIGH-LOW})$ or $\tilde{\sigma}_B(\text{HIGH-LOW}) > \tilde{\sigma}_B(\text{LOW-HIGH})$.

Table 5: 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, by definition, = 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

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

dummies, the positive effect of a rival player'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 ($p < 0.01$).

Result 7 suggests that players' expectations about their rival partner'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 *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 *across* matches. We focus on a player's experience in earlier matches (with other rival partners) 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 player'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) within a given match. In Table 5, we show the effect of the number of rounds played on player i's expectations $\tilde{\sigma}_j$ in any t of that match. In other words, given that player j has shared in Round $t-1$, player i updates her (posterior) beliefs 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.[36]

Result 8 confirms the conjecture we made in the previous section when we restricted our analysis in Tables 3 and 4 to player A's decision and expectations in Round 1. We showed that expecations $\tilde{\sigma}_j$ drive a player i's decision to share. The implication of Result 8 is that

[36]An alternative explanation stems from the payoff structure of the experiment. The further the game of word-of-mouth communication progresses, the smaller in size are the absolute and relative costs and benefits of sharing. Despite this practical convergence in payoffs, the theoretical net benefits of sharing are time-invariant in every treatment. However, the convergence in payoffs, especially in the later rounds of a given match, may cause subjects to feel as if the cost/benefit trade-off contained in the decision to share is less salient than it was in a match's earlier rounds.

the history of a match has an effect on a player's decision to share through her expectations of her rival partner's intentions.

In Result 8, we consider a player's updating of posterior beliefs within a match. For the last result, we turn to a player's updating of prior beliefs 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-partner's intentions.

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

In Table 6, 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 3 and capture past experience using two dummy variables. *Other Terminated* is equal to 1 if player A in a previous match had a rival-partner (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), we use the subsample of players A who vary their decisions across matches. In other words, we restrict our sample to players A who do not exhibit match-invariant decisions. We find that the effects of our feedback 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 indeed has an effect on player A's decision to share. For example, in model (I), if player A in an earlier match faced a rival partner who terminated the match by concealing an idea (*Other Terminated* = 1), then player A is 8.4 percentage points less likely to share in Round 1 of the given match,than otherwise.[37]

[37] 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.

Table 6: 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 3. The dependent variable is a dummy variable = 1 if player A shares in Round 1, and = 0 otherwise. Player A's expectations of receiving feedback 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, by definition, = 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

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

The results for *Other Terminated* are consistent with notions of awareness and revision of prior beliefs about a match partner's "type." At the beginning of each match, two players are randomly re-matched (from within each group) without replacement, and at the beginning of each match, a player forms her prior beliefs about the other player's intentions. Because in Table 6 the dependent variable reflects the first decision of a new match, there is no scope for updating these prior beliefs about the other player's type *within* a given match. Player A being more aware or cautious (resulting in lower expectations) explains the negative effect of *Other Terminated*.

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 a 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 share an idea. It is as if a player reveals her own type to herself as soon as she conceals 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.[38] If *Own Terminated* were to capture only a subject-fixed effect, 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 match partner's "type" through an effect analogous to "self-projection" in which a subject "project[s] her known behavior to guess others' behavior" (Lévy-Garboua 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 5, we indeed see a negative effect of *Other Terminated* on

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

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 partner) 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 player terminate a match. This result is in line with our earlier discussion of the salience of one's incentives in response to one's own termination, and the effect arises only if a player has not experienced termination before. Once a player has faced a rival player that terminated the match by choice, a player's own subsequent termination has little to no effect on her belief formation.

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 other industry sectors. 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 (the "potential future reciprocity") (Häussler et al., 2014).

We experimentally test whether a feedback loop that captures this very 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 establish a feedback loop in the laboratory and find 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; Schilling, 2016) and come with implications for the design of formal and informal platforms of knowledge exchange. In a network, 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). The results from our experiment highlight the importance of the strength of links in a social network relative to the number of links for the diffusion of knowledge within the network. To this end, reducing the strategic uncertainty about the behavior of linked members in one's network is more important than reducing the environmental uncertainty about other members' productive potential. In other words, the presence of reliable members is more conducive to a sustainable information exchange than the presence of members with high productive potential. In addition, negative prior experience can linger and have a lasting negative impact on the sustainability of knowledge transfer across or within organizations, highlighting the need for policies to address and neutralize this effect.

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. (forthcoming): "Here's an Opportunity: Knowledge Sharing among Competitors as a Response to Buy-in Uncertainty," *Organization Science*.

Bouty, I. (2000): "Journal Article 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.

Cohn, A., J. Engelmann, E. Fehr, and M. A. Maréchal (2015): "Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals," *American Economic Review*, 105, 860–885.

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.

Fauchart, E. and E. von Hippel (2008): "Norms-Based Intellectual Property Systems: The Case of French Chefs," *Organization Science*, 19, 187–201.

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.

Gerritzen, M., J. C. Jackwerth, and A. Plazzi (2016): "Birds of a Feather – Do Hedge Fund Managers Flock Together?" Research Paper 16-10, Swiss Finance Institute.

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.

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.

Mukherjee, A. and S. Stern (2009): "Disclosure or Secrecy? The Dynamics of Open Science," *International Journal of Industrial Organization*, 27, 449–462.

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.

Palfrey, T. R. and S. W. Wang (2009): "On Eliciting Beliefs in Strategic Games," *Journal of Economic Behavior & Organization*, 71, 98–109.

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.

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. (2016): "Innovation Networks," in *The Palgrave Encyclopedia of Strategic Management*, ed. by M. Augier and D. J. Teece, Palgrave Macmillan, 1–5.

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.

Appendix

A Model: Extended Presentation with Formal Proofs

We summarize the results for our extensions of the model by Stein (2008) in Propositions A.1 through A.4.

A.1 Decisions and Timing

Two players take turns in generating and sharing new ideas. Player A moves in odd rounds, and player B moves in even rounds. Player A begins in round $t = 1$ with an existing idea and must decide whether to share it with player B. In all future rounds $t \geq 2$, player $i = A, B$ then generates one new idea with success probability $p_{i,t}$ and must decide whether to share this new idea with player $j \neq i$ or conceal it. This probability $p_{i,t}$ depends on the previous round's action but is otherwise time-invariant:

$$p_{i,t} = \begin{cases} p_i & \text{if player } j \text{ shared an idea in } t-1 \\ 0 & \text{if no idea was shared in } t-1. \end{cases} \tag{A.1}$$

An assumption of strict complementarity in the generation of ideas (Stein, 2008; Hellmann and Perotti, 2011; Ganglmair and Tarantino, 2014) implies that communication continues until one player fails to generate a new idea (i.e., termination by chance) or decides to conceal a newly generated idea (i.e., termination by choice). Conditional on sharing, the idea-generation process is exogenous.

The timeline of the game and structure of the decision-making is depicted in Figure A.1. The hollow circle indicates the first round and the beginning of the game in which player A decides whether to share or conceal her initial idea. The triangle indicates a move by chance: once player A has shared, player B successfully generates a new idea with probability p_B but fails with probability $1 - p_B$. If player B succeeds, she decides whether to share or conceal the new idea. This decision is indicated by a solid circle. If a failure occurs, the game ends (indicated by a square). The communication continues for an indeterminate number of rounds but has a finite expected duration.

A.2 Payoffs

A player's payoffs are a function of her own stock of ideas n_i and her competitor's stock of ideas n_j. In the main text of the paper, we followed Stein and motivated the payoffs using

Figure A.1: Timeline of Word-of-Mouth Communication

The figure depicts the timeline and structure of the game of word-of-mouth communication. Player A initially hold one idea, and player B holds 0 ideas. Player A in $t = 1$ decides to share or conceal her initial idea. In all $t \geq 2$, player i generates a new idea with probability $p_{i,t}$ and decides to share this idea with player j or conceal the idea. The game continues until one of the players fails to generate a new idea or decides to conceal.

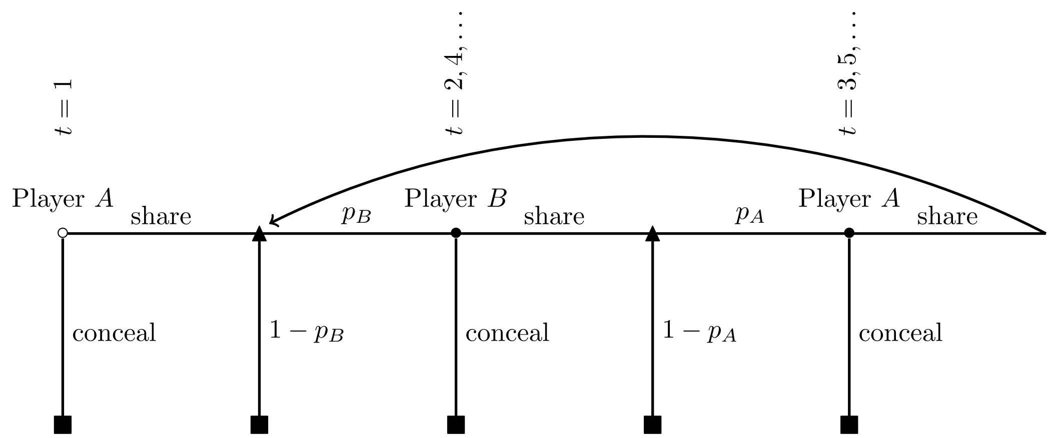

a model of product-market competition in which each player faces a market of size μ with unit-demand consumers (with valuation of 1). In fraction $1 - \theta$ of their respective markets, players priced their goods as monopolists; in fraction θ, the two players are competitors, engaging in price competition à la Bertrand. The parameter θ thus captures the degree of competition.

In this appendix, we motivate the same payoff functions used in the main text via the narrative given to the subjects in our computerized experiments. We depict the market structure in Figure A.2. We assume that ideas are for investment opportunities and that players are fund managers that compete for investors in the following way: Player i has captured all the investors in her own Segment i (with a market share $1-\theta$), and the payoffs from this side of the market depend on the stock of ideas n_i. These payoffs are represented by a function $v(n_i) = 1-\beta^{n_i}$ that increases in n_i at a decreasing rate with $v(0) = 1-\beta^0 = 0$. A fraction θ of a player's payoffs are generated in the *competitive segment* of the market. This Segment C contains new investors that both players compete to attract, and the player who finishes with the greater stock of ideas attracts all new investors. The payoffs from these new investors are greater when the difference between the two players' respective stocks of ideas is greater. In this model, the payoffs for player i from the competitive Segment C are $v(n_i) - v(n_j) = \beta^{n_j} - \beta^{n_i}$ if $n_i > n_j$, and zero otherwise.

The payoffs are realized after the game has ended. At this point, player i's total realized payoffs from her own Segment i and the competitive Segment C are

$$\mu \left[(1-\theta)\left(1-\beta^{n_i}\right) + \theta \max\left\{\beta^{n_j} - \beta^{n_i}, 0\right\} \right]. \tag{A.2}$$

Figure A.2: Competition and Payoffs

This figure depicts the two markets (of size μ) in which the players compete. Player i generates payoffs $1-\beta^{n_i}$ from her own Segment i, where n_i denotes player i's number of ideas (for investment opportunities). The player with more ideas also generates profits from the competitive Segment C. These payoffs for player i are positive if $n_i > n_j$ and increase in $n_i - n_j$.

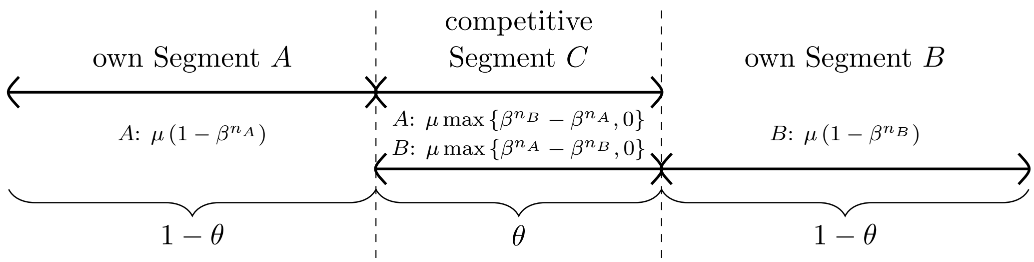

A.3 Incentive Compatible Communication

Formally, provided that stage t is reached (when both players have shared ideas in all $t-1$), if player i conceals the idea, then she holds $n_i = t$ ideas whereas j holds $n_j = t-1$ ideas. Player i's payoffs in any $t \geq 1$ are then:

$$\text{conceal}_t \;=\; \mu\left[(1-\theta)\left(1-\beta^t\right)+\theta\left[\beta^{t-1}-\beta^t\right]\right]=\mu\left[1-\beta^t-\theta\left(1-\beta^{t-1}\right)\right]. \quad \text{(A.3)}$$

If, instead, player i decides to share the new idea in $t \geq 1$, then her expected payoffs depend on the expected future stream of ideas, and therefore on how she expects player j to decide in future rounds. To keep things simple, we assume that both players play time-invariant strategies $\sigma_i \in [0,1]$ where $\sigma_i = \Pr(\text{player } i \text{ shares in } t)$ and let

$$\tilde{\sigma}_j \equiv E(\sigma_j) \quad \text{(A.4)}$$

denote player i's beliefs or expectations of player j's mixed strategy (i.e., player j's probability of sharing).

A player i who shares a newly generated idea in t enables the other player to generate (with probability p) and share (with expected probability $\tilde{\sigma}_j$) a new idea in $t+1$. With probability $\pi_j \equiv p_j \tilde{\sigma}_j$, player i will be able to generate and share yet another idea in $t+2$. Suppose that if player i shares in t with certainty, then she shares in all $t' > t$. In this case, her time-invariant strategy is $\sigma_i = 1$. Her expected payoffs from sharing in t (playing $\sigma_i = 1$) are equal to

$$\text{share}_{it} \;=\; \mu\left(1-\theta\right)\sum_{q=k}^{\infty} p_i^k \pi_j^k \left[\left(1-\pi_j\right)\left(1-\beta^{t+2k}\right)+\pi_j\left(1-p_i\right)\left(1-\beta^{t+1+2k}\right)\right]. \quad \text{(A.5)}$$

For the construction of these payoffs, first note that in period t, player i holds $n_i = t$ ideas. We construct the payoffs by determining the probabilities that player i has exactly $n_i = t + q$ ideas for $q = 0, \ldots, \infty$. With $t + q$ ideas, player i's payoffs are $v(t+q) = \mu\left(1 + \beta^{t+q}\right)$ in its own Segment i and $\mu \cdot \max\left\{\beta^{t+q-1} - \beta^{t+q}, 0\right\}$ in the competitive Segment C. We assume that once player i chooses to share in t, she shares in all future $t' > t$. Hence, $\sigma_i = 1$.

t: When player i shares an idea in t, both players have t ideas and player i's payoffs are $\mu\left(1-\theta\right)\left(1-\beta^{t}\right)$ with probability $1 - \pi_j = 1 - p_j\tilde{\sigma}_j$, that is, the probability that (i) player j fails to generate a new idea in $t+1$ (probability $1-p_j$); or (ii) player j generates a new idea but conceals it in $t+1$ (probability $p_j\left(1-\tilde{\sigma}_j\right)$).

$t+1$: Both players have $t+1$ ideas and player i's payoffs are $\mu\left(1-\theta\right)\left(1-\beta^{t+1}\right)$ with probability $\pi_j\left(1-p_i\right)$, that is, the probability that player j generates and shares a new idea in $t+1$ (probability π_j) but player i fails to generate a new idea in $t+2$ (probability $1-p_i$).

$t+2$: Player i has $t+2$ ideas, player j has at least $t+2$ ideas, and player i's payoffs are $\mu\left(1-\theta\right)\left(1-\beta^{t+2}\right)$ with probability $\pi_j p_i\left(1-\pi_j\right)$, that is, the probability that (i) player j generates and shares a new idea in $t+1$ (probability π_j), player i generates and shares a new idea in $t+2$ (probability p_i), but player j fails to generate a new idea in $t+3$ (probability $1-p_j$); or (ii) player j generates and shares a new idea in $t+1$ (probability π_j), player i generates and shares a new idea in $t+2$ (probability p_i), and player j generates a new idea but conceals it in $t+3$ (probability $p_j\left(1-\tilde{\sigma}_j\right)$).

$t+3$: Both players have $t+3$ ideas and player i's payoffs are $\mu\left(1-\theta\right)\left(1-\beta^{t+3}\right)$ with probability $\left(\pi_j\right)^2 p_i\left(1-p_i\right)$, that is, the probability that (i) player j generates and shares a new idea in $t+1$ (probability π_j), player i generates and shares a new idea in $t+2$ (probability p_i), player j generates and shares a new idea in $t+3$ (probability π_j), but player i fails to generate a new idea in $t+4$ (probability $1-p_i$).

$t+4$: Player i has $t+4$ ideas, player j has at least $t+4$ ideas, and player i's payoffs are $\mu\left(1-\theta\right)\left(1-\beta^{t+4}\right)$ with probability $\left(\pi_j\right)^2\left(p_i\right)^2\left(1-\pi_j\right)$, that is, the probability that (i) player j generates and shares a new idea in $t+1$ (probability π_j), player i generates and shares a new idea in $t+2$ (probability p_i), player j generates and shares a new idea in $t+3$ (probability π_j), player i generates and shares a new idea in $t+4$ (probability p_i), but player j fails to generate a new idea in $t+5$ (probability $1-p_j$); or (ii) player j generates and shares a new idea in $t+1$ (probability π_j), player i generates and shares a new idea in $t+2$ (probability p_i), player j generates and shares a new idea in $t+3$ (probability π_j), player i generates and shares a new idea in $t+4$ (probability p_i), player j generates a new idea but conceals it in $t+5$ (probability $p_j(1-\tilde{\sigma}_j)$).

$t=5$: Player i's payoffs are $\mu(1-\theta)(1-\beta^{t+5})$ with probability $(\pi_j)^3(p_i)^2(1-p_i)$.

$t=6$: Player i's payoffs are $\mu(1-\theta)(1-\beta^{t+6})$ with probability $(\pi_j)^3(p_i)^3(1-\pi_j)$.

$t=\ldots$: etc.

Continuing in this fashion and summing up player i's payoffs for each $q=0,\ldots,\infty$ weighted by the respective probability yields the expression for player i's expected payoffs from sharing.

At any given t, player i shares the idea if

$$\text{share}_t \geq \text{conceal}_{it}, \tag{A.6}$$

given her expectations $\tilde{\sigma}_j$ of player j's future actions. This condition depends on both players' probabilities of success (i.e., *abilities*), as well as player j's expected behavior (i.e., *intentions*). After some algebra, condition (A.6) can be simplified and rewritten to yield condition (7) in the main text:

$$\tilde{\phi}_i(\tilde{\sigma}_j) \equiv \mu\left[\frac{1+\beta p_i}{1+\beta\pi_j}\beta\pi_j - \theta\right] \geq 0. \tag{7}$$

A.4 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-partner 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 in expression (7) (for a given $\tilde{\sigma}_j$) as follows:

$$\tilde{\sigma}_j \geq \frac{\theta}{(1-\theta+\beta p_i)\,\beta p_j}. \tag{A.7}$$

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:

$$s_i(\tilde{\sigma}_j) = \begin{cases} \text{share} & \text{if } \tilde{\sigma}_j \geq \dfrac{\theta}{(1-\theta+\beta p_i)\,\beta p_j} \\ \text{conceal} & \text{if } \tilde{\sigma}_j < \dfrac{\theta}{(1-\theta+\beta p_i)\,\beta p_j}. \end{cases} \tag{A.8}$$

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 in expression (7) 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 $\sigma_i = 1$ 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 $\text{conceal}_{t-1} = \mu\,[(1-\beta^{t-1}) - \theta\,(1-\beta^{t-2})]$ when it 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*

$$\sigma_i^* = \frac{\theta}{(1-\theta+\beta p_j)\,\beta p_i}. \tag{A.9}$$

Proof. In equilibrium, a player's expectations about the rival's strategy are correct, that means, $\tilde{\sigma}_j = \sigma_j$. 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.8), 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.

Assuming time-invariant strategies σ_i for $i = A, B$, we can calculate the expected duration of our game of word-of-mouth communication.

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 (as depicted in Figure A.1).

- 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

$$\begin{aligned}D &= \sum_{q=0}^{\infty} \delta_{q+1} (q+1) \\ &= \sum_{q=0}^{\infty} (\sigma_A p_B)^q (\sigma_B p_A)^q [(1 - \sigma_A p_B)(1+q) + \sigma_A p_B (1 - \sigma_B p_A)(2+q)] \\ &= 1 + \frac{\sigma_A p_B}{1 - \sigma_A \sigma_B p_A p_B}. \qquad \text{(A.10)}\end{aligned}$$

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,

$$\frac{\partial D}{\partial p_B} > \frac{\partial D}{\partial p_A} \iff \sigma_A \sigma_B p_B^2 < 1, \qquad \text{(A.11)}$$

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

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

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

$$\frac{\partial \tilde{\phi}_i}{\partial p_j} > \frac{\partial \tilde{\phi}_i}{\partial p_i} \iff \frac{1+\beta p_i}{1+\beta p_j\tilde{\sigma}_j} > \beta p_j. \tag{A.12}$$

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.

B Robustness Results

Our results concerning how the probability of feedback (measured by ability and intentions) affects sharing 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 B.1. We provide detailed descriptions and summary statistics for these control variables in Table C.1.

B.1 Personal Connections

Indicators of personal connections or social bonds (i.e., number of people a participant recognizes in the experimental session [*Acquaintances*] and number of people in the session a participant considers friends [*Friends*][39]) do not affect our results for the probability of feedback (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.

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.[40] In order to

[39]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.

[40]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).

Table B.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 receiving feedback 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 C.1 for more detailed definitions and descriptive statistics of these co-variates. We reproduce model (V) from Table 3 with the main results in the first column. The number of observations is the number of Round 1 decisions by FM 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 3:(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

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

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' perception 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 effects 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).[41] 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.

[41] 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).

C Additional Figures and Tables

Figure C.1: Player A's Expectations

This figure provides histograms (left scale; bars) and kernel density estimates (right scale; curve) of player A's expectations in Round 1. The percentage numbers indicate the size of three subgroups of expectations: "low" expectations for $\tilde{\sigma}_B \in [0\%, 33\%]$, "medium" expectations for $\tilde{\sigma}_B \in (33\%, 66\%]$, and "high" expectations for $\tilde{\sigma}_B \in (66\%, 100\%]$.

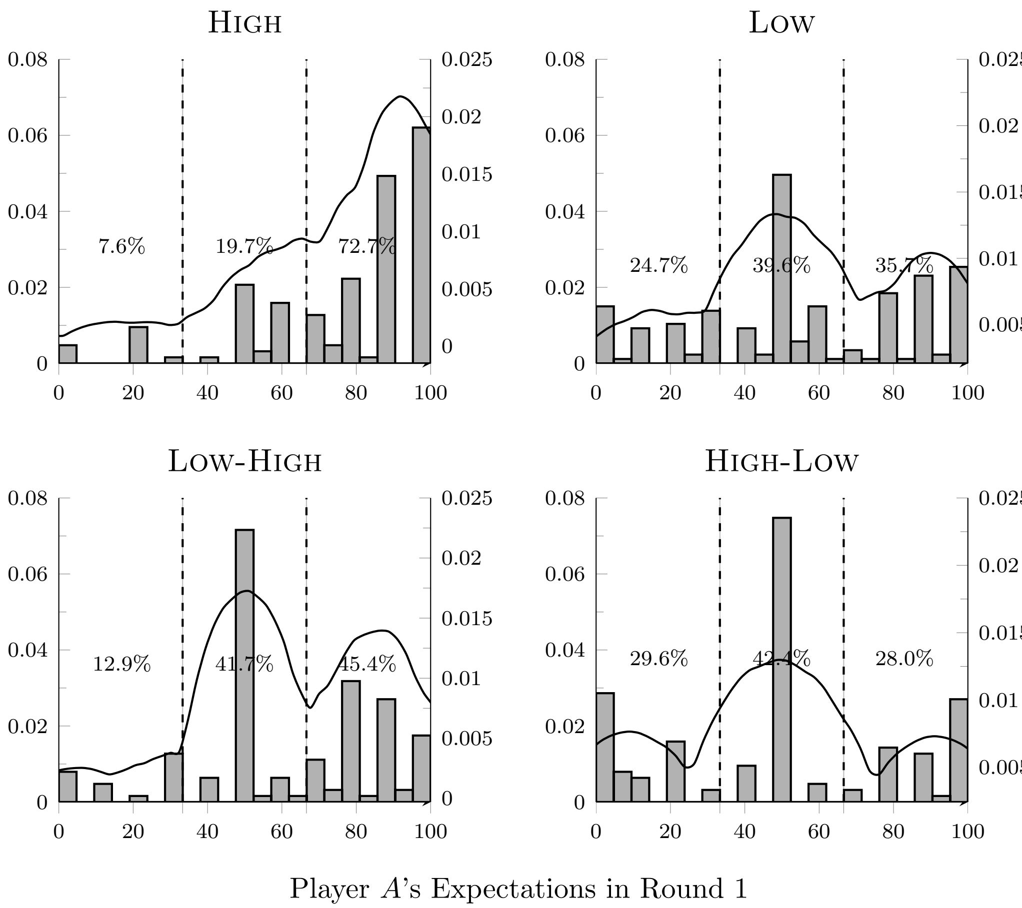

Table C.1: Definitions and Summary Statistics of Explanatory Variables

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 match partner's success probability.
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 match partner's success probability.
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 match partner who terminated the match by choice (i.e., concealed an idea) 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 (i.e., concealed an idea) 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, that means, that selected 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