

NBER WORKING PAPER SERIES

CONVERSATIONS AMONG COMPETITORS

Jeremy C. Stein

Working Paper 13370

<http://www.nber.org/papers/w13370>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

September 2007

This research is supported by the National Science Foundation. I am grateful to Sam Hanson for outstanding research assistance. Thanks also to Oliver Hart and seminar participants at Dartmouth University for helpful comments. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Jeremy C. Stein. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Conversations Among Competitors
Jeremy C. Stein
NBER Working Paper No. 13370
September 2007
JEL No. D82,D83

ABSTRACT

I develop a model of bilateral conversations in which players may honestly exchange ideas with their competitors. The key to incentive compatibility is a strong form of complementarity in the information structure: a player can only generate a useful new insight on a given topic if he has access to his counterpart's previous thoughts on the topic. I then embed this model into a linear social network in which player A first can have a conversation with player B, then player B can have a conversation with player C, and so on. I show that relatively underdeveloped ideas can travel long distances over the network and thus be shared by many agents. More valuable ideas, by contrast, tend to remain localized among small groups of agents.

Jeremy C. Stein
Department of Economics
Harvard University
Littauer 209
Cambridge, MA 02138
and NBER
jeremy_stein@harvard.edu

I. Introduction

Conversation is a central part of economic life. A wide range of information gets passed along from one person to another via word-of-mouth communication, and a number of authors have argued that the nature and extent of connections between people—i.e., the structure of social networks—can have a profound influence on how far and how efficiently information spreads across the economy.¹ For example, in the context of financial markets, Shiller (2000, p. 155) writes: “Word-of-mouth transmission of ideas appears to be an important contributor to day-to-day or hour-to-hour stock market fluctuations...”

In many sorts of conversations, it can be taken more or less for granted that the participants will communicate honestly with one another. If one friend asks another for her opinion of a restaurant, movie, vacation spot, etc., it is hard to see why the response would be anything other than completely truthful. However, there are also a number of important cases where it is less immediately obvious that an honest exchange of information can be expected. Often the participants in a conversation are competitors with one another, at least along some dimension, which means that when one player’s information set is improved, the other player may be made worse off. Consider the following illustrations:

1. *Professional money managers*: Hong, Kubik and Stein (2005) document that mutual-fund managers in a given city tend to have trading behavior that covaries more strongly with other managers in the same city, as opposed to with managers in different cities. Cohen, Frazzini and Malloy (2007) uncover a similar correlation structure in the

¹ See, e.g., Ellison and Fudenberg (1995) for a theoretical treatment of word-of-mouth communication. Jackson (2005) provides a recent survey of the literature on social networks.

trades of mutual-fund managers who went to college together. In the spirit of the Shiller quote above, these findings can be interpreted as evidence of word-of-mouth communication among those fund managers who are socially connected to each other. There is also a good deal of anecdotal evidence that hedge-fund managers share ideas and information with one another.² Indeed, recent theoretical work is beginning to explore the asset-pricing implications of information-sharing along social networks, while maintaining the assumption that such information-sharing is feasible (Colla and Mele (2005), Ozsoylev (2005)). But given that professional money managers have strong incentives to care about relative performance, it is not clear why, e.g., one would tell another honestly about an attractive trading opportunity that he has discovered.³

2. *Knowledge spillovers in Silicon Valley*: Many observers argue that the free flow of ideas across firms has been a key factor in fostering the high rate of technological progress in the Silicon Valley. For example, Saxenian (1994, p. 2-3) writes:

“Silicon Valley has a regional network-based industrial system that promotes collective learning and flexible adjustment among specialist producers of a complex of related technologies...*Companies compete intensely while at the same time learning from one another about changing markets and technologies through informal communications...*” (Italics added.)

Yet if Silicon Valley firms do in fact compete intensely with one another, wouldn't any one executive be tempted either to withhold valuable information from his peers in other firms, or to actively mislead them, in the course of an informal conversation?

² See, e.g., Fennell (2005), who writes of hedge funds: “In a close-knit business where lots of people know each other well, it is inevitable that gossip, rumors and hard facts are being exchanged all the time.”

³ To be more precise, such an honest exchange of information would seem particularly puzzling if it happens before either party has taken a position in the asset in question. Of course, if one manager already has a long position in a stock, he might wish to find a way to credibly communicate this to other managers, so that their buying pushes up the price of his position.

In this paper, I propose a theory of incentive-compatible information exchange among players who, as in the above examples, are in competition with one another. The theory has two key ingredients. First, information flows in both directions—from player A to player B and vice-versa—during the course of a conversation, as players quite literally take turns bouncing ideas off of one another. This differs from the classic framework of Crawford and Sobel (1982). Like I do, Crawford and Sobel pose the question of whether “cheap talk” can be credible in a situation where the two parties involved have partially conflicting interests. But in their model, one party is always the better-informed “sender”, and the other is always the less-well-informed “receiver”, so there is no scope for two-way communication.⁴

The second key ingredient in the model is a strong form of complementarity in the information structure. In particular, player B can only come up with the next useful idea in stage t of a conversation if player A has honestly disclosed his idea from stage $t-1$.⁵ This complementarity, when combined with the two-way nature of the conversation, is what allows for incentive-compatible dialogue. On the one hand, if player A deviates from a truth-telling equilibrium at stage $t-1$ by lying to player B, he makes B less well-informed—which player A prefers all else equal, since B is his competitor. On the other hand, he also makes it impossible for B to come up with a further idea at stage t which builds on A’s previous insight, and which might have been passed back to A had the conversation continued in an honest fashion. This is obviously a direct cost to player A.

⁴ The same can be said about a number of recent theoretical papers that study the process of verbal communication, e.g., Krishna and Morgan (2004), and Dewatripont and Tirole (2005). Although they substantially enrich the basic Crawford-Sobel (1982) model, they continue to maintain the sender-receiver dichotomy, so that information only ever flows in one direction.

⁵ Hellman and Perotti (2005) also emphasize the importance of complementarities in the production of ideas, though their focus is on a different set of issues.

I show that if this latter effect is strong enough, it can be incentive-compatible for players to share their ideas truthfully with one another at all stages of a conversation.

The model can also be used to study the dynamics of information diffusion along social networks that encompass more than two players. Suppose that after A and B conclude their conversation, B can choose whether or not to tell a third player C what he has learned from the conversation, in the hopes that C may be able to further elaborate on this idea. I show that B will only initiate such a conversation with C if, loosely speaking, his informational advantage at the outset is not too big—i.e., so long as B's prior conversation with A was not too productive.

This logic leads to a novel account of information diffusion in which relatively underdeveloped ideas can travel long distances and be shared by many agents, but in which more fully realized—and hence more valuable—ideas tend to remain localized among just a few agents. In other words, even if all agents in the economy are connected to the same network, incentive-compatibility considerations create an endogenous tendency for the most valuable ideas to stay confined to small groups. I argue that this puts an upper bound on the extent of conformity among agents in a large network. For example, in a stock-market setting, it is possible that a given signal is spread by word-of-mouth to a large number of fund managers. But in this case, the signal must be a relatively weak one, so that it only induces a modest covariance in the trading behavior of these managers. In contrast, if a signal is shared by a small number of fund managers, it can be a much stronger one that leads their trades to covary more closely.

The remainder of the paper is organized as follows. Section II presents a basic model of a single conversation. Here I consider two players A and B, whose interests are

partially at odds with one another, and ask under what circumstances they will be able to sustain an honest exchange of ideas. Section III develops an extension of the basic framework in which there are many players A, B, C, D, etc., who are arrayed along a straight-line network, and who can have sequential conversations with their most proximate neighbors. Thus after A and B conclude a conversation, B can initiate one with C, and so forth. This extension allows one to ask how far a given set of ideas can travel along the network, and how the distance traveled relates to the quality of the ideas. Section IV discusses the model's implications for Shiller's (2000) hypothesis, namely that word-of-mouth transmission of ideas can have significant consequences for asset prices. Finally, Section V concludes.

II. The Basic Model: A Single Conversation

A. Information structure

There are two players, A and B, engaged in a conversation. They take turns moving, so that A moves at times 1, 3, 5, etc., and B moves at times 2, 4, 6, etc. At time 1, player A has access to an initial idea χ_1 that he can choose to communicate to player B. This idea has two uses. First, anybody observing χ_1 can eventually “decode” it—i.e., can map it into a payoff-relevant signal $s_1(\chi_1)$. This decoding is assumed to happen after the conversation has ended and A and B have gone their separate ways.

Second, if player B gains access to χ_1 , he has a probability p of coming up with another idea χ_2 at time 2 that builds on and refines χ_1 , and that can be mapped into a second payoff-relevant signal $s_2(\chi_2)$. Crucially, if idea χ_1 is not truthfully revealed to player B, he can never come up with χ_2 . This assumption embeds a strong form of

complementarity into the production function for ideas: a useful new idea can only be produced by an agent who has access to the prior idea.

The process can potentially continue back and forth indefinitely, until one player fails to come up with an idea. In particular, if at any odd date t , player A has observed player B's prior-period idea χ_{t-1} , he has a probability p of coming up with another new idea χ_t , which he can then choose to communicate to player B. Again, it is impossible for A to generate idea χ_t without having been exposed to a complete and truthful description of idea χ_{t-1} . And whoever observes χ_t can, after the conversation has ended, go off and generate the signal $s_t(\chi_t)$.

Of course, a player in possession of an idea need not communicate it truthfully. In particular, if player A has come up with an idea χ_t at time t , he may either: i) not say anything to player B, effectively pretending that he was unable to come up with an idea; or ii) pass along a bogus (i.e., informationally useless) version of the idea, χ_t^b . I assume that the former option is costless, but that the latter imposes on A an arbitrarily small positive cost of ε .⁶ Moreover, in the latter case, player B does not immediately recognize that he has been lied to—he takes the statement of χ_t^b at face value and attempts to come up with a further insight that builds on it. It is only after the conversation ends, when B attempts to decode the ideas that he has collected, that he learns that χ_t^b is useless: at that point, he is unable to turn χ_t^b into a payoff-relevant signal.

⁶ This is a technical assumption whose purpose will become clear shortly. One can motivate it by saying that it takes a little time and effort to come up with a good lie, or that lying is personally distasteful.

B. Payoffs

At the conclusion of the conversation, the players attempt to decode the ideas they have obtained, i.e., they try to turn these ideas into payoff-relevant signals. Let n_A represent the total number of signals that player A is able to produce (which equals the total number of non-bogus ideas that he has collected), and let n_B represent the total number of signals that player B is able to produce. To capture the notion that the players are competitors, A's payoff should be an increasing function of n_A , and a decreasing function of n_B , and vice-versa.

In order to generate payoff functions with transparent microfoundations, I consider a specific competitive setting with a very simple market structure. As will become apparent, this setting connects more directly to the Silicon Valley production-technology example sketched in the Introduction than to the fund-manager example.

Think of the signals generated by a conversation as elements of a recipe for building a new product more cost-efficiently. In particular, if player A has access to n_A signals, he can manufacture the product for a cost of $(1 - h(n_A))$, and symmetrically for player B. Here $h(n)$ is an increasing function that captures the total cost savings associated with n signals. I will put more structure on the $h(n)$ function below, but for the moment I leave its general form unspecified, and just impose the conditions that $h(0) = 0$, and $h(\infty) < 1$.

I assume that player A and player B each face a unit mass of customers for the new product, and that all customers have a reservation value of one. Moreover, there is a fractional overlap of θ in A's and B's customer bases, with $0 < \theta < 1$. In other words, A has a monopoly on a fraction $(1 - \theta)$ of his customers, but must compete with B for the

remaining fraction θ . I further assume that the products are otherwise undifferentiated, so that when competition between A and B does occur, it is à la Bertrand. In what follows, θ can be thought of as a proxy for the degree of competition between A and B.

Taken together, the assumptions imply that the payoffs to player A and player B are respectively given by:

$$U_A = (1 - \theta)h(n_A) + \theta \max\{0, h(n_A) - h(n_B)\} \quad (1)$$

$$U_B = (1 - \theta)h(n_B) + \theta \max\{0, h(n_B) - h(n_A)\}. \quad (2)$$

The first part of (1) corresponds to the fact that for a fraction $(1 - \theta)$ of his customers, A is a monopolist and charges the full reservation value of one; with a cost of $(1 - h(n_A))$, his profits per customer are thus $h(n_A)$. On the remaining fraction θ of his customers, where A and B overlap, Bertrand competition implies that A only makes a profit to the extent that his costs are strictly below those of B.⁷

To give a concrete example, suppose that player A has been exposed to ideas χ_1 through χ_5 , while player B has only been exposed to ideas χ_1 through χ_4 ; this would happen if, after having observed χ_5 , A decided either to remain quiet, or to make a dishonest report of χ_5 to B. It follows that player A can generate the first five signals s_1 through s_5 , while player B can only generate the first four signals s_1 through s_4 . Thus we have $U_A = h(5) - \theta h(4)$, and $U_B = (1 - \theta)h(4)$.

⁷ To be clear, I am assuming that each player can price discriminate—i.e., A can charge a different price to those customers who are exclusively his, as opposed to those he shares with B.

C. Benchmark case: no complementarities in idea production

Before proceeding, it is useful to establish a benchmark result, namely that conversations can never arise absent complementarities in idea production of the sort described above. The following proposition summarizes the result.

Proposition 1: Suppose that a player's probability of having an idea χ_t at time t —and hence of being able to generate the signal $s_t(\chi_t)$ —is fixed at p , and is independent of whether he has been given a truthful report of the other agent's prior idea χ_{t-1} . In this case, it can never be part of an equilibrium for players to truthfully exchange ideas with one another.

Proof: Assume to the contrary that there does exist an equilibrium in which the players truthfully exchange ideas with one another over at least some interval. Taking player B's honest reporting strategy as given, at any time t player A will always want to deviate from the proposed equilibrium, and mislead player B with a dishonest report of his idea—i.e., A will want to report χ_t^b instead of χ_t . To see why, note that such a dishonest report makes it impossible for B to generate the signal $s_t(\chi_t)$, thereby leaving B with one less signal than A. Moreover, since B does not know (until after the conversation has ended) that the report was dishonest, his reporting strategy cannot punish A for the deviation. Thus A does not experience any reduction in his total number of signals n_A as a result of deviating; the only cost he bears is the ε cost of inventing the bogus report. With n_A unchanged, and n_B reduced to a value strictly lower than n_A , (and

with ε assumed to be arbitrarily small), equation (1) tells us that the deviation strictly increases player A's payoff U_A .

Intuitively, without complementarities in idea production, player A can only gain by deviating from the honest equilibrium and lying to player B about his ideas—such a lie leaves B less well-informed, and taking B's truthful-reporting strategy as fixed, does not degrade the quality of the information that A expects to get back from B. The fact that it is impossible to sustain a cooperative outcome in this setting, in spite of an infinite horizon, distinguishes the no-complementarities version of the model from, e.g., an infinitely-repeated prisoner's dilemma game. The key distinction is that in the prisoner's dilemma case, a deviation from the cooperative strategy is immediately observable and hence can be punished in a tit-for-tat fashion, while the same is not true here, since lies cannot be detected until after the conversation has concluded.

This logic also suggests why complementarities in idea production are necessary to support the honest exchange of ideas. In a setting with complementarities, if A misleads B about his idea χ_i , A also stands to bear a cost. This is because without access to an honest report of χ_i , B can no longer come up with the next useful idea, which—assuming that B is playing an honest-reporting strategy—could otherwise have been bounced back to A.

D. Sustaining a conversation: necessary conditions

I now return to the case where there are complementarities in idea production, and explore the conditions under which a cooperative conversational equilibrium can be

sustained. To do so, I hypothesize that such an equilibrium exists, and then check that at each date t , neither player has an incentive to deviate from the proposed equilibrium.

In a fully cooperative equilibrium, each player truthfully reports any idea that he has to the other player. This continues until one player fails to come up with an idea, at which time he truthfully announces this as well, and the conversation concludes, with each player having access to the same total number of ideas.

Now suppose it is date t , and that player i has just come up with an idea χ_t . If he reports it honestly, and the game continues along the equilibrium path until somebody fails to come up with an idea, player i 's expected payoff, $EU_i(\text{continue}@t)$, is given by:

$$EU_i(\text{continue}@t) = (1 - \theta)H(t), \quad (3)$$

where:

$$H(t) = \sum_{i=0}^{\infty} p^i (1 - p)h(t + i) \quad (4)$$

In words, $H(t)$ is the expected value of the total cost reduction $h(\)$ that both players will realize if player i shares χ_t and they continue to play cooperatively from that point on. For example, with probability $(1 - p)$, there will be no further ideas after time t , so each player will wind up with t signals that have a total cost-cutting value of $h(t)$; with probability $p(1 - p)$, there will be exactly one further idea after t , so each player will wind up with $(t + 1)$ signals that have a total cost-cutting value of $h(t + 1)$; and so forth.

By contrast, suppose that player i considers deviating from the proposed cooperative equilibrium at time t . In principle, this deviation could take the form of i either not reporting his idea at all to the other player, or inventing a bogus report. However, unlike in the no-complementarities case of Proposition 1, now if i is going to deviate, it will always be a dominant strategy for him to do so by simply not reporting anything. This is because in a world with complementarities, *any* deviation effectively ends the game from his perspective—once he stops feeding honest information to his counterpart, he can never get anything of value back. So there is no point in spending the ε cost of inventing a bogus report, when keeping quiet is a costless way to achieve the same outcome.

Thus we can focus our attention on deviations which take the form of player i putting an end to the conversation by not reporting an idea χ_i that he is in possession of at time t . In this case, his payoff, $U_i(\text{stop}@t)$, is given by:

$$U_i(\text{stop}@t) = h(t) - \theta h(t - 1). \quad (5)$$

This expression reflects the fact that if i deviates, he keeps idea χ_i to himself, and therefore winds up with one more signal than the other player (t signals versus $(t - 1)$). This allows him to not only earn a profit of $(1 - \theta)h(t)$ from the part of the market that he has all to himself, but also a profit of $\theta(h(t) - h(t - 1))$ on the part of the market that he has to compete for with the other player, since he now has a cost advantage.

In order for the cooperative equilibrium to hold together, we require that: $EU_i(\text{continue}@t) \geq U_i(\text{stop}@t)$ for all values of t . Substituting in from equations (3) and (5), and doing some rearranging, this requirement can be expressed as:

$$\frac{(H(t) - h(t-1))}{(h(t) - h(t-1))} \geq \frac{1}{(1 - \theta)} \quad (6)$$

Thus we have:

Proposition 2: If condition (6) is satisfied for all values of t , it is possible to sustain a conversation in which each player truthfully reports every idea he has to the other player, until one of the two players fails to come up with a further idea.

Note that once player i does fail to come up with an idea, he will always announce this failure truthfully—thereby ending the conversation—as opposed to inventing a bogus idea and passing it on to the other player. Again, this is because he can derive no further benefit from the other player either way, and inventing a lie entails an extra cost of ε .

E. A parametric example: geometric decay in the value of successive ideas

As currently stated, the necessary condition in (6) is not transparent, because it depends on the shape of the $h(n)$ function, whose form I have not yet specified. To make things more intuitive, it is helpful to put some structure on the $h(n)$ function. One particularly tractable choice is the following:

$$h(n) = (1 - \beta^n) \tag{7}$$

with $0 < \beta < 1$. This formulation corresponds to the assumptions that: i) as a conversation goes on, each new idea is of less value than the previous one; and ii) that this decay is geometric in nature. In particular, the incremental cost savings associated with the n th signal is given by $\beta^{n-1}(1 - \beta)$. Thus the first signal lowers costs by $(1 - \beta)$, the second signal further lowers costs by $\beta(1 - \beta)$, and so forth.

With this functional form, condition (6) can be re-written more specifically as:

$$\frac{(\beta^{t-1} - \sum_{i=0}^{\infty} p^i (1-p)\beta^{t+i})}{(\beta^{t-1} - \beta^t)} \geq \frac{1}{(1-\theta)} \tag{8}$$

Condition (8) in turn can be greatly simplified to yield:

$$\frac{1}{(1-p\beta)} \geq \frac{1}{(1-\theta)}, \tag{9}$$

or, alternatively,

$$p\beta \geq \theta. \tag{10}$$

The inequality in (10) captures in a very simple way the three main factors that make it possible for conversations to be sustained. First, and most obviously, it helps if

the competition parameter θ is small, so that one player only suffers a little bit when the other's information set is improved. Second, it also helps if p is large, which means that when one player truthfully reveals an idea to the other at time t , there is a relatively high likelihood that another idea will be bounced back to him at time $t+1$; this makes continuing the conversation more attractive relative to deviating and cutting it short. And finally, incentives for cooperation are also stronger when the decay parameter β is closer to one; this means that the marginal value of incremental ideas declines slowly, which also increases the appeal of continuing the conversation as opposed to cutting it off.

It is worth noting that the simple condition in (10) is independent of t . In other words, the incentive to continue a conversation does not depend on how far into it the players are. This is true even though the amount of information that is yet to be gained decreases with t . Intuitively, there are two offsetting effects. On the one hand, as time passes, the absolute appeal of continuation declines. On the other hand, so does the absolute appeal of deviating, and thereby staying one step ahead of one's competitor. Since all that matters is the ratio of these two quantities, they offset each other exactly, causing t to drop out of the comparison. This can be seen explicitly by looking at (8), where the numerator of the left-hand-side (which reflects the appeal of continuation) and the denominator (which reflects the appeal of deviating) are both proportional to β^{t-1} , causing β^{t-1} to drop out when the expression is simplified.

III. Sequential Conversations: Information Diffusion On a Social Network

A. Extending the basic model

I now augment the model so that there are multiple players (A, B, C, D, etc.) arrayed along the simplest possible social network, namely a straight line. I assume that each player can only interact with his most immediate neighbors, and that this interaction happens sequentially, starting with players A and B. Thus after A and B conclude their conversation, B must choose whether or not to initiate a conversation with player C. If this happens, then when B and C conclude their conversation, C can choose whether to start one with player D, and so on.

In order to generate the players' payoff functions, I consider the natural extension of the market structure from the two-player case. Specifically, every player in the interior of the network (i.e., every player but A) now has a customer base that overlaps partially with his neighbors on either side. For example, of the unit mass of customers facing player B, a fraction θ overlaps with player A, another fraction θ overlaps with player C, and a fraction $(1 - 2\theta)$ are customers who are exclusive to B—i.e., customers over whom B has monopoly power. (Note that I must now assume that $\theta < 1/2$, which is innocuous.)

It follows that player A's payoff function remains as in equation (1), while the payoff function for player B is modified to:

$$U_B = (1 - 2\theta)h(n_B) + \theta \max\{0, h(n_B) - h(n_A)\} + \theta \max\{0, h(n_B) - h(n_C)\}. \quad (11)$$

An analogous expression holds for the payoffs of each of the players who come after B.

B. Necessary conditions for a second conversation to get started

Assume for the moment that the first conversation—that between A and B—is itself incentive-compatible. As we have seen, this conversation will go on until one of the two players is unable to come up with a further refinement. At this point, A and B will part ways, with each of them being symmetrically informed. Let k denote the random number of signals that A and B are able generate after their conversation concludes. That is, both A and B have access to idea χ_k at the end of their conversation, and can therefore both manufacture signals s_1 through s_k .

Next, B has to decide whether or not to turn around and share idea χ_k with player C. As before, the potential advantage of doing so is that player C will, with probability p , be able to come up with a further refinement χ_{k+1} , and that the conversation between B and C will continue on for several more stages from there. The disadvantage to B of starting a conversation with C is that by giving C access to idea χ_k , B effectively gives C *all* of the first k signals s_1 through s_k —i.e., B repeats everything of value that was learned during the A-B conversation. The implicit assumption here is that the refined idea χ_k also embodies all of the cumulative knowledge in the previous-stage ideas χ_1 through χ_{k-1} .

Player B's incentive to initiate a conversation with C is therefore lower than A's original incentive to initiate a conversation with B. This is because when A first speaks to B, he gives away a brand-new idea that is equivalent to just one signal, in the hopes of getting one more signal in return. By contrast, when B first speaks to C, he gives away a more-fully-developed idea that is equivalent to k signals, again in the hopes of getting one more signal in return. Clearly, it is less desirable to give away a more fully-developed idea to one's competitor.

To make this all precise, observe that at the conclusion of the A-B conversation, with both A and B in possession of k signals, we can rewrite B's payoff as:

$$\begin{aligned} U_B(k) &= (1 - 2\theta)h(n_B) + \theta \max\{0, h(n_B) - h(k)\} + \theta \max\{0, h(n_B) - h(n_C)\} \\ &= (1 - \theta)h(n_B) - \theta h(k) + \theta \max\{0, h(n_B) - h(n_C)\}. \end{aligned} \quad (12)$$

Based on (12), player B reasons as follows. If he reveals his information to player C, and the conversation continues from there, player B's expected payoff from initiating the B-C conversation is given by:⁸

$$EU_B(\text{initiate@BC}) = (1 - \theta)H(k) - \theta h(k). \quad (13)$$

Note that conditional on completing a conversation with player C, player B does not care if C then goes ahead and improves his information set via a subsequent conversation with player D. This is because once B and C have access to the same number of signals, B can never earn a profit from those customers they have in common; this does not change if C eventually becomes even better informed than B.

If instead player B decides not to speak to player C, thereby eliminating the potential for any further conversations along the network, he walks away with a payoff equal to:

⁸ One still has to check the incentive-compatibility conditions that ensure that once B and C have started a conversation, they will both want to continue it until they run out of ideas. However, it is easy to show that these conditions are implied by (6), the necessary condition for sustaining the first conversation between A and B. Also, note that once a conversation is started, C will always have a stronger incentive to continue it than B. In other words, the incentive constraint will be slacker for C than it is for B. This is because C has the option value associated with the possibility of moving on to a further conversation with D.

$$U_B(\text{withdraw@BC}) = (1 - \theta)h(k). \quad (14)$$

Therefore, a necessary condition for player B to initiate a conversation with C is:

$$\frac{H(k)}{h(k)} \geq \frac{1}{(1 - \theta)} \quad (15)$$

Comparing (15) to (6), it is apparent that, for $k = t$, (15) implies (6), so that it is harder to satisfy (15), all else equal. This just formalizes the intuition stated above: at any given point in time t , it is easier to get a player to continue an existing conversation than it is to get a player with an equally well-developed idea to initiate a new conversation: in the former case, the player is only being asked to give away one incremental signal, while in the latter, he is being asked to start things off by giving away k signals.

C. When does information diffusion comes to a stop?

For concreteness, the remainder of the analysis in the paper focuses on the parametric specification described above, that where the value of new signals decays geometrically. With this specification, condition (15) boils down to:

$$\frac{1 - (\beta^k (1 - p) / (1 - p\beta))}{(1 - \beta^k)} \geq \frac{1}{(1 - \theta)}. \quad (16)$$

Observe that the previous necessary condition for continuing a conversation in the two-player model, as given by (10), is just a special case of (16) with $k = 1$. In other words, if player B winds up his conversation with A in possession of just one signal, his decision of whether to initiate a new conversation with C is identical to his decision of whether to continue an existing conversation in the two-player version of the model, since either choice amounts to giving away one signal in exchange for the same expected informational return.

If (16) is in fact satisfied for a given realization of k , then the conversation between B and C gets off the ground. This conversation then continues stochastically until one of the two is unable to generate a further refinement. Suppose that this happens at the point where B and C can each generate k' signals, with $k' \geq k$. Now C faces the exact same problem in deciding whether to initiate a conversation with D, so he applies the same necessary condition as B did previously, except that k is replaced in (16) with k' .

Using this sequential logic, we can describe the dynamics of information diffusion in the multi-player version of the model. Note that the left-hand side of (16) is strictly decreasing in k . Now define k^* as that value of k for which (16) holds with equality:

$$k^* = \frac{\log(1 - p\beta) + \log(\theta) - \log(\theta(1 - p) + p(1 - \beta))}{\log(\beta)}. \quad (17)$$

The following proposition characterizes the outcome for the geometric-decay case:⁹

⁹ It is possible to show that an analog to Proposition 3 holds for a broader class of payoff functions, $h(n)$. Specifically, suppose that $h(n)$ is a strictly increasing and strictly concave function such that $h(0)=0$, and

Proposition 3: Consider two adjacent players J and K in the linear network, and assume that the condition in (10) is satisfied—i.e., that $p\beta \geq \theta$ —so that once initiated, conversations can always be sustained. Assume further that J and K conclude a conversation with each having the ability to produce k_{JK} signals. If $k_{JK} \leq k^*$, player K then initiates a new conversation with the next player L in the network, and the process of information diffusion continues forward. If $k_{JK} > k^*$, no further conversations are initiated, and information diffusion along the network comes to a stop.

Example 1: Suppose that $\beta = 0.95$, $\theta = 0.1$, and $p = 0.5$. Then according to (17), $k^* = 6.95$. This implies that if a conversation between J and K ends with each player having the ability to produce 6 or fewer signals, K will initiate a new conversation with L. However, if a conversation between J and K ends with each player having the ability to produce 7 or more signals, no further conversations will be initiated.

D. Implications for number of participants, quality of ideas, and conformity

With Proposition 3 in hand, it is possible to establish a number of properties of the sequential model. The statements that follow center on two items of interest. First, the random variable c is defined as the number of conversations that occur in a given play of the game *subsequent* to the initial A-B conversation (which, under the assumption that (10) is satisfied, always takes place). Second, the random variable k , which we have

$\lim_{n \rightarrow \infty} h(n) < 1$, and that the general condition in (6) for sustaining a conversation is satisfied for all t . One can then show that there exists a k^* with $1 \leq k^* < \infty$ such that the conclusion in the proposition continues to apply. Proof available on request.

already encountered, is the duration of the initial A-B conversation. One way to think about k is that it measures the quality of the idea that emerges from this first conversation. Alternatively, k can be interpreted as measure of the degree of pairwise conformity among participants in the conversational chain. Specifically, for any given play of the game, k is equal to the number of signals that the initial player A holds in common with any other player J who has participated in the chain. Thus k is a summary statistic for the extent to which A's and J's production technologies (or forecasts, depending on the setting) resemble one another.

As described in the Introduction, one of the main insights of this paper is that when ideas travel further, they tend to be of lower quality, and hence to induce a weaker degree of conformity among those who have been exposed to them. There are two ways that this insight can be expressed in the language of the model. The first, and perhaps most natural way to do so, is to fix the parameters p , β , and θ , and to make statements about the correlation between the ex-post realizations of c and k . The appendix establishes the following proposition:

Proposition 4: Let ρ_{ck} denote the correlation between the random variables c and k . For any given values of the parameters p , β , and θ such that (10) holds, we have that $\rho_{ck} < 0$.

The logic behind the proposition is straightforward. Suppose that the initial conversation between A and B is blessed by good luck, and continues for many rounds, so that the number of signals k is relatively high. In this case, k is closer to k^* , which

represents the cutoff point beyond which no new conversations are started; it is therefore more likely that the whole process will soon come to a stop. Said a bit differently, the expected number of subsequent conversations conditional on k , $E(c|k)$, is declining in k .

Example 2: Suppose, as in Example 1, that $\beta = 0.95$, $\theta = 0.1$, and $p = 0.5$. Then, using equation (A.5) in the appendix, $\rho_{ck} = -0.402$.

An alternative way to think about the relationships between the number of conversations, idea quality, and conformity, is to compute the ex ante expected values of c and k , and then to see how these expected values vary as we change one of the exogenous parameters, say p . The appendix shows that there are simple closed-form expressions for the expected values of c and k :

Proposition 5: Assume that (10) holds, and define $I(k^*)$ as the greatest integer that is less than or equal to k^* . Then we have: $E(c) = I(k^*)(1-p)/p$, and $E(k) = 1/(1-p)$.

The comparative statics that follow from Proposition 5 are for the most part intuitively obvious. First, it is immediate that $dE(k)/dp > 0$. Since one can interpret p as a proxy for the talent of the players in a given network—more talented players are less likely to draw a blank at any given point in a conversation—this just says that both the expected quality of the idea emerging from the initial A-B conversation, as well as conformity, are increasing in talent. Second, it can be shown that $E(c)$ is weakly

decreasing in θ , which means that the expected number of participants in a conversational chain increases when players are less competitive with one another.

The one comparative static that is a little tricky is $dE(c)/dp$. As it turns out, $E(c)$ is not monotonic in p , because of the discreteness associated with the integer-valued function $I(k^*)$.¹⁰ However, if we get rid of this discreteness issue, it is possible to make a more clear-cut statement. Specifically, let $\Phi\{E(c)\}$ denote the upper bound on $E(c)$ defined by:

$$\Phi\{E(c)\} = k^*(1-p)/p \tag{18}$$

It can then be shown that $d\Phi\{E(c)\}/dp < 0$, as illustrated in Figure 1, which plots both $E(c)$ and $\Phi\{E(c)\}$ as functions of p for $\beta = 0.95$ and $\theta = 0.1$. In other words, as p increases, the upper bound on the expected number of conversations falls. More precisely, an increase in p has two competing effects on $\Phi\{E(c)\}$. First, the expected length of any conversation goes up, and so it becomes more likely that any given cutoff point k^* is reached after a smaller number of conversations. Second, however, the cutoff point k^* is itself an increasing function of p , since the desirability of initiating new conversations rises with talent. As it turns out, the former effect dominates the latter, so that $d\Phi\{E(c)\}/dp < 0$. This is established formally in the appendix.

Thus there is a second sense in which we can think of the number of conversations as being negatively related to both idea quality and conformity. As we

¹⁰ To be more specific, $E(c)$ is monotonic in p almost everywhere—except on the measure-zero set of points where k^* takes on an integer value.

increase the talent p of the players, the expected quality of the idea generated in the initial A-B conversation goes up, as does conformity; this is just the statement that $dE(k)/dp > 0$. And the same increase in talent p also lowers the upper bound on the expected number of conversations, as we have just seen. So if there is variation in p across social networks, it will trace out the aforementioned negative relationships. For example, if one social network is composed of highly-talented managers, and another network is composed of lower-ability managers, the model would predict that conversational chains would be on average shorter in the former case, but that they would tend to transmit higher-quality ideas, and to induce more conformity among those managers that end up participating in the conversations.

Example 3: First, suppose that, as in the previous examples, $\beta = 0.95$, $\theta = 0.1$, and $p = 0.5$. Straightforward calculation based on Proposition 5 yields: $E(c) = 6$, $\Phi\{E(c)\} = 6.95$, and $E(k) = 2$. If we keep everything else the same but set $p = 0.8$, we obtain instead: $E(c) = 4.25$, $\Phi\{E(c)\} = 4.47$, and $E(k) = 5$.

IV. Conversations and Asset Prices

For expositional purposes, I have framed the discussion of the model around a specific example, one in which “ideas” represent technological cost-savings innovations for a particular good, and in which the “competitors” are rival producers of this good. Of course, this is not the only potentially relevant setting. As noted in the Introduction, Shiller (2000) hypothesizes that word-of-mouth transmission of ideas can play an important role in generating asset-price volatility. In this context, “ideas” presumably

represent signals about the expected returns to particular investment strategies, and the “competitors” are professional investors, e.g., hedge-fund managers.

It is interesting to examine Shiller’s (2000) conjecture in light of the model of this paper. A first question that arises is whether the model’s basic information structure—with its strong complementarities in idea production—can be thought of as fitting naturally in a financial-markets setting. This “standing-on-the-shoulders” formulation is familiar in work on technological innovation (see, e.g., Aghion and Howitt (1992)). However, it represents something of a departure relative to the finance literature, where information is typically modeled in terms of a series of additive signals about the payoffs to various securities.

A. Complementarities in the production of investment ideas: an example

In an effort to make the case that the model’s information structure can actually be quite relevant for thinking about financial markets, I present the following imaginary conversation between two would-be arbitrageurs.¹¹

1. Trader A: “Here’s an idea for you. Did you see that Pauly Pipelines (PPL) announced a management-led LBO bid at \$120/share? Of course their bonds got killed—spreads widened by 90 basis points. But the funny thing is that the bonds of Pauly Exploration (PEX) also got hammered—spreads widened by almost 70 basis points. That makes no sense. Sure, they have the same name and were founded by the same guy fifty years ago, but they are now totally separate legal entities. And there’s no way anybody is going to try to do an LBO of PEX. So I can’t see why a big added debt load for PPL has any implications whatsoever for PEX debt. The market must have just gotten the names confused. I’m thinking of making a big long bet on PEX bonds—spreads ought to tighten back up within a few days.”

¹¹ Although the conversation is purely fictitious, it is loosely inspired by events surrounding the announcement of a management-led buyout bid at pipeline operator Kinder Morgan in May of 2006.

2. Trader B: “I didn’t notice the move in the PEX bonds. But are you sure about this? They may be separate legal entities, but I think there is still a lot of overlap on their boards. I’m pretty sure the Pauly family trust has a couple of seats on each. What I’d worry about is this: if PPL is struggling with the debt load, they might try to make a deal to sell some of their pipeline assets to PEX. If that happens, and PEX finances the purchase with debt, then they wind up very leveraged too. In which case it makes sense for PEX debt to be trading down.”

3. Trader A: “I guess that’s possible. I mean, it makes sense if current PPL management stays in control—they’d be in a good position to cut an asset-sales deal with PEX. But what if this is just the first shot in a bidding war and some third party winds up doing the LBO instead and taking control of PPL? That would still be just as bad for PPL bonds, but since the third party is unlikely to then play ball with PEX, it would mean that PEX bonds should not get hurt.”

4. Trader B: “Does the market expect another bidder to come in? You said the management bid was for \$120, and PPL stock went to—wait, let me check my screen—\$135. So yeah, you’re right, the market does seem to be building in a real chance of a bidding war....Oh wait, I got it. Here’s the trade you want to do. You’re saying that PEX bonds have a better chance of reverting if a bidding war breaks out for PPL stock, but that you can get killed if PPL management stays in control at \$120. So you want to be long PEX bonds and also long a bunch of puts on PPL stock with a strike at say \$135. That way, if no other bidder shows up and PPL falls back to \$120, you’re covered.”

5. Trader A: “Yeah, something like that. That sound like a trade you might put on?”

6. Trader B: “I might. I need to think about it some more. But it doesn’t sound totally crazy.”

The above conversation has exactly the structure envisioned in the model. First, when trader A initiates the conversation at time 1, he gives trader B a valuable piece of information, namely the fact PEX bonds appear to be underpriced. In principle, B could end the conversation right there, and simply go off and exploit this information on his own. To the extent that A and B are rivals, this would clearly be harmful to A. However,

access to A's information allows B to generate a further insight at time 2 (about the potential for PPL to sell assets to PEX), which increases the value of A's original idea and which B then passes back to A; had A not shared his idea honestly with B at time 1, he would never have been able to realize this time-2 improvement in his information set.

Similarly, trader B's decision to share his time-2 insight with A is validated by the fact that this in turn prompts A to come up with yet another refinement at time 3 (about the possibility of a bidding war and a third party gaining control of PPL), which A then bounces back to B. This process continues back and forth through time 4, at which point the two traders have jointly developed an arbitrage strategy that is more sophisticated than either one could have arrived at alone.

B. Can conversations have large price effects?

If one accepts the model's information structure as a reasonable description of some financial-market interactions, it follows that the model can help to rationalize the empirical findings of Hong, Kubik and Stein (2005), and Cohen, Frazzini and Malloy (2007). These papers show that there are clusters of correlation in the trades of those mutual-fund managers who are either located in the same city, or who went to college together—patterns that are strongly suggestive of word-of-mouth diffusion of investment ideas. And notably, these effects show up among agents who have clear incentives to care about their performance relative to one another.

Beyond such local-correlation patterns in trading behavior, can the mechanism in the model also generate economically interesting price effects of the sort envisioned by Shiller (2000)? Here the answer is less clear-cut. It stands to reason that a given idea

will have more of an impact on stock prices to the extent that two conditions are satisfied: i) the idea reaches a large number of investors; and ii) its information content is substantial, so that each investor who comes into contact with the idea revises his expectations by a meaningful amount.

However, a central message from the model is that there can be a fundamental tension between these two conditions: those ideas that reach the largest numbers of investors are precisely those that tend to have relatively little real information content.¹² Indeed, in the limit, such far-traveling ideas may not be much more substantial than mere gossip. This in turn would seem to imply an endogenous upper bound on the magnitude of any word-of-mouth effect.

At the same time, the model also highlights the conditions under which the upper bound is likely to be a loose one, so that Shiller's hypothesis is on stronger ground. First, and most transparently, as the competition parameter θ approaches zero, we have seen that the expected number of conversations becomes large, without any offsetting effect on the quality of ideas. This suggests that conversations among individual investors (who are likely to be less directly competitive with one another) could potentially wind up having bigger equilibrium effects on prices than conversations among professional

¹² To map the model more literally into a financial-markets setting, suppose that each player is interested in forecasting the payoff V on a security, which can be written as the following infinite sum: $V = s_1 + \dots + s_\infty$. Suppose further that s_i is a mean-zero normal random variable with a variance of β^i . This formulation mirrors the set-up in the text: the marginal value of each successive signal—here in terms of the fraction of the variance of V that it explains—declines geometrically. And a low-information-content, far-traveling idea is one that is comprised of just a small number of signals (i.e., is given by $s_1 + \dots + s_k$, for k not much larger than one), and hence that explains only a small fraction of the overall variance of V .

money managers, even though the latter sorts of conversations do occur, and may be quite informative when they do.¹³

A second observation is that the upper bound argument relies critically on the premise that all investors are rational Bayesians, so that the stock-price impact associated with a given idea is proportional to its true informational content. Or said a bit differently, the assumption is that investors do not overreact to the news that they gather from their conversations with others. In reality, however, it may be that information that is obtained via direct face-to-face interaction is viewed by investors as particularly salient and compelling, especially when compared to, e.g., relatively dry earnings releases that they might read in a newspaper. If so, even low-quality rumors might have significant consequences for prices, which clearly helps Shiller's story.

V. Conclusions

The transmission of information by word-of-mouth occurs in a wide range of economic settings, including those in which agents have somewhat competing interests. I have argued that in such competitive environments, an information structure with strong complementarities can be the glue that holds conversations together. Specifically, each participant in a conversation will have an incentive to honestly disclose a given idea to his counterpart, in the hopes that the counterpart will be able to take it one step further, and will then bounce the more fully-developed idea right back to him.

¹³ An obvious counterpoint is that professional managers control much larger portfolios, so even if an idea reaches only a few of them, it can translate into more trading volume than if the same idea had reached hundreds of individual investors.

While an information structure with complementarities makes conversations among competitors feasible, it does not follow that it removes all barriers to communication. Thus while ideas can travel across networks of competitors who are connected to one another, they do not necessarily travel very far. Indeed, the more successfully an idea develops in early-stage conversations, the more likely it is to remain localized among the handful of players who were its originators. In contrast, low-quality ideas—those not much more informative than gossip—may ultimately be very widely diffused.

Appendix

The proof of Proposition 4 builds on that of Proposition 5, so these two results are proved in reverse order.

A. Proof of Proposition 5: Consider a given path of play in the equilibrium described in Proposition 3. Until there are more than k^* signals available for decoding, a new conversation is initiated each time a conversant fails to generate an additional signal. Since player A is assumed to start with an initial signal, c is the random number of failures in a series of independent Bernoulli (p) trials before there are $I(k^*)$ successes. A random variable such as c with

$$\Pr(c = i) = \binom{i + I(k^*) - 1}{I(k^*) - 1} p^{I(k^*)} (1 - p)^i \text{ for } i = 0, 1, 2, \dots \quad (\text{A.1})$$

is said to follow the negative binomial distribution. It is then straightforward to show that $E(c) = I(k^*)(1 - p)/p$ (see Casella and Berger (2002), p. 95-96).

Next, note that the random variable k follows a geometric distribution with

$$\Pr(k = i) = (1 - p)p^{i-1} \text{ for } i = 1, 2, \dots \quad (\text{A.2})$$

A simple calculation confirms that $E(k) = 1/(1 - p)$ (see Casella and Berger (2002), p. 97).

B. Proof of Proposition 4: First, recall that $\text{cov}(c,k) = E(ck) - E(c)E(k)$. By Proposition 5, we have $E(c)E(k) = I(k^*)/p$. To calculate $E(ck)$, we can use the fact that $E(ck) = E[E(c|k)k]$ where $E(c|k)$ denotes the conditional expectation of c given k . Next note that:

$$E(c|k) = \begin{cases} 1 + \left(I(k^*) + 1 - k \right) \frac{(1-p)}{p} & k \leq I(k^*) \\ 0 & k \geq I(k^*) + 1. \end{cases} \quad (\text{A.3})$$

Thus, for $k \leq I(k^*)$, we carry out Bernoulli (p) trials until there are $I(k^*) + 1 - k$ successes, so the expression for $E(c|k)$ follows from Proposition 5 (we must add one since we are effectively conditioning on $c \geq 1$); while for $k \geq I(k^*) + 1$, B will not initiate a conversation with C so that $E(c|k) = 0$. Note that $E(c|k)$ is decreasing in k . Calculating, we then obtain:

$$E[E(c|k)k] = \sum_{i=1}^{I(k^*)} \left[\left(1 + \left(I(k^*) + 1 - i \right) \frac{(1-p)}{p} \right) i \right] (1-p) p^{i-1} = \frac{I(k^*)}{p} - \frac{1-p^{I(k^*)}}{1-p}, \quad (\text{A.4})$$

so that $\text{cov}(c,k) = -(1-p^{I(k^*)})/(1-p) < 0$. After computing $\text{var}(c)$ and $\text{var}(k)$ (see Casella and Berger (2002), p. 95-97), and rearranging the resulting expression for

$\rho_{ck} = \text{cov}(c,k) / \sqrt{\text{var}(c) \text{var}(k)}$, we have:

$$\rho_{ck} = -\left(1 - p^{I(k^*)} \right) \sqrt{\frac{p}{I(k^*)(1-p)}} < 0. \quad (\text{A.5})$$

C. Comparative statics properties of $E(c)$ and $\Phi\{E(c)\}$: Let us begin with the comparative statics of k^* with respect to p and θ . Let

$$F(k, p, \beta, \theta) = \left[\frac{\theta(1-p) + p(1-\beta)}{1-p\beta} \right] \beta^k - \theta \quad (\text{A.6})$$

and note that the function $k^*(p, \beta, \theta)$ is defined implicitly by $F(k^*(p, \beta, \theta), p, \beta, \theta) = 0$.

Since $\partial F(k, p, \beta, \theta) / \partial k < 0$, applying the Implicit Function Theorem and noting that

$$\frac{\partial F(k^*, p, \beta, \theta)}{\partial p} = \frac{(1-\theta)(1-\beta)}{(1-p\beta)^2} \beta^{k^*} > 0, \text{ and} \quad (\text{A.7})$$

$$\frac{\partial F(k^*, p, \beta, \theta)}{\partial \theta} = \frac{(1-p)}{(1-p\beta)} \beta^{k^*} - 1 < 0, \quad (\text{A.8})$$

we see that $\partial k^* / \partial p > 0$ and $\partial k^* / \partial \theta < 0$. Since $I(k^*)$ is weakly increasing in k^* , it follows immediately from the above results that $E(c) = I(k^*)(1-p)/p$ is weakly decreasing in θ .

Finally, recall that the upper bound on $E[c]$ is given by $\Phi\{E[c]\} = k^*(1-p)/p$.

Differentiating with respect to p yields

$$\frac{\partial \Phi\{E[c]\}}{\partial p} = \frac{\partial k^*}{\partial p} \frac{(1-p)}{p} - \frac{k^*}{p^2}. \quad (\text{A.9})$$

Computing $\partial k^* / \partial p$ and noting that $(1-p)/(1-p\beta) < 1$, we have

$$\begin{aligned}
\frac{\partial k^*}{\partial p} \frac{(1-p)}{p} &= \left[-\frac{1}{\ln[\beta]} \frac{(1-\theta)(1-\beta)}{(1-p\beta)[\theta(1-p)+p(1-\beta)]} \right] \frac{(1-p)}{p} \\
&< -\frac{1}{p \ln[\beta]} \left[\frac{(1-\theta)(1-\beta)}{\theta(1-p)+p(1-\beta)} \right].
\end{aligned} \tag{A.10}$$

Finally since $-\ln[x] > 1-x$ for $0 < x < 1$, it follows that

$$\begin{aligned}
\frac{k^*}{p^2} &= \frac{1}{p^2} \frac{\log \left[\frac{(1-p\beta)\theta}{\theta(1-p)+p(1-\beta)} \right]}{\ln[\beta]} \\
&> -\frac{1}{p^2 \ln[\beta]} \left[1 - \frac{(1-p\beta)\theta}{\theta(1-p)+p(1-\beta)} \right] \\
&= -\frac{1}{p \ln[\beta]} \left[\frac{(1-\theta)(1-\beta)}{\theta(1-p)+p(1-\beta)} \right].
\end{aligned} \tag{A.11}$$

Combining these inequalities, it follows that $\partial \Phi\{E[c]\} / \partial p < 0$.


References

- Aghion, Philippe, and Peter Howitt, 1992, "A Model of Growth Through Creative Destruction," *Econometrica* 60, 323-351.
- Casella, George, and Roger L. Berger, 2002, *Statistical Inference*, Duxbury Press.
- Cohen, Lauren, Andrea Frazzini and Christopher Malloy, 2007, "The Small World of Investing: Board Connections and Mutual Fund Returns," working paper, Harvard Business School.
- Colla, Paola, and Antonio Mele, 2005, "Information Linkages and Correlated Trading," working paper, London School of Economics.
- Crawford, Vincent, and Joel Sobel, 1982, "Strategic Information Transmission," *Econometrica* 50, 1431-1451.
- Dewatripont, Mathias, and Jean Tirole, 2005, "Modes of Communication," *Journal of Political Economy* 113, 1217-1238.
- Ellison, Glenn D., and Drew Fudenberg, 1995, "Word of Mouth Communication and Social Learning," *Quarterly Journal of Economics* 110, 93-125.
- Fennell, Edward, 2005, "Keeping Out of the Shadows", posted July 5 to Times Online, <http://www.timesonline.co.uk>.
- Hellmann, Thomas, and Enrico Perotti, 2005, "The Circulation of Ideas: Firms versus Markets," working paper, University of British Columbia.
- Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein, 2005, "Thy Neighbor's Portfolio: Word-of-Mouth Effects in the Holdings and Trades of Money Managers," *Journal of Finance* 60, 2801-2824.
- Jackson, Matthew O., 2005, "The Economics of Social Networks," in Richard Blundell, Whitney Newey and Torsten Persson eds., *Proceedings of the 9th World Congress of the Econometric Society*, Cambridge University Press.
- Krishna, Vijay, and John Morgan, 2004, "The Art of Conversation: Eliciting Information From Experts Through Multi-Stage Communication," *Journal of Economic Theory* 117, 147-179.
- Ozsoylev, Han N., 2005, "Asset Pricing Implications of Social Networks," working paper, Said Business School, University of Oxford.

Saxenian, AnnaLee, 1994, *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*, Harvard University Press.

Shiller, Robert J., 2000, *Irrational Exuberance*, Princeton University Press.

Figure 1:
 $E(c)$ and $\Phi\{E(c)\}$ as Functions of p



Note: The above figure plots both $E(c) = I(k^*)(1 - p)/p$, as well as its upper bound, $\Phi\{E(c)\} = k^*(1 - p)/p$, as functions of p for $\beta = 0.95$ and $\theta = 0.1$.