

American Economic Association

Local Discouragement and Global Collapse: A Theory of Coordination Avalanches

Author(s): Thomas D. Jeitschko and Curtis R. Taylor

Reviewed work(s):

Source: *The American Economic Review*, Vol. 91, No. 1 (Mar., 2001), pp. 208-224

Published by: [American Economic Association](#)

Stable URL: <http://www.jstor.org/stable/2677905>

Accessed: 22/05/2012 15:00

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at

<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *The American Economic Review*.

<http://www.jstor.org>

Local Discouragement and Global Collapse: A Theory of Coordination Avalanches

By THOMAS D. JEITSCHKO AND CURTIS R. TAYLOR*

We study a dynamic game in which all players initially possess the same information and coordinate on a high level of activity. Eventually, players with a long string of bad experiences become inactive. This prospect can cause a coordination avalanche in which all activity in the population stops. Coordination avalanches are part of Pareto-efficient equilibria; they can occur at any point in the game; their occurrence does not depend on the true state of nature; and allowing players to exchange information may merely hasten their onset. We present applications to search markets, organizational meltdown, and inefficient computer upgrades. (JEL D83)

“We have nothing to fear, but fear itself.”
Franklin D. Roosevelt,
1st Inaugural Address,
Washington, DC, March 1933.

Economic activity often requires a significant degree of coordination among agents. In some settings, such as labor and product markets, prices aggregate the information held by individuals and automatically coordinate supply and demand behavior.¹ There are, however, many settings in which coordination is important, but the information held by individuals is *not* aggregated. In these settings, coordination among agents is only implicit, supported by nothing more than the belief by each individual that others will act appropriately. For instance, in the absence of deposit insurance, the only thing that keeps an individual from running on his bank is the belief that other depositors will refrain from running on it. Similarly, an impor-

tant reason that an individual purchases a fashion good is the belief that others will also find it fashionable.

Such faith-based coordination is, however, fragile. Specifically, even in states of nature under which coordinating on a certain action is optimal, some individuals may privately receive erroneous negative information that discourages them from acting appropriately. Moreover, this *local discouragement* can spill over and generate a *global collapse* of coordination. Indeed, merely the fear that some individuals have become discouraged can trigger a *coordination avalanche* in which all agents simultaneously abandon a possibly optimal course of action. That is, global collapse may arise from nothing more than “fear itself.”

In this paper, we present a model where this type of coordination failure may occur in a Pareto-efficient equilibrium. Specifically, we study a stochastic dynamic version of a Stag-Hunt game with many players.² In each period of the game, each player must decide between taking a risky action (investing) or a safe action (not investing). A player who does not invest receives a contemporaneous payoff of zero. If a player decides to invest, then he experiences a *success* with probability p and a *failure* with

* Jeitschko: Department of Economics, Texas A&M University, College Station, TX 77843 (e-mail: thosd@econ.tamu.edu); Taylor: Department of Economics, Duke University, P. O. Box 90097, Durham, NC 27708 (e-mail: crtaylor@econ.duke.edu). We thank Parikshit Ghosh, Farshid Vahid, Juuso Valimaki, and especially Steffen Huck for helpful comments. Jeitschko’s research was supported in part by the Bush Program in the Economics of Public Policy and Taylor’s in part by the Alfred P. Sloan Foundation and the National Science Foundation (Grant No. SBR-9810858).

¹ The processes by which prices aggregate information are only beginning to be studied. For pioneering work in this area, see Wolfgang Pesendorfer and Jeroen M. Swinkles (2000).

² The Stag-Hunt game exhibits the type of strategic complementarities necessary to yield multiple Pareto-rankable equilibria. On this point, see Russell Cooper and Andrew John (1988).

probability $1 - p$. These outcomes are independent across players and over time.

There are three key ingredients in our model. First, players are initially uncertain about the actual value of p . In particular, there are states of nature (high values of p) in which it is optimal to invest and other states (low values of p) in which investment is not optimal. Players possess a common prior under which the expected value of p is high enough to warrant initial investment.

The second crucial ingredient is the presence of a network externality or complementarity in actions. Specifically, while a failed investment always yields a negative payoff, the return to a successful investment depends positively on the number of other players also investing. Hence, the expected return to investing depends both on the state of nature and on the degree of coordination among players.

The final important component of the model is that information is not aggregated. In particular, agents observe only their own investment outcomes which they use to update their personal beliefs about p .

To see how these three ingredients can lead to a global collapse of coordination, suppose that it is individually rational for all agents to invest for the first τ periods of the game, and consider the set of agents who have experienced no successes by this date. For τ sufficiently large, these unsuccessful agents will become discouraged, i.e., they will believe that p is small and will rationally decide to stop investing. Moreover, the fact that these agents stop investing after date τ lowers the expected return to investing for everyone else in the population because of the network externality. Indeed, the fact that agents with zero successes stop investing after date τ can induce agents with one success to stop; the fact that agents with zero and one success stop investing can induce agents with two successes to stop; and so on. In the end, even the most optimistic agents who have experienced τ successes and no failures can be induced to stop investing by the fact that all of the less fortunate agents stop. In other words, a coordination avalanche can generate complete coordination failure at date $\tau + 1$ even when p is actually very high and investing is optimal.

Although our model is stylized, its three critical features—uncertainty, complementarities

in actions, and dispersed information—occur widely in economic and social contexts. Above we mentioned bank runs and purchase of fashion goods as settings in which the logic of our model has bite. Other potential applications include: initial public offers of stock, monetary crises, and industrial investment under imperfect competition, among others. Indeed, below we show how our model can be extended to explain the collapse of a search market, the sudden departure of several key members of an organization, and the mass adoption of a computer software upgrade. Each of these phenomena can be viewed as a coordination avalanche which occurs independently of the true state of nature and may, therefore, actually be quite inefficient.

Although coordination breaks down in our model because information is not aggregated, it should be emphasized that we are not studying a setting of social learning or herding.³ In other words, our notion of a coordination avalanche has little to do with the theory of informational cascades developed by Bikhchandani et al. (1992). An informational cascade occurs in a setting where agents initially possess differential information regarding the wisdom of taking a certain action. Agents move sequentially and observe the actions taken by all players ahead of them. At some point in the game, each remaining player may rationally decide to ignore his private information and take the same action as the player ahead of him in the queue; i.e., a cascade may occur. Moreover, because information is not being aggregated, it is possible that agents in a cascade will all take the wrong action.

In our model, by contrast, agents initially possess identical information. They do not observe the actions of other players, but update their beliefs based upon their personal experiences. At some point, a group of agents may become discouraged, and this possibility can generate a chain reaction throughout the population which leads all agents simultaneously to abandon a particular course of action.

A key feature of an informational cascade is its lack of depth. In other words, a cascade is

³ The herding literature began with Abhijit V. Banerjee (1992) and Sushil Bikhchandani et al. (1992).

easily stopped or reversed in the face of new information or expert players. Our notion of a coordination avalanche, on the other hand, involves wide-scale coordination failure which may be very difficult to stop or reverse. Hence, informational cascades are a good explanation for transitory spells of social conformity, while coordination avalanches are probably better at explaining more catastrophic and permanent changes.⁴

A paper that complements ours is Andrew Caplin and John Leahy (1994). Caplin and Leahy study a three-stage model of market dynamics. In the first stage (Business as Usual), market participants obtain private information regarding a parameter of common interest. Acting on the basis of this information is, however, costly, which leads to a spell of inertia. In the second stage (Market Crashes), some agents eventually become privately convinced that action is required. Their actions are observed and, hence, release information to other market participants which can lead them to suddenly follow suit. In the third stage (Wisdom After the Fact), agents lament the fact that the information released by the crash was known in the aggregate very early in the game (after the first period).

We are—not surprisingly—sympathetic to this story. We, however, do not believe that crashes are always precipitated by the observable actions of a set of agents. While herd-like stampedes of this kind are undoubtedly important, we believe that crashes are sometimes triggered by nothing more than the fear that some agents have become discouraged. Thus, while a crash occurs in Caplin and Leahy's model only if investment is inefficient, a coordination avalanche occurs in our setting whether or not it is economically justified. In fact, most crashes are probably precipitated by some combination of herding and fear. Caplin and Leahy (1994) identify the herd aspect of crashes, and in this paper we attempt to model the fear component.

⁴ In Ho Lee (1998) studies a financial market model in which the price of an asset may rise in a cascade until an expert arrives and sells it short. The price of the asset then drops precipitously in what Lee terms as an "informational avalanche," a concept that is clearly very different from ours. Amil Dasgupta (1999) investigates a hybrid model where herding occurs in coordination games.

Our paper is also related to a branch of the game-theoretic literature on common beliefs and almost common knowledge pioneered by Dov Monderer and Dov Samet (1989) and Ariel Rubinstein (1989). In particular, Hans Carlsson and Eric van Damme (1993) also study a model in which agents are uncertain about the payoffs in the game they are playing.⁵ These authors show that asymmetric or differential information and introspection will lead players to select the risk-dominant equilibrium over the Pareto-efficient one in a class of 2×2 coordination games. There are, however, several important differences between their analysis and ours. First, Carlsson and van Damme study a one-shot game in which players initially possess differential information which leads them to play secure strategies. By contrast, we study a dynamic game in which agents are initially identically informed and are presumed to play cooperative strategies in early periods. Also, whereas Carlsson and van Damme study a class of two-player games, we analyze "population games" in which the expected payoff from choosing the risky action depends positively on the fraction of other players also choosing it. Lastly, Carlsson and van Damme present an elegant contribution to the theory of equilibrium selection, while our orientation in this paper is more applied.

Finally, our analysis shares some common features with two papers in the macroeconomic theory of regime changes, namely Peter Howitt and R. Preston McAfee (1992) and Christophe Chamley (1999). Both of these papers also feature models built on Stag-Hunt games with multiple Pareto-rankable equilibria. In both Howitt and McAfee (1992) and Chamley (1999), the economy oscillates randomly between a high and a low equilibrium. In Howitt and McAfee (1992), agents learn to play sunspot equilibria because of spurious correlation between sunspots and economic fundamentals. This spurious correlation is similar to the feature in our model that even in good states of nature, there are always agents who will ultimately become discouraged. In Chamley (1999), transitions between the equilibria are triggered by an under-

⁵ See also Hyun Song Shin and Timothy Williamson (1996) and Stephen Morris and Shin (2001).

lying random process that governs agent heterogeneity.

In the following we investigate the causes and timing of avalanches. Specifically, we show that coordination avalanches will occur when network externalities are sufficiently strong or information differences between agents are sufficiently weak. We also show that an avalanche can happen at any point in the game. In Section IV we consider a setting in which agents acquire additional information by “talking with their friends,” and demonstrate that such information exchange can cause a coordination avalanche to occur sooner than it otherwise would have. We present two extensive applications of the theory in Section V. First, organizational meltdowns in which the most talented members of a productive group all simultaneously leave the organization for alternative employment. Second, we present a simple model of mass adoption of a computer upgrade. Some brief concluding remarks appear in Section VI. The proofs of all results either appear in Appendix A or have been omitted in the interest of space (and are available from the authors). In Appendix B we generalize the model somewhat and show that if a coordination avalanche occurs, then no agent ever experiments in equilibrium.

I. The Model

There is a continuum of *ex ante* identical agents with mass normalized to one. Each agent is a risk-neutral expected-utility maximizer who lives forever but (for analytic convenience) is assumed to possess a discount factor of $\delta = 0$.

In each period, $t = 1, 2, \dots$, each agent chooses an action from the set {YES, NO}. If an agent selects NO, then he receives a contemporaneous payoff of zero. If he selects YES, then his payoff depends on what other players in the population choose and on a move by nature. Specifically, if an agent plays YES, then he experiences a *success* with probability p and a *failure* with probability $1 - p$. An agent who experiences a success in period t receives a payoff of $u(m_t) - c$ and an agent who experiences a failure receives $-c$, where m_t is the mass of agents in the population who play YES in period t , $u(\cdot)$ is a strictly increasing function normalized so that $u(0) = 0$ and $u(1) = 1$, and $c \in (0, 1)$ is a constant. Hence, playing YES

involves an “investment” of c and an uncertain return. The fact that $u(\cdot)$ is increasing captures the network externality or complementarity that is an essential part of our story.

Another key element of the model is that agents are initially uncertain regarding the true value of p . In particular, they initially possess common prior beliefs over the value of p represented by the probability distribution $F(\cdot)$ which may have a density denoted by $f(\cdot)$. The expectation operator with respect to the prior is denoted by $E[\cdot]$. Agents (in the basic model) observe only their own histories of successes and failures which they use to update their beliefs on p at the end of each period.⁶ Finally, all aspects of the environment are assumed to be common knowledge.

To begin the analysis, consider a one-shot version of the game. Denote the expected value of p under the initial beliefs, $F(\cdot)$, by $\bar{p} \equiv E[p]$.

LEMMA 1 (ONE-SHOT EQUILIBRIUM): *If $\bar{p} > c$, then there are three Bayesian-Nash equilibrium outcomes to the one-shot game.*

1. A locally stable outcome in which all agents play YES and receive expected payoffs of $u(1)\bar{p} - c = \bar{p} - c$;
2. An unstable outcome in which a proportion $m = u^{-1}(c/\bar{p})$ of the agents play YES and all agents receive expected payoffs of zero;
3. A locally stable outcome in which none of the agents play YES and all agents receive payoffs of zero.

If $\bar{p} = c$, then there are two equilibrium outcomes corresponding to 2 and 3 above. And finally, if $\bar{p} < c$, the unique (globally stable) equilibrium corresponds to 3.

When $\bar{p} > c$, equilibrium outcome 1 obviously Pareto-dominates the others. Indeed, as the lemma makes clear, $\bar{p} > c$ is necessary and sufficient for the existence of multiple equilibria that are Pareto rankable. As this is fundamental to our analysis, we assume $\bar{p} > c$ below. Also, in games of longer duration, we assume that agents coordinate on a Pareto-efficient perfect

⁶ A variant of the model in which agents share information is considered in Section IV.

Bayesian equilibrium (PBE). This means that all agents play YES in the first period and continue playing YES so long as it is a mutual best response to do so. Finally, because we wish to analyze coordination avalanches that start when some agents ultimately become discouraged and quit playing YES, it is necessary to assume that $F(c - \varepsilon) > 0$, for some $\varepsilon > 0$. In other words, the prior must place nonzero weight on values of p for which investment is never efficient; i.e., agents must initially believe that such states are possible.

II. Coordination Avalanches

The expected return to an agent from playing YES at any point in the game depends on two things, namely, on the true value of p and on the mass of agents in the population playing YES at that point. As noted in the previous section, the mass of agents who play YES in the first period is $m_1 = 1$.

Consider an agent who experiences a failure in the first period. This agent will update his beliefs on p and be somewhat less optimistic about playing YES in the second period. Indeed, after a sufficiently long string of failures without a success, this agent will be so pessimistic about the value of p that he will stop playing YES (i.e., he will have a dominant strategy to play NO), and once he stops playing YES, he will never start again because an agent who plays NO receives no new information.

This is formalized in the following lemma.

LEMMA 2 (THE STOPPING DATE): *Define τ to be the smallest integer such that*

$$(1) \quad E[(p - c)(1 - p)^\tau] < 0.$$

Then $1 \leq \tau < \infty$, and in an efficient PBE all agents play YES through date τ , and any agent who experiences τ failures plays NO following that date.

All agents in the population *know* that the mass of $(1 - p)^\tau$ agents who have never had a success will quit investing after date τ . However, individuals with different experiences have different beliefs about how large the mass

of discouraged agents is likely to be. Specifically, individuals who have experienced relatively few successes over the τ periods believe that $(1 - p)^\tau$ is likely to be big, while individuals who have experienced relatively few failures believe it is likely to be small. Hence, calculating an optimal strategy for each possible history involves a rather subtle assessment about the likely experiences of other agents and how they will react to these experiences.

Let $B(k|t, p)$ be the binomial distribution with parameters t and p ; i.e., $B(k|t, p)$ is the probability of experiencing k or fewer successes in t periods when the probability of a success in each period is p . Suppose that all agents with k or fewer successes decide to stop playing YES following date τ . Then, the expected gross return to an agent who plays YES at date $\tau + 1$ is $pu(1 - B(k|\tau, p))$. In general, no agent knows the value of this return. Instead, each agent has beliefs about its value that are derived from his personal experiences. For instance, whether an agent with $k + 1$ successes should also quit playing YES depends on the expected value of $pu(1 - B(k|\tau, p))$ given $k + 1$ successes in τ trials. If this expected value is less than c , then agents with $k + 1$ successes should also quit playing YES, and the question then is whether agents with $k + 2$ successes should quit, given their beliefs.

Keeping track of agents' beliefs requires some additional notation. We denote the history of an agent who experiences k successes and $t - k$ failures by $h = g^k b^{t-k}$. (Note that if an agent stops playing YES at some point, then he stops acquiring information and his history stops evolving.) Define

$$r_{g^k b^{t-k}} \equiv pu(1 - B(k - 1|t, p)), \quad k = 0, \dots, t,$$

where $B(-1|t, p) = 0$. Then, $r_{g^k b^{t-k}}$ is the expected gross return to an agent who plays YES when the true success rate is p and all agents with fewer than k successes play NO. Also, define the expected value of r_h given history h by $\bar{r}_h \equiv E_h[r_h]$.⁷ In other words, \bar{r}_h is

⁷ It is straightforward to verify that

$$\bar{r}_{g^k b^{t-k}} = \frac{E[pu(1 - B(k - 1|t, p))p^k(1 - p)^{t-k}]}{E[p^k(1 - p)^{t-k}]}.$$

the expected gross return from playing YES to an agent with history h when all agents with worse histories play NO and all agents with the same or better histories play YES. Finally, define the expected value of p conditional on history h by $\bar{p}_h = E_h[p]$. The following example helps to fix this notation and illustrates how a coordination avalanche occurs.

*Example 1 (A Diamond Search Market).*⁸ Consider a setting in which agents may search for a new trading partner in each period $t = 1, 2, \dots$. If an agent searches in period t , then he finds a partner with probability pm_t , where p is interpreted as a macroeconomic friction parameter. When two agents meet, they transact a trade worth 1 to each of them [i.e., $u(m) = m$]. Suppose that the initial (common prior) beliefs about p are given by a Beta distribution with density $f(p) \equiv \phi(p|\alpha, \beta)$, where $\alpha = 1$ and $\beta = 2$, and suppose that the cost of search in each period is $c = 2/7$.⁹ The initial beliefs yield

$$\bar{p} = \frac{\alpha}{\alpha + \beta} = \frac{1}{3} > c.$$

So, in the efficient PBE, all agents search in the first period. The measure of agents who do not find a partner in the first period and receive payoffs of $-c$ is $B(0|1, p) = 1 - p$. The posterior belief of these agents is $f_b(p) = \phi(p|\alpha, \beta + 1)$, which yields

$$\bar{p}_b = \bar{r}_b = \frac{1}{4} < c.$$

Hence, these individuals have a dominant strategy to stop searching in the second period. How does this affect agents who found a trading partner in the first period? The posterior belief of these agents is $f_g(p) = \phi(p|\alpha + 1, \beta)$, which yields

$$\bar{p}_g = \frac{1}{2} > c.$$

When deciding on whether to search in the second period, however, these agents must account for the expected attrition of the discouraged agents. In particular, the probability of finding a second-period partner is

$$\begin{aligned} r_g &= p(1 - B(0|1, p)) \\ &= p(1 - (1 - p)) = p^2. \end{aligned}$$

Thus, agents who experienced a successful first-period search must compute the expectation of p^2 given their beliefs about p ; i.e., they must compute

$$\bar{r}_g = \int_0^1 p^2 f_g(p) dp = \frac{3}{10} < c.$$

Hence, agents who had good experiences in the first period are induced to stop searching in the second period by the expected attrition of the agents who had bad experiences. In the words of the following definition, a *coordination avalanche* occurs at date $\tau + 1 = 2$.

Definition 1 (Coordination Avalanche): If in a Pareto-efficient PBE of the game, all agents play YES through period τ and all agents play NO from period $\tau + 1$ on, then a coordination avalanche is said to occur in period $\tau + 1$.

PROPOSITION 1 (NECESSARY AND SUFFICIENT CONDITIONS): *A coordination avalanche occurs in period $\tau + 1$ of the game if and only if $\bar{p}_{b^{\tau-1}} \geq c > \bar{p}_{b^\tau}$ and $\bar{r}_{g^{k_b^{\tau-k}}} < c$ for $k = 1, \dots, \tau$.*

To understand the intuition underlying a coordination avalanche, suppose that $\bar{r}_{g^{j_b^{\tau-j}}} < c$ for all $j = 0, \dots, k$, so that agents with k or fewer successes quit playing YES in period $\tau + 1$. Whether agents with $k + 1$ successes will also quit depends on the magnitudes of two opposing forces, the *informational effect* and the *marginal network externality*. To isolate these forces, write

⁸ See Peter A. Diamond (1982).
⁹ Although our theory applies to any distribution for which $\bar{p} > c$ and $F(c - \epsilon) > 0$, Beta distributions are analytically very convenient because they are a conjugate family for the Bernoulli distributions. This greatly simplifies Bayesian updating (see Morris H. DeGroot, 1970 p. 160).

$$\begin{aligned}
(2) \quad & \bar{r}_{g^{k+1}b^{\tau-k-1}} - \bar{r}_{g^k b^{\tau-k}} \\
& = \{E_{g^{k+1}b^{\tau-k-1}}[pu(1 - B(k - 1|\tau, p))] - \\
& \quad - E_{g^k b^{\tau-k}}[pu(1 - B(k - 1|\tau, p))]\} + \\
& \quad + \{E_{g^{k+1}b^{\tau-k-1}}[p(u(1 - B(k|\tau, p)) - \\
& \quad - u(1 - B(k - 1|\tau, p)))]\}.
\end{aligned}$$

By assumption, $\bar{r}_{g^{k+1}b^{\tau-k-1}} < c$. Hence, if the difference on the left side of (2) is negative, then agents with $k + 1$ successes will join the agents with k or fewer successes in playing NO in period $\tau + 1$. Now, observe that the first term in braces on the right side of (2) is positive. This is the informational effect: Agents with $k + 1$ successes are more optimistic than those with k successes. Hence, they believe both that the probability of success, p , is likely to be larger and that the mass of quitters with $k - 1$ or fewer successes, $B(k - 1|\tau, p)$, is likely to be smaller. This effect can stop a coordination avalanche because it indicates that agents with $k + 1$ successes have a greater incentive to play YES than agents with k successes.

On the other hand, the second term in braces on the right side of (2) represents the marginal network externality which promotes a coordination avalanche. In particular, when deciding whether to quit playing YES, agents with k successes only have to worry about the attrition of agents with $k - 1$ or fewer successes. Agents with $k + 1$ successes, however, must also worry about the attrition of the agents with k successes. The marginal network externality given in (2) is what the agents with $k + 1$ successes perceive as the reduction in their expected payoff from playing YES due to the attrition of the agents with k successes. There is no reason a priori to suppose this effect is small. To the contrary, Bayesian agents will generally believe that there is a relatively large number of individuals with experiences similar to their own. Hence, while agents with τ successes believe that the mass of agents with no successes is likely to be small, agents with only one success believe that the mass of agents with no successes is likely to be large, and the perception of agents with only one success is critical in starting a coordination avalanche.

In the introduction, we suggested that a coordination avalanche can be interpreted as the manifestation of “fear itself.” In other words, a total collapse might occur even when the underlying fundamentals actually favor investment (i.e., even when $p > c$). In this regard, perhaps the most striking aspect of Proposition 1 is what it does not say. Specifically, the conditions under which a coordination avalanche occurs have nothing to do with the true value of p . That is, whether an avalanche ultimately occurs depends only on the shape of $u(\cdot)$, and the prior, $F(\cdot)$. In other words, if a coordination avalanche occurs in period $\tau + 1$, then it does so for any realization of p , even $p = 1$.

To illustrate, consider once again Example 1 in which a coordination avalanche occurs at $t = 2$. Note that agents who experienced a first-period success are not induced to quit searching due to the actual number, $1 - p$, of discouraged agents. No one observes this number. Rather, the decision of an agent with a first-period success to quit searching at $t = 2$ is based solely on his updated beliefs. The actual distribution of experiences in the population is not observable and, therefore, plays no part in an agent’s decision-making process. Hence, an agent with a first-period success will quit searching in the second period, even if the actual success rate is $p = 1$, and no one actually experienced a first-period failure.

The point here is not that a coordination avalanche is the inevitable consequence of initial beliefs. This is simply an artifact of the stylized model we have constructed to illustrate the phenomenon. The deterministic nature of coordination avalanches would undoubtedly be altered in a more complicated but less tractable model. What we do wish to emphasize is that coordination may collapse even in situations when the actual experiences of most agents do not justify quitting. In other words, it is the fear (not necessarily the fact) that some agents have become discouraged that generates a total collapse, and in this sense, coordination avalanches can be said to derive from nothing but fear itself.

III. The Causes and Timing of Avalanches

As noted above, whether the discouragement of the most unfortunate agents triggers an ava-

lanche depends only on the shape of $u(\cdot)$ and the prior. For instance, if $u(\cdot)$ is concave, then the marginal network externality generated by the first group of quitters is small and an avalanche is, therefore, unlikely. On the other hand, if $u(\cdot)$ is convex, then the first group of quitters imposes a relatively large marginal network externality on the rest of the agents in the population which can easily trigger an avalanche. This is formalized in the following proposition.

PROPOSITION 2 (LARGE EXTERNALITY): *If $u(\cdot)$ is sufficiently convex and $F(\cdot)$ is continuous at $p = 1$, then a coordination avalanche will occur.*

The intuition here is easily grasped. If $u(\cdot)$ is very convex, then successful investment requires a lot of teamwork (coordination). In this case, the discouragement of even a small number of agents has dire consequences for all the others. Hence, an avalanche will occur at date $\tau + 1$ unless every agent in the population is fairly certain that no one actually has become discouraged, i.e., $(1 - p)^\tau = 0$ or equivalently that $p = 1$. This, however, requires the prior to have a probability atom at $p = 1$; i.e., $F(\cdot)$ must be discontinuous at this point.

Proposition 2 shows how a large marginal network externality gives rise to a coordination avalanche for any initial beliefs except those that place significant probability on $p = 1$. The following result demonstrates that if the initial beliefs are sufficiently concentrated around $p = c$, then the informational effect will be small and local discouragement will lead to global collapse for any increasing function $u(\cdot)$.

PROPOSITION 3 (SMALL INFORMATIONAL EFFECT): *Suppose that the prior is given by*

$$F(p) \equiv \begin{cases} \lambda \hat{F}(p), & \text{if } p < c \\ 1 - \lambda(1 - \hat{F}(p)), & \text{if } p \geq c, \end{cases}$$

for $0 < \lambda \leq 1$, where $\hat{F}(\cdot)$ is any probability distribution satisfying $\bar{p} > c$ and $\hat{F}(c - \varepsilon) > 0$.

- (i) *The stopping date τ , defined in equation (1) does not depend on λ .*
- (ii) *If λ is sufficiently small, then a coordination avalanche occurs.*

This proposition provides some guidance about the type of initial beliefs that lead to a coordination avalanche. Specifically, an avalanche will occur at date $\tau + 1$ if the initial beliefs are sufficiently concentrated around $p = c$.¹⁰ To understand the intuition, recall that when the prior is concentrated, no single success or failure is very informative. (Indeed, in the limit, beliefs are degenerate and no updating occurs at all.) Thus, a concentrated prior means that the difference in beliefs between an agent with k and $k + 1$ successes is minor; i.e., the informational effect is small. To see why the prior must be concentrated around $p = c$, note that if it were very concentrated around $p < c$, then \bar{p} would be less than c , and no agent would play YES at the outset. On the other hand, if the prior were very concentrated around $p > c$, then it would take many periods before any agents became discouraged. When τ is large, the expected mass of agents with k successes and $\tau - k$ failures is small for any $k = 0, \dots, \tau$, which gives rise to a modest marginal network externality. This small externality may not necessarily dominate the small informational effect for agents who have had only few failures. Indeed, given the strong prior with $\bar{p} > c$, agents with only few failures will continue to think that investing is worthwhile. Conversely, when the prior is concentrated around $p = c$, even if τ is large, the marginal network externality always outweighs the informational effect, since even under the prior investments yield only small expected net returns. This leads to a total collapse of coordination of $\tau + 1$.

The preceding discussion should not be interpreted as saying that a necessary condition for the occurrence of an avalanche is that τ be small. This is not true. What is critical is the relative magnitudes of the informational effect and the marginal network externality. For instance, Proposition 2 indicates that an avalanche will occur very late in the game if $F(\cdot)$ is concentrated around some point $p > c$ and $u(\cdot)$ is very convex. It is possible, however, for an avalanche to occur very late in the game even in the important setting where $u(\cdot)$ is linear.

¹⁰ While it is technically convenient to concentrate the prior around $p = c$ by transferring probability mass to this point, the result holds more generally as the proof of the next proposition illustrates.

PROPOSITION 4 (AVALANCHES ANY TIME): *Suppose $u(m) = m$. Then, for any $\tau = 1, 2, \dots$ and for any $c \in (0, 1)$, there exists a prior, $F(\cdot)$, such that a coordination avalanche occurs in period $\tau + 1$.*

The intuition for this result is similar to that of Proposition 3. Specifically, if an avalanche is to occur at a very late date, then the marginal network externality associated with the attrition of any group will be quite small. This means that the informational effect must also be small; i.e., the distribution must be fairly concentrated.

Coordination avalanches that occur after many periods are especially dramatic. Agents play YES round after round until at some point they all simultaneously switch to NO in a rather spectacular collapse of cooperation. At this point in the game, agents with no successes quit investing based on their beliefs about p , and agents with $k + 1 > 0$ successes are pushed into quitting by their beliefs about agents with k successes much like a row of dominoes toppling or a chain reaction, until even the agents with only good experiences opt to play NO. This is illustrated in the following example.

Example 2: Suppose $u(m) = m$, the prior is a Beta distribution with $\alpha = 4$ and $\beta = 26$, and $c = 0.1$. In accordance with Proposition 4, the initial beliefs yield $\bar{p} = 0.1333$, which is close to c , and the variance is 0.0038, which is small. Moreover, it can easily be verified that a coordination avalanche occurs in this setting in the eleventh period. That is, through the tenth round, all agents play YES, and then switch to NO forever. At this point all agents who have never had a failure (and this is the entire population if $p = 1$) are more than 99.99 percent certain that $p > c$.

IV. Information Exchange

Coordination avalanches occur not because all agents decide that the actual success rate, p , is too low, but because they are unable to observe the population dynamics and fear that there may be significant attrition. Stated another way, avalanches occur because agents act on the basis of their personal experiences rather than aggregate information. Note that the true value of p is reflected in the experiences of the pop-

ulation as a whole after the first round of play. For example, suppose all agents reported their first-period outcomes to an information center which then publicized the success rate. In this case, information would be efficiently aggregated, and agents would continue to play YES from the second period on if and only if the true value of p exceeded c .

Even without an information center, agents may often have access to sources of information apart from their own experiences. For instance, they might directly observe the outcomes of a set of other agents or a group of individuals might share their experiences with each other. Because a coordination avalanche is generated by a failure to aggregate information, it seems reasonable that providing agents with additional unbiased signals would help deter an avalanche. In this section, we show that this is not necessarily so. Indeed, the situation may actually deteriorate when agents receive more information because this can cause an avalanche to occur at an earlier date than it otherwise would have.

We model the availability of additional information as follows. Suppose that at the end of each period each agent meets with the same group of $n - 1 \geq 1$ other agents (his friends). Suppose that each individual in a group truthfully reports his experience in each period to his friends.¹¹ Hence, so long as an agent continues to play YES, he receives n independent observations in each period (one from his own project outcome and $n - 1$ from his friends). For ease of reference, we refer to this setting as a game with communication and the original setting as a game without communication.

PROPOSITION 5 (AVALANCHES WHEN AGENTS COMMUNICATE): *If a coordination avalanche occurs in period $\tau + 1$ in the game without communication and if τ/n is an integer, then a coordination avalanche occurs in period $\tau/n + 1$ in the game with communication.*

Although this result is somewhat surprising, the intuition behind it is quite straightforward. When

¹¹ Of course, an agent is actually indifferent about what he reports to other group members.

agents talk with their friends, the acquisition of observations speeds up by a factor of n . Specifically, if all agents have played YES through period t , then the distribution of observations in the population at this point is $B(k|nt, p)$ for $k = 0, \dots, nt$. In other words, the mass of agents who have never observed a success (through personal experience or that of a friend) is $(1 - p)^{nt}$; the mass of agents who have observed one success is $ntp(1 - p)^{nt-1}$; and so on. Hence, if a coordination avalanche occurs in period $\tau + 1$ in the game without communication, then it occurs in period $\tau/n + 1$ in the game with communication, provided that τ/n is an integer.

As an illustration, recall Example 2 from the end of the previous section in which a coordination avalanche occurs in period $\tau + 1 = 11$ in a game without communication. Proposition 5 indicates that if agents communicate with $n - 1 = 1$ friend at the end of each period, then an avalanche occurs in period $\tau/n + 1 = 6$. If agents communicate with $n - 1 = 4$ friends, then an avalanche occurs in period $\tau/n + 1 = 3$. And, if agents communicate with $n - 1 = 9$ friends, then an avalanche occurs in period $\tau/n + 1 = 2$. Thus, while it is true that good news spreads more quickly in a game with communication, so does bad news, and it is the fear that bad news is spreading that can trigger an avalanche.

It is, of course, not true that communication always makes things worse. In particular, suppose that an avalanche occurs in period $\tau + 1$ of a game without communication, but that τ/n is not an integer (or sufficiently close to one). Then it may be possible in a game with communication to “jump” over the avalanche.

To illustrate, recall Example 1 from Section II in which an avalanche occurs in the second period in a game without communication. After the first period there are $(1 - p)$ agents who have experienced a failure and p agents who have experienced a success, and the expected attrition of the unsuccessful set is more than enough to induce the successful set also to quit. To see how communication can help, suppose that at the end of each period, agents share their experiences with one friend (so $\tau/n = 1/2$). In this case, there are three sets of agents after the first period. Specifically, there are $(1 - p)^2$ agents who have not observed a success, $2p(1 - p)$ agents who have observed one success and one failure, and p^2 agents who have

observed two successes. Agents in the first set will certainly quit at this point. The expected attrition of these agents is, however, not quite enough to induce agents in the second set to quit. In particular, it is straightforward to verify that $\bar{r}_{gb} = 2/7 = c$.¹² Thus, in this case, communication definitely improves matters because it stops a total collapse of coordination that would otherwise occur following the first round of play.

The fact that communication can stop an avalanche is not surprising. As noted above, coordination avalanches occur because information is not aggregated. In this light, it is somewhat more striking that giving agents additional information can actually make things worse by hastening the onset of a total collapse.

V. Applications and Extensions

In the introduction, we suggested that our theory is applicable in a wide variety of settings. Specifically, a coordination avalanche may occur whenever there is a network externality and information is not aggregated. In this section we outline two applications which are intended to highlight the breadth of the theory.

A. Organizational Meltdown

It is common in many organizations where complementarities are important (e.g., academic departments, theater and music groups, and sports teams) to observe several of the most talented individuals simultaneously abandoning the organization for alternative employment. Often, the new positions taken by these individuals are no better than the ones they left. Indeed, the primary reason for leaving may not have been the attraction of the new position, but the fear that other talented team members were about to quit and damage the performance of the organization. In the end, of course, this is exactly what happens. Our theory of coordination avalanches can easily be modified to explain the sudden departure of several key members of an organization as we now illustrate.

¹² Note that indifferent agents who have experienced at least one failure will not play NO in a Pareto-efficient PBE because agents with fewer failures are strictly better off if they play YES.

Suppose that an organization is initially composed of $n_1 > 1$ productive agents each of whom is one of two possible types, $\theta_i = 1$ (talented) or $\theta_i = 0$ (untalented) for $i = 1, \dots, n_1$. An individual's type is private unverifiable information. In each period $t = 1, 2, \dots$ each agent in the organization undertakes a project that either succeeds, yielding him a payoff of one, or fails, yielding him a payoff of zero (for simplicity, the outcome of each project is also assumed to be private unverifiable information). The probability that agent i 's project succeeds in period t is $\gamma(\theta_i)p_t$, where

$$\gamma(\theta) \equiv \theta + (1 - \theta)\eta$$

for some small $\eta > 0$ and where

$$p_t \equiv \frac{\sum_{i=1}^{n_t} \theta_i}{n_t}$$

is the average level of talent in the organization at date t . Suppose that at the beginning of each period each agent in the organization has the option of quitting and receiving an expected payoff flow of $\theta_i c$, where $c \in (0, 1)$. Let $F(\cdot)$ denote the common prior over p_1 and suppose that $\bar{p}_1 > c$ and $F(c - \varepsilon) > 0$.

First, observe that the expected date- t payoff to an untalented agent from staying in the organization is $\eta E_t[p_t | \theta = 0]$ while his payoff from quitting is zero. Hence, an untalented agent will leave the organization only if he is certain that there are no talented members left.

What about talented members? The expected payoff to a talented agent in the first period if all other talented agents stay in the organization is \bar{p}_1 and his opportunity cost is c . Hence, in an efficient PBE, all agents remain in the organization for the first period, and $p_2 = p_1$. Indeed, by analogy with Lemma 2, $p_t = p_1$ for all $t \leq \tau$, where τ is the smallest integer satisfying

$$E[(p_1 - c)(1 - p_1)^\tau | \theta = 1] < 0.$$

Following period τ , however, a talented agent who has experienced no successes will believe that the average level of talent in the organization is so low that he should quit. This, of

course, can trigger a coordination avalanche in which all talented individuals "jump ship" in period $\tau + 1$. Note also that such an avalanche may be based on fear rather than fact. Specifically, the organization may actually contain a very high fraction of talented agents none of whom was actually discouraged, but each of whom was worried about the possible discouragement of his colleagues. Indeed, it is possible that all individuals in a small organization might actually be talented, in which case a coordination avalanche results in the complete dissolution of the group.

B. Avalanches in Computer Upgrades

Another area in which network externalities are important is the purchase and use of computer hardware and software. Indeed, a major reason computer users upgrade their computers is to stay compatible with friends and co-workers who are also upgrading. In fact, the fear that a significant number of other computer users have switched (or are about to switch) to a new standard can actually induce a sudden wide-scale shift to the new program even when switching is not efficient. We illustrate this possibility with a simple modification of our model.

Suppose that there is a continuum of computer users with total mass of one. There are two types of individuals in this population, type L computer users for whom $\theta = \theta_L \geq 0$ and type H computer users for whom $\theta = \theta_H > \theta_L$. The proportion of type L individuals in the population is denoted by p . There is initial uncertainty regarding the actual value of p represented by the common prior $F(\cdot)$.

The benefits to an individual from using a software program derive from two sources. Specifically, in each period $t = 1, 2, \dots$ each agent i receives a stand-alone benefit from using program j of $\rho_j \theta_i$ where $\rho_j > 0$ is a productivity parameter. Also, in each period, each agent is uniformly randomly matched with another agent in the population to work on a joint project. If two matched agents use the same software program, then they each receive a payoff of one, and if they use different programs, then they each receive zero.

All agents initially own the original version of the program denoted by $j = 0$. At $t = 1$, an upgrade denoted by $j = 1$ comes on the market,

where $\rho_1 > \rho_0$. The total cost to any individual of switching to the new program is $\kappa > 0$. The benefits from switching to the new program, of course, depend in part on the number of other individuals who switch.

The following key inequalities are assumed to hold:

$$(3) \quad \theta_H > \frac{\kappa + 1}{\rho_1 - \rho_0},$$

$$(4) \quad \theta_L < \frac{\kappa - (1 - \bar{p})}{\rho_1 - \rho_0},$$

$$(5) \quad \theta_L > \frac{\kappa - 1}{\rho_1 - \rho_0}.$$

Inequality (3) says that the stand-alone benefit to a type H computer user from adopting the new program is so high that he should upgrade even if no other individuals do so (i.e., upgrading at $t = 1$ is a dominant strategy). This is not true for type L individuals. Indeed, (4) indicates that (given initial beliefs), type L individuals are better off jointly staying with the old program rather than jointly switching to the new one. Inequality (5), however, implies that if a type L individual believes that a large enough fraction of computer users have switched to the new program, then he should follow suit.

The potential for a coordination avalanche in this setting is clear. In particular, (3) and (4) imply that there is a Pareto-efficient PBE in which all type H and no type L computer users switch to the new program at the beginning of $t = 1$. Hence, in the first period, $p(1 - p)$ type L computer users will suffer compatibility problems from being matched with type H individuals. Accordingly, these type L computer users will update their beliefs on p and be more inclined to switch to the new program at $t = 2$.¹³ Specifically, in the context of this example, define

$$c \equiv \frac{1}{2}((\rho_1 - \rho_0)\theta_L + 1 - \kappa),$$

¹³ Type H computer users also update their beliefs on p , but (3) ensures that they never wish to switch back to the old program.

and let τ be defined as in (1). Then, a type L computer user who has been matched to do projects with τ type H and no type L individuals will switch to the new program in period $\tau + 1$. This can, of course, trigger an upgrade avalanche in which all type L individuals switch. Indeed, Proposition 3 indicates that if $F(\cdot)$ is sufficiently concentrated around c , then an avalanche is inevitable.

Observe once again that a coordination avalanche may be based on fear rather than fact. In particular, the number, $1 - p$, of type H individuals who value the new program highly may actually be quite small (even zero), and an avalanche may, nevertheless, occur. The net change in social surplus when all individuals switch to the new program is

$$(\rho_1 - \rho_0)(p\theta_L + (1 - p)\theta_H) - \kappa.$$

Inequality (3) implies that this is positive if p is sufficiently small, but inequality (4) implies that it is negative if p is sufficiently large. Hence, a coordination avalanche may induce all agents to switch to the new program even though upgrading is socially inefficient and the number of agents who benefit from upgrading is minute.

VI. Conclusion

The theory of coordination avalanches presented in this paper is novel in several respects. Specifically, we have shown that it is possible for coordination among agents to break down endogenously at any point in time. Moreover, this collapse does not depend on agents' arbitrary beliefs about sunspots or on social herding. Rather, it derives from uncertainty about payoffs, complementarities in actions, and dispersed information. We believe that these three characteristics are common in many economic and social settings.

We presented brief applications of the theory to search markets, organizational meltdowns, and the wide-scale adoption of a computer upgrade. We also provided some technical conditions under which a coordination avalanche would occur. Intuitively, coordination avalanches are likely when there are strong network externalities or weak informational differences among agents with different experiences.

An intriguing feature of a coordination

avalanche is that its occurrence does not depend on the true state of nature. In other words, an avalanche occurs because agents fear that some of their fellow players have become discouraged. Hence, coordination avalanches are a product of agents' beliefs rather than observable actions. As we noted above, we do not think that real-world collapses are typically based only on beliefs, but we do think that they are often an important

part of the story. In this light, our analysis points to two policy prescriptions. First, any mechanism that raises the confidence of individuals will help to avert a coordination avalanche. Second, credibly publicizing aggregate information can also help. A coordination avalanche is triggered by fear and uncertainty, and anything that eases fear or reduces uncertainty can, therefore, keep coordination from collapsing.

APPENDIX A

PROOF OF LEMMA 2 (THE STOPPING DATE):

Let \hat{t} be the date after which an agent with \hat{t} failures and no successes will stop playing YES in a Pareto-efficient PBE. Since $u(1) = 1$, \hat{t} is the smallest integer such that the expected value of $1 \cdot p$ given \hat{t} failures and no successes is less than c :

$$\frac{\int_0^1 p(1-p)^{\hat{t}} dF(p)}{\int_0^1 (1-p)^{\hat{t}} dF(p)} < c.$$

Multiplying both sides by the denominator of the left side and consolidating yields

$$E[(p-c)(1-p)^{\hat{t}}] < 0.$$

Hence, $\hat{t} = \tau$. Next, note that $\tau = 0$ implies $Ep < c$ violating the assumption that $\bar{p} > c$. Hence, it remains only to be shown that $\tau < \infty$ exists. This follows easily from the assumption that $F(c - \varepsilon) > 0$ for some $\varepsilon > 0$. The details are omitted in the interest of space.

PROOF OF PROPOSITION 2 (LARGE EXTERNALITY):

Since $u(0) = 0$, $u(1) = 1$, and $u(\cdot)$ is increasing, the limiting case of the convex function $u(\cdot)$ is

$$u(m) = \begin{cases} 0, & \text{if } 0 \leq m < 1 \\ 1, & \text{if } m = 1. \end{cases}$$

By Lemma 2, agents with no successes will quit playing YES in period $\tau + 1$. At this point the payoff from playing YES to an agent who has had only good experiences is

$$E_{g^\tau}[u(m_{\tau+1})p|m_{\tau+1} < 1] + \lim_{x \rightarrow 1} (u(1)1[F_{g^\tau}(1) - F_{g^\tau}(x)]) - c.$$

If $F(\cdot)$ is continuous at 1, then so are the posterior beliefs, so that the second term is equal to zero. For $u(\cdot)$ sufficiently convex, the first term is arbitrarily close to zero, so that the total expected payoff from playing YES is negative and a coordination avalanche occurs in $\tau + 1$.

PROOF OF PROPOSITION 3 (SMALL INFORMATION EFFECT):

Under the initial beliefs, the first agents will become discouraged following date $\hat{\tau}$ which is the smallest integer satisfying

$$\frac{\lambda \int_0^1 p(1-p)^{\hat{\tau}} d\hat{F}(p) + (1-\lambda)c(1-c)^{\hat{\tau}}}{\lambda \int_0^1 (1-p)^{\hat{\tau}} d\hat{F}(p) + (1-\lambda)(1-c)^{\hat{\tau}}} < c.$$

Multiplying through by the denominator on the left and collecting terms renders this as

$$\lambda \int_0^1 (p-c)(1-p)^{\hat{\tau}} d\hat{F}(p) < 0.$$

Comparing this with (1) confirms that $\hat{\tau} = \tau$ which does not depend on λ . Now, let $\hat{E}[\cdot]$ denote the expectation operator with respect to $\hat{F}(\cdot)$, and recall (2). The (positive) informational effect which retards an avalanche can be written as

$$\begin{aligned} & \left(\frac{\lambda \hat{E}[pu(1-B(k-1|\tau, p))p^{k+1}(1-p)^{\tau-k-1}]}{\lambda \hat{E}[p^{k+1}(1-p)^{\tau-k-1}] + (1-\lambda)c^{k+1}(1-c)^{\tau-k-1}} + \right. \\ & \quad \left. + \frac{(1-\lambda)cu(1-B(k-1|\tau, c))c^{k+1}(1-c)^{\tau-k-1}}{\lambda \hat{E}[p^{k+1}(1-p)^{\tau-k-1}] + (1-\lambda)c^{k+1}(1-c)^{\tau-k-1}} \right) - \\ & - \left(\frac{\lambda \hat{E}[pu(1-B(k-1|\tau, p))p^k(1-p)^{\tau-k}]}{\lambda \hat{E}[p^k(1-p)^{\tau-k}] + (1-\lambda)c^k(1-c)^{\tau-k}} + \right. \\ & \quad \left. + \frac{(1-\lambda)cu(1-B(k-1|\tau, c))c^k(1-c)^{\tau-k}}{\lambda \hat{E}[p^k(1-p)^{\tau-k}] + (1-\lambda)c^k(1-c)^{\tau-k}} \right), \end{aligned}$$

for $k = 0, \dots, \tau - 1$. It is easy to check that as λ approaches zero, this difference becomes arbitrarily small. On the other hand, as λ approaches zero, the (negative) marginal network externality converges to $c(u(1-B(k-1|\tau, c)) - u(1-B(k|\tau, c)))$, which is strictly negative because $c > 0$, $\tau < \infty$, and $u(\cdot)$ is increasing.

PROOF OF PROPOSITION 4 (AVALANCHES ANY TIME):

Let τ and c be given. Then by Proposition 1, it must be shown that there exists a common prior such that $\bar{p}_{b^{\tau-1}} > c > \bar{p}_{b^\tau}$, and $\bar{r}_{g^k b^{\tau-k}} < c, \forall k = 1, \dots, \tau$.

Suppose that the common prior is a Beta distribution with parameters α and β ; i.e., $f(p) = \phi(p|\alpha, \beta)$. Then we seek $\alpha > 0$ and $\beta > 0$ such that a coordination avalanche occurs in period $\tau + 1$.

Assume $\alpha > [c/(1-c)](\tau - \varepsilon) > 0$ for arbitrary $\varepsilon \in (0, 1)$ and set $\beta = [(1-c)/c]\alpha - \tau + \varepsilon$. Then $\alpha > 0$ and $\beta > 0$. Moreover, period $\tau + 1$ is the first period in which any agents play NO. That is,

$$\bar{p}_{b^{\tau-1}} = \frac{\alpha}{\alpha + \beta + \tau - 1} = \frac{c\alpha}{\alpha - c(1 - \varepsilon)} > c,$$

and

$$\bar{p}_{b^\tau} = \frac{\alpha}{\alpha + \beta + \tau} = \frac{c\alpha}{\alpha + c\varepsilon} < c.$$

It now must be shown that there exists $\alpha > [c/(1-c)](\tau - \varepsilon)$ and $\beta = [(1-c)/c]\alpha - \tau + \varepsilon$, such that agents with $k = 1, \dots, \tau$ good experiences also play NO.

Notice that for all $k = 1, \dots, \tau$,

$$\begin{aligned} \bar{r}_{g^k b^{\tau-k}} &= E_{g^k b^{\tau-k}}[p(1 - B(k - 1|\tau, p))] \\ &= \int_0^1 \left[p - \sum_{j=0}^{k-1} \binom{\tau}{j} p^{j-1} (1-p)^{\tau-j} \right] f_{g^k b^{\tau-k}}(p) dp. \end{aligned}$$

An agent with k good experiences after playing YES for τ periods has a posterior about p given by $f_{g^k b^{\tau-k}}(p) = \phi(p|\alpha + k, \beta + \tau - k)$. Thus, $\bar{r}_{g^k b^{\tau-k}}$ becomes

$$\begin{aligned} &\frac{\Gamma(\alpha + \beta + \tau)\Gamma(\alpha + k + 1)}{\Gamma(\alpha + k)\Gamma(\alpha + \beta + \tau + 1)} - \\ &- \sum_{j=0}^{k-1} \binom{\tau}{j} \frac{\Gamma(\alpha + \beta + \tau)\Gamma(\alpha + k + j + 1)\Gamma(\beta + 2\tau - k - j)}{\Gamma(\alpha + k)\Gamma(\beta + \tau - k)\Gamma(\alpha + \beta + 2\tau + 1)}. \end{aligned}$$

Substituting $\beta = [(1 - c)/c]\alpha - \tau + \varepsilon$ from above, and recalling that $\Gamma(z + 1) = z\Gamma(z)$, this simplifies to

$$\frac{\alpha + k}{\frac{1}{c}\alpha + \varepsilon} - \sum_{j=0}^{k-1} \binom{\tau}{j} \frac{\prod_{u=0}^j (\alpha + k + u) \prod_{v=0}^{\tau-j-1} \left(\frac{1-c}{c}\alpha + \varepsilon - k + v\right)}{\prod_{w=0}^{\tau} \left(\frac{1}{c}\alpha + \varepsilon + 1 + w\right)}.$$

Thus, for all k , $\bar{r}_{g^k b^{\tau-k}}$ is the sum of ratios of polynomials and hence continuous in α . Since

$$\begin{aligned} \lim_{\alpha \rightarrow \infty} \bar{r}_{g^k b^{\tau-k}} &= \frac{\alpha}{\frac{1}{c}\alpha} - \sum_{j=0}^{k-1} \binom{\tau}{j} \frac{\alpha^{j+1} \left(\frac{1-c}{c}\alpha\right)^{\tau-j}}{\left(\frac{1}{c}\alpha\right)^{\tau+1}} \\ &= c - \sum_{j=0}^{k-1} \binom{\tau}{j} (1-c)^{\tau-j} c^{j+1}, \end{aligned}$$

for sufficiently large α , $\bar{r}_{g^k b^{\tau-k}} < c, \forall k = 1, \dots, \tau$.

PROOF OF PROPOSITION 5 (AVALANCHES WHEN AGENTS COMMUNICATE):

Since an avalanche occurs in period $\tau + 1$ of the game without communication:

(A1) $\bar{p}_{b^t} > c > \bar{p}_{b^\tau}, \quad \forall t < \tau$

and

(A2) $\bar{r}_{g^k b^{\tau-k}} < c, \quad k = 1, \dots, \tau.$

Now, note that as long as all agents play YES in the game with communication, the distribution of observations in the population evolves over time according to $\{B(k|nt, p)\}_{t=1}^{\infty}$. Similarly, as long as all agents play YES in the game without communication, the distribution of experiences in the population evolves over time according to $\{B(k|t, p)\}_{t=1}^{\infty}$. The first sequence is thus a subsequence of the second.

Consider a period $j < \tau/n$ in the game with communication. If all agents play YES through period j , then the beliefs of agents who have never observed a success (their own or a friend's) yield \bar{p}_{b^j} . This is greater than c by (A1). Hence, all agents play YES through period $j < \tau/n$ in a Pareto-efficient equilibrium. However, in period $j = \tau/n$, (A1) indicates that the most pessimistic agents will quit playing YES. Moreover, since $B(k|nj, p) = B(k|\tau, p)$, an avalanche occurs in the game with communication at this point by (A2).

APPENDIX B

Suppose that agents do not fully discount the future and let $\delta > 0$ be the discount factor. Also let V_h be the expected discounted payoff to an agent with history h in a Pareto-efficient PBE. In this context, experimentation means that it is optimal to play YES for an agent with history h even when $\bar{r}_h - c < 0$, which obviously cannot be the case if $\delta = 0$. The following result shows that the analysis presented in the main text generalizes to a setting with nonmyopic agents so long as δ is not too large.

PROPOSITION B1 (AVALANCHES AND NONMYOPIC AGENTS): *If a coordination avalanche occurs in period $\tau + 1$ of the game when $\delta = 0$, then there exists $\bar{\delta} > 0$ such that a coordination avalanche occurs in period $\tau + 1$ of the game when $0 < \delta < \bar{\delta}$.*

The intuition behind this result is fairly obvious. Specifically, because an avalanche occurs in the first period when all agents would receive a myopic expected payoff strictly less than zero from playing YES (i.e., when $\bar{r}_{g^k b^{t-k}} < c$ for $k = 0, \dots, t$), an avalanche will still occur at this point if agents are not too patient. The next result shows that avalanches and experimentation are incompatible.

PROPOSITION B2 (AVALANCHES AND EXPERIMENTATION): *Agents experiment in an efficient PBE only if a coordination avalanche does not occur.*

PROOF:

We prove the converse. Suppose an avalanche occurs in period $\tau + 1$. Then we want to show that no agent ever plays YES when his expected current payoff is negative. First, because an avalanche occurs in period $\tau + 1$, $V_{g^k b^{\tau-k}} = 0$ for $k = 0, \dots, \tau$. Move back one period and write the Bellman equation for the agents who have had no successes

$$V_{b^{\tau-1}} = \max \{0, \bar{p}_{b^{\tau-1}} - c + \delta \times 0\}.$$

Since an avalanche occurs in period $\tau + 1$, the most pessimistic agents still play YES in period τ , so $\bar{p}_{b^{\tau-1}} - c \geq 0$. Since $\bar{p}_{h^{\tau-1}} \geq \bar{p}_{b^{\tau-1}}$, the result follows.

The intuition here is a bit more subtle. First, note that an avalanche is always triggered by the expected attrition of agents who have never had a success. Now, consider an agent who has experienced only failures one period before the avalanche occurs. It is useless for him to experiment at this juncture because the game will collapse in the next period whether or not he has a success. Hence, he must earn a nonnegative expected current payoff from playing YES at this point. But, no agent in the current period or any previous one has ever had lower expectations, so no one ever experiments.

The message of Proposition B2 is clear. Agents will not experiment in the face of a coordination avalanche. Moreover, while myopic agents are probably the most likely to fall victim to an avalanche, Proposition B1 indicates that avalanches can knock out patient players as well.

REFERENCES

- Banerjee, Abhijit V.** "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, August 1992, 107(3), pp. 797–817.
- Bikhchandani, Sushil; Hirschleifer, David and Welch, Ivo.** "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy*, October 1992, 100(5), pp. 992–1026.
- Caplin, Andrew and Leahy, John.** "Business as Usual, Market Crashes, and Wisdom after the Fact." *American Economic Review*, June 1994, 84(3), pp. 548–65.
- Carlsson, Hans and van Damme, Eric.** "Global Games and Equilibrium Selection." *Econometrica*, September 1993, 61(5), pp. 989–1018.
- Chamley, Christophe.** "Coordinating Regime Switches." *Quarterly Journal of Economics*, August 1999, 114(3), pp. 869–905.
- Cooper, Russell and John, Andrew.** "Coordinating Coordination Failures in Keynesian Models." *Quarterly Journal of Economics*, August 1988, 103(3), pp. 441–63.
- Dasgupta, Amil.** "Social Learning with Payoff Complementarities." Mimeo, Yale University, November 1999.
- DeGroot, Morris H.** *Optimal statistical decisions*. New York: McGraw-Hill, 1970.
- Diamond, Peter A.** "Aggregate-Demand Management in Search Equilibrium." *Journal of Political Economy*, October 1982, 90(5), pp. 881–94.
- Howitt, Peter and McAfee, R. Preston.** "Animal Spirits." *American Economic Review*, June 1992, 82(3), pp. 493–507.
- Lee, In Ho.** "Market Crashes and Informational Avalanches." *Review of Economic Studies*, October 1998, 65(4), pp. 741–59.
- Monderer, Dov and Samet, Dov.** "Approximating Common Knowledge with Common Beliefs." *Games and Economic Behavior*, June 1989, 1(2), pp. 170–90.
- Morris, Stephen and Shin, Hyun Song.** "Private versus Public Information in Coordination Problems," in Mathias Dewatripont, Lars Peter Hansen, and Stephen Turnovsky, eds., *Advances in economic theory and econometrics: Proceedings of the 8th World Congress of the Econometric Society*, 2001 (forthcoming).
- Pesendorfer, Wolfgang and Swinkels, Jeroen M.** "Efficiency and Information Aggregation in Auctions." *American Economic Review*, June 2000, 90(3), pp. 499–525.
- Rubinstein, Ariel.** "The Electronic Mail Game: Strategic Behavior under 'Almost Common Knowledge'." *American Economic Review*, June 1989, 79(3), pp. 385–91.
- Shin, Hyun Song and Williamson, Timothy.** "How Much Common Belief Is Necessary for a Convention?" *Games and Economic Behavior*, April 1996, 13(2), pp. 252–68.