Influential Opinion Leaders*

Antoine Loeper† Jakub Steiner‡ Colin Stewart§

Universidad Carlos III de Madrid Northwestern University University of Toronto

October 31, 2012

Abstract

We present a two-stage coordination game in which early choices of experts with special interests are observed by followers who move in the second stage. We show that the equilibrium outcome is biased toward the experts' interests even though followers know the distribution of expert interests and account for it when evaluating observed experts' actions. Expert influence is fully decentralized in the sense that each individual expert has a negligible impact. The bias in favor of experts results from a social learning effect that is multiplied through a coordination motive. We show that the total effect can be large even if the direct social learning effect is small. We apply our results to the diffusion of products with network externalities and the onset of social movements.

1 Introduction

When a large group of agents seek to coordinate their behavior in an uncertain environment, it is common for individuals to look to better informed experts for guidance. The preferences of these experts may not coincide with those of the agents who observe their choices. In light of this conflict, do the experts' preferences influence mass opinion and behavior? We show that the choices of expert early movers can have a large effect on outcomes, biasing the results toward their

*We thank David Austen-Smith, Mehmet Ekmekci, Yosh Halberstam, Stephen Morris, Alessandro Pavan, Carolyn Pitchik, and Matt Turner for helpful comments. Loeper acknowledges the financial support from grant ECO2010-19596 from the Spanish "Ministerio de Ciencia e Innovación". Stewart is grateful to SSHRC for financial support of this research.

†email: aloeper@eco.uc3m.es

‡email: j-steiner@kellogg.northwestern.edu

§email: colin.stewart@utoronto.ca

own preferences. The effect arises even though our model features Bayesian decision-makers who know the distribution of experts' biases, and each expert can influence only a negligible share of the population.

One setting to which our model naturally applies is the diffusion of products with positive network externalities. Since Katz and Lazarsfeld (1955), the empirical marketing literature has shown that friends and relatives with superior knowledge about a product are often viewed as the most reliable source of information by prospective consumers, and that the diffusion process is overwhelmingly driven by these well informed, visible individuals (e.g., Weimann 1991). Consistent with this literature, our results show that if each potential adopter observes the early choices of a few experts, then the equilibrium coordination outcome disproportionately reflects the experts' preferences. Moreover, this result holds even if experts are known to face different relative prices. In particular, offering a low price to a small group of early adopters can lead to widespread inefficient adoption even if later buyers know the past prices and all market participants have good information about the quality of the good. In line with the marketing literature, opinion leaders become natural marketing targets.

Social movements offer another important example of a large-scale coordination problem in which well informed individuals can play an influential role within their social network. Empirically, knowing an activist involved in a social movement is one of the main determinants of mobilization (e.g., McAdam and Paulsen 1993, or Opp and Gern 1993). In the context of a social movement, our model can explain how a small vanguard of protesters who are well informed about the status quo regime can spark a massive popular uprising, even if the participants in early protests are not representative of the population, and the population knows their preference distribution.

Systematic manipulation of decision-makers' actions by experts' interests may appear to be at odds with rational choice. A Bayesian decision-maker accounts for experts' biases when evaluating their advice, potentially offsetting the experts' influence. For instance, in the cheap talk literature the bias of the informed agent typically results in a limitation on credible communication rather than consistent manipulation of the principal. In contrast, expert influence can arise naturally in models of social learning in which a follower observes choices made by experts whose preferences may differ from her own. We show how a coordination motive may multiply this social learning effect. Moreover, the multiplication can be so large as to create a sizeable total effect even when the direct social learning effect is vanishingly small. As a result, the determinants of the experts' influence in this model are qualitatively different from a standard social learning model.

We identify a novel channel through which experts' biases, despite being known, influence the coordination outcome. Agents must choose one of two actions, A or B. The coordination outcome of the game is A if A is chosen by sufficiently many agents, and is B otherwise. Each agent is assumed to have an intrinsic bias for a particular action, which she trades off against her preference for choosing the action that becomes the coordination outcome. The set of agents is composed of a continuum of experts and followers, each of whom possesses private information about the fraction of choices that is needed for A to become the coordination outcome instead of B. In the first stage of the game, experts simultaneously choose their actions. In the second stage, followers observe a private sample of experts' actions and simultaneously choose their actions.

As in the global games literature, private information ensures equilibrium uniqueness and gives rise to strategic uncertainty. In equilibrium, the coordination outcome is determined by a combination of both the experts' and the followers' biases, enabling us to quantify the influence of experts' preferences. Since experts' beliefs are based on their private information, their actions provide information to followers about the future coordination outcome, thereby influencing those followers. The model exhibits a subtle interplay between the beliefs and actions of experts and those of followers. Each expert, unable to affect the outcome on her own, treats the outcome as given (albeit uncertain). Each follower also treats the outcome as given, and accounts for the distribution of experts' biases when forming beliefs based on observed actions. We show that the interplay between experts and followers skews the outcome toward the average bias of experts. The experts become opinion leaders.

Even though followers are aware of the biases of the experts, their beliefs about the coordination outcome become skewed in the direction of the experts' biases, at least sometimes. While the beliefs of rational followers cannot be manipulated systematically (in the sense that they are correct when averaged across states), they can be affected by experts' biases in some states. The set of states in which the followers fail to filter out the experts' biases turns out to be small, but these happen to be the states that are pivotal for the equilibrium outcome.

Consider those experts and followers whose biases are weak enough that, if they were certain about the coordination outcome, they would choose the corresponding action. In particular, the optimal actions of agents in these two groups depend on their respective biases only when they are uncertain about the outcome, and are aligned otherwise. When experts are better informed than followers, followers do not know when experts are uncertain about the outcome, and thus believe that experts' action are likely to be independent of their bias. Hence followers effectively ignore experts' biases when evaluating their actions. However, in those contingencies in which experts are uncertain of the outcome, experts' actions reflect their bias; followers, not knowing that experts are uncertain, ignore the effect of the experts' biases. Consequently, in these contingencies, followers' choices comply with experts' interests.

Even though contingencies in which experts cannot predict the coordination outcome are rare when experts are well informed, strategic complementarities can multiply the effect so as to make the action preferred by experts considerably more likely to be adopted. Starting from an equilibrium of the coordination game without experts, introducing experts leads to more followers choosing the experts' preferred action in contingencies where the outcome would otherwise be very close to a tie. This in turn leads to more choices of that action in other nearby contingencies, with followers adopting it more often, and so on, multiplying the effect. The size of the effect at each step of this contagion vanishes as experts become very well informed, but the total effect generally remains large.

We explicitly characterize the equilibrium of the game with experts and followers. The characterization shows that the presence of experts generally affects the likelihood of coordination on each of the two actions. The direction of experts' influence depends on their biases in a non-monotone way. As the share of experts who are moderately biased in favor A increases, followers become more likely to coordinate on A. However, as the share of experts who are partisan for A—those who choose action A irrespective of their belief—increases, followers can become more likely to coordinate on B. In line with the empirical marketing literature, the magnitude of experts' influence depends primarily on their informational advantage and their visibility. When the experts are sufficiently visible and better informed than followers, their influence is maximal: the coordination outcome is independent of the state and of the followers' preferences, it depends only on the experts' preferences.

The predictions of our model can be applied to a wide class of settings that combine coordination with social learning. As noted earlier, this combination arises naturally in the diffusion of goods with network externalities and in the onset of social movements. At a technical level, our model is also related to the literature on social learning and that on global games. We discuss applications and related literature in Section 5.

2 Model

The model features two sets of players: a continuum of experts and a continuum of followers, each with unit measure. Followers are indexed by i and experts are indexed by j. Both types of players choose one of two actions, denoted A and B. The *coordination outcome* of the game is A if the fraction of followers who choose action A is greater than a stochastic parameter $\theta \in [0,1]$, and it is B otherwise.[1] Although we normalize the measure of both groups of agents to 1, the group of experts should be thought of as being small enough as to have a negligible direct impact on the majority action. This is reflected in the assumption that the coordination outcome is determined purely by the action choices of the followers.

Followers' payoffs depend on their own action and the coordination outcome. Hence, the preferences of a follower i are described by $\left(\pi^i_{AA}, \pi^i_{AB}, \pi^i_{BA}, \pi^i_{BB}\right)$, where, π^i_{XY} is her payoff when she plays $a^i_f = X$ and the coordination outcome is Y. To capture the coordination motive, we assume that there is a stronger incentive to choose A when the coordination outcome is A, that is, $\pi^i_{AA} - \pi^i_{AB} + \pi^i_{BB} - \pi^i_{BA} > 0$. Under this assumption, follower i's best response takes a simple form: she plays $a^i_f = A$ if she believes that the coordination outcome is A with a probability greater than $p^i_f = \frac{\pi^i_{BB} - \pi^i_{AB}}{\pi^i_{AA} - \pi^i_{AB} + \pi^i_{BB} - \pi^i_{BA}}$, and she plays $a^i_f = B$ otherwise. The greater is the cutoff p^i_f, the more inclined follower i is to choose action B. In what follows, we refer to p^i_f as the *bias* of follower i and we assume that p^i_f is a measurable function of i.

For all i such that $p^i_f < 0$ or $p^i_f > 1$, A or B, respectively, is a dominant action. For $X \in \{A, B\}$, we denote by s^X_f the fraction of followers for whom X is weakly dominant. We refer to these followers as X-partisans and to nonpartisan followers as independents. In the sequel, we assume that $s^A_f > 0$, $s^B_f > 0$, and $s^A_f + s^B_f < 1$.

Experts' payoffs are assumed to have the same structure. For an expert j, we denote by p^j_e the critical belief that characterizes her best response, and by s^A_e and s^B_e the fraction of A-partisan and

[1] In case of a tie, the coordination outcome may be chosen arbitrarily.

B-partisan experts, respectively.

If $\theta < s_f^A$, since A-partisans always choose A, the coordination outcome is A irrespective of the actions of the other players. Similarly, if $\theta > 1-s_f^B$, the coordination outcome is B. We focus on the coordination outcome for $\theta \in \left[s_f^A, 1 - s_f^B\right]$, where it depends on the independent followers' actions. In this region, the coordination motive creates a coordination problem among the independent followers.

In order to capture the idea that agents may be uncertain about the actions of others, we assume that experts and followers possess private information. More specifically, each follower receives private information about the state θ consisting of two parts: an exogenous signal and a collection of observations of experts' actions.

The information structure and timing are as follows. First, the state θ is drawn from a uniform distribution on $[0,1]$. Then each independent expert j receives a private signal $x_e^j = \theta + \sigma_e \varepsilon_e^j$, and each independent follower i receives a private signal $x_f^i = \theta + \sigma_f \varepsilon_f^i$. The experts' errors ε_e^j and the followers' errors ε_f^i are drawn from continuous distributions F and G, with densities f and g, respectively, and with support $[-1,1]$.[2] Errors are independent across agents and independent of θ. To simplify the exposition, we assume that $\sigma_e + \sigma_f \leq \frac{1}{2} \min\left(s_f^A, s_f^B\right)$.[3]

After the signals have been observed and before followers choose actions, experts simultaneously choose their actions a_e^j. In addition to her private signal x_f^i, each follower observes a random sample of n expert actions, where, for simplicity, $n \in \mathbb{N}$ is fixed across followers. The sample is private and taken with uniform probability over all experts, regardless of type. Followers do not observe the biases or signals of the experts in their sample (see Section 6 for a discussion of the role of this assumption). After observing expert actions, the followers simultaneously choose actions a_f^i.

A strategy for an independent expert maps each signal x_e^j to an action $a_e^j \in \{A, B\}$. Letting $\lambda^i \in \{0, \ldots, n\}$ denote the number of actions A in follower i's sample, a strategy for an independent follower maps each pair (x_f^i, λ^i) to an action $a_f^i \in \{A, B\}$. A strategy for an expert is monotone if there is some threshold signal above which she chooses B and below which she chooses A. A strategy s^i for a follower is monotone if (i) $s^i(x, \lambda) = B$ implies that $s^i(x', \lambda') = B$ whenever $x' \geq x$

[2] For most of our results, the bounded support of the error terms simplifies exposition but is not necessary.

[3] The role of the assumption $\sigma_e + \sigma_f \leq \frac{1}{2} \min\left(s_f^A, s_f^B\right)$ is to avoid boundary issues when deriving players' posterior beliefs about θ. The assumption that θ is uniformly distributed allows us to concentrate on the updated beliefs of the players conditional on their signals without taking into account the information contained in the prior distribution. Similar results can be obtained for more general continuous priors in the limit case $(\sigma_e, \sigma_f) \to (0,0)$.

and $\lambda' \leq \lambda$, and (ii) $s^i(x, \lambda) = A$ implies that $s^i(x', \lambda') = A$ whenever $x' \leq x$ and $\lambda' \geq \lambda$. We restrict attention to monotone strategies. All parameters of the model, including all distributions, are common knowledge.

3 Coordination without Opinion Leaders

Before we solve the main model, we consider the coordination game in which the followers do not observe the experts' actions (in the notation of Section 2, $n = 0$). In this case, the game reduces to a simultaneous move game among the independent "followers". Using standard global games techniques, we derive the unique monotone Bayesian Nash equilibrium, which has the property that the coordination outcome is A when θ is smaller than a critical threshold θ_0^*, and B when θ is greater than θ_0^*. In the sequel, we refer to θ_0^* as the pivotal state. Since the A-partisan followers always play A and the B-partisan followers always play B, $\theta_0^* \in \left[s_f^A, 1 - s_f^B\right]$.

Given the threshold θ_0^*, let $\pi_f\left(x_f^i, \theta_0^*, \sigma_f\right)$ be the posterior belief that follower i assigns to A being the coordination outcome after receiving the signal x_f^i; that is, $\pi_f\left(x_f^i, \theta_0^*, \sigma_f\right) = \Pr\left(\theta < \theta_0^* | x_f^i\right)$. A straightforward application of Bayes' rule gives that for all $\theta_0^* \in \left[s_f^A, 1 - s_f^B\right]$,

$$\pi_f\left(x_f^i, \theta_0^*, \sigma_f\right) = 1 - G\left(\frac{x_f^i - \theta_0^*}{\sigma_f}\right).$$

Independent follower i chooses A if and only if her posterior belief $\pi_f\left(x_f^i, \theta_0^*, \sigma_f\right)$ exceeds the critical probability p_f^i. By the definition of the pivotal state θ_0^*, the outcome is a tie when $\theta = \theta_0^*$. Since θ is defined to be the share of followers choosing A required for a tie, θ_0^* must equal the share of followers choosing A in the pivotal state. Taking into account both partisan and independent followers, the *pivotal condition* is given by

$$\theta_0^* = s_f^A + \left(1 - s_f^A - s_f^B\right) \Pr\left(\pi_f\left(x_f^i, \theta_0^*, \sigma_f\right) > p_f^i \mid \theta_0^*\right), \tag{1}$$

where i is a uniformly drawn independent follower.

When the pivotal state is realized, followers' beliefs p_f^i reflect only the noise in their signals rather than useful information about the coordination outcome. As a result, their beliefs are diffuse in this contingency, as indicated by the following lemma.

Lemma 1. *Posterior beliefs in the pivotal state* $\theta_0^* \in \left[s_f^A, 1 - s_f^B\right]$ *are distributed uniformly on* $[0,1]$ *regardless of the noise distribution. Thus for any* $p \in [0,1]$ *and any* i*, we have*

$$\Pr\left(\pi_f\left(x_f^i, \theta_0^*, \sigma_f\right) > p \mid \theta_0^*\right) = 1 - p. \tag{2}$$

The uniform property of posterior beliefs in the lemma has been used in Guimaraes and Morris (2007) and Steiner (2006). For convenience, we include the proof in the appendix.

Integrating (2) across the population of independent followers, the pivotal condition (1) implies that

$$\theta_0^* = s_f^A + \left(1 - s_f^A - s_f^B\right)(1 - p_f).$$

where p_f denotes the average bias of the independent followers. In the absence of experts, the equilibrium outcome aggregates the preferences of followers in a natural way. The above expression shows that θ_0^* is increasing in s_f^A and decreasing in s_f^B and in p_f. This means that action A becomes the coordination outcome if it is dominant for sufficiently many followers and/or independent followers are sufficiently biased in its favor.

The channel through which the independent followers' biases affect the outcome is best understood through the pivotal condition (1). The pivotal state θ_0^* is determined by the best responses of the followers *in the pivotal state*. In this state, the independent followers receive inconclusive signals making them unsure about the coordination outcome, thereby suppressing the significance of the coordination motive. Consequently, their individual choice is affected by their individual bias p_f^i and the aggregate action is a monotone function of their average bias p_f.

The analysis in the next section, where followers observe expert actions, also focuses on behavior in the pivotal state, in which there is considerable strategic uncertainty. In the pivotal state, the behavior of experts who are uncertain about the coordination outcome is affected by their intrinsic biases and it turns out that followers do not filter out the experts' biases. As a result, the experts' biases affect the equilibrium outcome.

4 Coordination with Opinion Leaders

We now return to the model of Section 2 in which each follower observes a random sample of $n \geq 1$ expert actions. We restrict attention to weak perfect Bayesian equilibria in monotone strategies.

As above, any monotone equilibrium gives rise to a pivotal state θ_n^*, such that the coordination outcome is A for $\theta < \theta_n^*$ and B for $\theta > \theta_n^*$. The equilibrium analysis below has the same structure as the analysis of the benchmark game. We take the value of θ_n^* as given, compute the best responses of both the experts and the followers to θ_n^*, and then use the requirement that the outcome in the pivotal state is a tie.

4.1 Experts' Behavior

We begin by considering the best responses of experts. Given the threshold θ_n^*, independent expert j chooses A if and only if her posterior belief $\pi_e\left(x_e^j, \theta_n^*, \sigma_e\right)$ that $\theta < \theta_n^*$ exceeds a critical probability p_e^j. Let $l\left(\theta, \theta_n^*, \sigma_e\right)$ denote the probability that a randomly selected expert chooses A in state θ given the threshold θ_n^*. Taking into account both partisan and independent experts, we have

$$l\left(\theta, \theta_n^*, \sigma_e\right) = s_e^A + \left(1 - s_e^A - s_e^B\right) \Pr\left(\pi_e\left(x_e^j, \theta_n^*, \sigma_e\right) > p_e^j | \theta\right), \tag{3}$$

where j is a randomly chosen independent expert.

The analysis of experts' behavior is particularly simple if the realized state is sufficiently far from the pivotal state θ_n^*. In that case, every independent expert correctly forecasts the coordination outcome and chooses the corresponding action. Thus we have

$$l\left(\theta, \theta_n^*, \sigma_e\right) = \begin{cases} 1 - s_e^B \text{ if } \theta \leq \theta_n^* - 2\sigma_e, \\ s_e^A \text{ if } \theta \geq \theta_n^* + 2\sigma_e. \end{cases}$$

The analysis of the experts' behavior is also relatively simple in the pivotal state θ_n^*. By Lemma 1, experts' posterior beliefs are uniformly distributed on $[0,1]$ when $\theta = \theta_n^*$. Therefore, the ex ante probability that independent expert j chooses A in state θ_n^* is $1 - p_e^j$, and the same reasoning as in the benchmark case $n = 0$ gives the following lemma.

Lemma 2. *Let p_e denote the average bias of the independent experts. For any $\theta_n^* \in \left[s_f^A, 1 - s_f^B\right]$*

and any $\sigma_e > 0$, we have

$$l\left(\theta_n^*, \theta_n^*, \sigma_e\right) = s_e^A + \left(1 - s_e^A - s_e^B\right)\left(1 - p_e\right).$$

In particular, in the pivotal state, the share of experts choosing action A is strictly decreasing in p_e and is independent of θ_n^* and σ_e.

We have made two observations: (i) in typical states—those outside a σ_e-neighborhood of θ_n^*—the independent experts' biases do not influence the distribution of expert actions, and (ii) in the pivotal state, the share of experts choosing A decreases with the independents' average bias. Both observations are important for the analysis of followers' behavior below. When σ_e is small relative to σ_f, for any signal realization x_f^i, follower i believes with high probability that the state is typical, and because of (i) followers effectively neglect the experts' biases when evaluating their actions; given followers' information, contingencies in which experts' biases affect expert choices are unlikely. Followers neglect the experts' biases even in the pivotal state in which, because of (ii), experts' biases do shape their actions. Since the equilibrium is determined by the followers' behavior in the pivotal state, the equilibrium outcome reflects the independent experts' bias p_e even though p_e is commonly known and followers correctly account for it when forming beliefs.

4.2 Followers' Behavior

Next we analyze the followers' behavior. Let $p_f\left(x, \lambda, \theta_n^*, \sigma_e, \sigma_f\right) = \Pr_{\sigma_e, \sigma_f}\left(\theta < \theta_n^* | x, \lambda\right)$ denote the posterior probability that a follower assigns to A becoming the coordination outcome after observing a signal x and a number λ of experts choosing A (given the threshold θ_n^*). Bayes' rule gives that for all $\theta_n^* \in \left[s_f^A, 1 - s_f^B\right]$,

$$p_f\left(x, \lambda, \theta_n^*, \sigma_e, \sigma_f\right) = \frac{\int_{\theta \leq \theta_n^*} g\left(\frac{x-\theta}{\sigma_f}\right) \Pr(\lambda|\theta) d\theta}{\int_{\theta} g\left(\frac{x-\theta}{\sigma_f}\right) \Pr(\lambda|\theta) d\theta}. \tag{4}$$

The distribution of observed experts' behavior $\Pr(\lambda|\theta)$ depends on the realized state θ and on the experts' strategies. Conditional on θ, λ is binomially distributed with sample size n and success probability $l\left(\theta, \theta_n^*, \sigma_e\right)$.

Let $v\left(\theta, \theta_n^*, \sigma_e, \sigma_f\right) \in [0, 1]$ denote the share of independent followers choosing A in state θ when all agents play best responses given θ_n^*. We have

$$v(\theta, \theta_n^*, \sigma_e, \sigma_f) = \Pr\left(p_f\left(x_f^i, \lambda^i, \theta_n^*, \sigma_e, \sigma_f\right) > p_f^i \mid \theta\right).$$

As in the benchmark game, θ_n^* must satisfy the condition

$$\theta_n^* = s_f^A + \left(1 - s_f^A - s_f^B\right) v\left(\theta_n^*, \theta_n^*, \sigma_e, \sigma_f\right), \tag{5}$$

which states that, in the pivotal state θ_n^*, action A is chosen by exactly the right proportion of followers as to make the outcome a tie.

Due to the symmetry of the model with respect to θ, $v(\theta_n^*, \theta_n^*, \sigma_e, \sigma_f)$ is independent of θ_n^* and depends only on $\frac{\sigma_e}{\sigma_f}$. It follows from (5) that the pivotal state is uniquely determined.

Proposition 1. *The game has a unique monotone equilibrium. The equilibrium pivotal state θ_n^* depends on (σ_e, σ_f) only through the ratio $\frac{\sigma_e}{\sigma_f}$.*

Proofs are in the appendix.

4.3 Experts' influence

The uniqueness of the equilibrium allows us to quantify the influence of experts as a function of their informational advantage $\frac{\sigma_f}{\sigma_e}$, their degree of partisanship $\left(s_e^A, s_e^B\right)$, and their visibility n.

The following proposition states that experts can be influential only when they have information that could be of value to the followers.

Proposition 2. *As $\frac{\sigma_f}{\sigma_e} \to 0^+$, the pivotal threshold θ_n^* converges to the pivotal threshold θ_0^* of the game without experts.*

The intuition is as follows. When followers are better informed than experts, followers can tell from their signals when the experts cannot predict the coordination outcome. Therefore, followers know when experts' actions are biased by their preferences, and they can factor out this bias.

Let us now focus on the limit as $\frac{\sigma_f}{\sigma_e} \to +\infty$, in which experts' signals are much more precise than the followers'. Without loss of generality, we fix σ_f and let $\sigma_e \to 0^+$. In this limit, followers' posterior beliefs are relatively simple to compute. Let $\pi_f\left(x_f^i, \theta_n^*, \sigma_f\right) = \Pr_{\sigma_f}\left(\theta < \theta_n^* | x_f^i\right)$ denote the “pre-expert” probability that action A prevails evaluated by follower i conditioning only on her private signal x_f^i (as opposed to the “post-expert” probability $p_f\left(x_f^i, \lambda^i, \theta_n^*, \sigma_e, \sigma_f\right)$). For any $\theta \neq \theta_n^*$, and sufficiently small σ_e, all independent experts choose the action that matches the coordination outcome. Thus, for $\theta > \theta_n^*$, $\lim_{\sigma_e \to 0^+} l(\theta, \theta_n^*, \sigma_e) = s_e^A$ and for $\theta < \theta_n^*$, $\lim_{\sigma_e \to 0^+} l(\theta, \theta_n^*, \sigma) = 1 - s_e^B$. Hence, using Bayes rule, $p_f(x, \lambda, \theta_n^*, \sigma_e, \sigma_f)$ converges to

$$\frac{\pi_f(x, \theta_n^*, \sigma_f)\binom{n}{\lambda}\left(1 - s_e^B\right)^\lambda \left(s_e^B\right)^{n-\lambda}}{\pi_f(x, \theta_n^*, \sigma_f)\binom{n}{\lambda}(1 - s_e^B)^\lambda (s_e^B)^{n-\lambda} + (1 - \pi_f(x, \theta_n^*, \sigma_f))\binom{n}{\lambda}(s_e^A)^\lambda (1 - s_e^A)^{n-\lambda}}.$$

Straightforward algebraic manipulation leads to the following lemma.

Lemma 3. *For every x, λ, and θ_n^*, we have*

$$\lim_{\frac{\sigma_f}{\sigma_e} \to +\infty} p_f(x, \lambda, \theta_n^*, \sigma_e, \sigma_f) = \frac{\pi_f(x, \theta_n^*, \sigma_f)}{\pi_f(x, \theta_n^*, \sigma_f) + (1 - \pi_f(x, \theta_n^*, \sigma_f))\left(\frac{s_e^A}{1 - s_e^B}\right)^\lambda \left(\frac{1 - s_e^A}{s_e^B}\right)^{n-\lambda}}.$$

According to the lemma, in the limit, followers treat the experts’ choices as informative signals but ignore the independent experts’ incentives when evaluating these signals. In particular, the posterior belief increases in the number λ of experts’ choices of action A, but does not depend on the bias of the independent experts.

Next we characterize the pivotal state in the limit as $\frac{\sigma_f}{\sigma_e} \to +\infty$ using condition (5). Lemma 1 implies that followers’ beliefs $\pi_f\left(x_f^i, \theta_n^*, \sigma_f\right)$ before observing experts’ choices are uniformly distributed on $[0, 1]$, and Lemma 2 determines the distribution of experts’ actions in the pivotal state. Finally, Lemma 3 describes, in the limit, followers’ beliefs after observing experts’ choices. Combining these lemmas yields the following proposition.

Proposition 3. *As $\frac{\sigma_f}{\sigma_e} \to +\infty$, the pivotal threshold θ_n^* converges to*

$$\theta_n^{**} = s_f^A + \left(1 - s_f^A - s_f^B\right) \Pr\left(\frac{\pi}{\pi + (1 - \pi)\left(\frac{s_e^A}{1 - s_e^B}\right)^\lambda \left(\frac{1 - s_e^A}{s_e^B}\right)^{n-\lambda}} > p_f^i\right),$$

where π, λ and p_f^i are independent random variables with $\pi \sim U[0, 1]$, $\lambda \sim B\left(n, s_e^A + \left(1 - s_e^A - s_e^B\right)(1 - p_e)\right)$, and i is a randomly chosen independent follower.[4]

[4]Here $B(n, p)$ denotes the binomial distribution for n draws with probability p.

Our main results follow from this proposition. First, the proposition implies that the independent experts' bias p_e has an unambiguous effect on the pivotal state θ_n^{**}. Since the distribution of λ is decreasing in p_e (in the sense of first-order stochastic dominance), and the posterior belief $\pi / \left(\pi + (1-\pi) \left(\frac{s_e^A}{1-s_e^B} \right)^{\lambda} \left(\frac{1-s_e^A}{s_e^B} \right)^{n-\lambda} \right)$ is increasing in λ, A becomes more likely to be the coordination outcome if independent experts' biases shift in its favor.

Corollary 1. *The pivotal threshold θ_n^{**} is strictly decreasing in the independent experts' average bias p_e.*

The impact of the experts' bias becomes large when the number n of observed expert actions becomes large. Consider $\lim_{n \to \infty} \theta_n^{**}$, corresponding to the equilibrium outcome in the ordered limit in which first $\frac{\sigma_f}{\sigma_e} \to +\infty$ and then $n \to \infty$. As the following corollary indicates, in this limit, the experts' influence is maximal: in all states in which the independent followers face a coordination problem—that is, in all $\theta \in \left[s_f^A, 1 - s_f^B \right]$—the coordination outcome is independent of the state and of the independent followers' preferences; it depends only on the experts' preferences.

Corollary 2. *For each s_e^A and s_e^B, there exists p_e^* such that $lim_{n \to \infty} \theta_n^{**} = s_f^A$ if $p_e > p_e^*$ and $lim_{n \to \infty} \theta_n^{**} = 1 - s_f^B$ if $p_e < p_e^*$.*

As the sample of observed experts' actions increases in size, followers view their samples as increasingly reliable indicators of the coordination outcome. For instance, if partisan experts are evenly distributed, i.e. $s_e^A = s_e^B$, then the followers believe that the general population will coordinate on the action supported by the majority of experts in their private sample. In this case, $p_e^* = 1/2$, so the outcome favored by the average expert prevails in all states in which there is a coordination problem.

The preceding analysis highlights the influential role of independent experts. Partisan experts can also affect the outcome, but the nature of their influence is qualitatively different. The reason is that the strategy of partisan experts does not depend on the state, so followers can correct for their bias even though they do not know the state. In particular, if all experts become partisans, their influence vanishes. Formally, as $s_e^A + s_e^B \to 1$, $l(\theta, \theta_n^*, \sigma_e) \to s_e^A$ in all states θ, so the distribution of λ^i is independent of θ. Therefore, the "post-expert" belief $p_f \left(x_f^i, \lambda^i, \theta_n^*, \sigma_e, \sigma_f \right)$ tends to the "pre-expert" belief $\pi_f \left(x_f^i, \theta_n^*, \sigma_f \right)$, and θ_n^* tends to θ_0^*.

Partisan experts cannot be influential on their own, but they can affect the influence of independent experts. However, unlike independent experts, partisan experts do not necessarily move the outcome in their preferred direction. Proposition 3 implies that $\theta_n^{**} \to 1 - s_f^B$, as $s_e^A \to 0$ and $n \to \infty$. Hence, as the share of A-partisans vanishes, coordination outcome A becomes more likely. The intuition is as follows. When experts are better informed than followers, followers always believe with high probability that experts know the coordination outcome, and therefore that all independent experts choose the same action. In the pivotal state, followers typically observe incongruous signals in their expert samples. When A-partisans are rare, followers conclude from these incongruous signals that all independent experts played A and that the B actions are due to B-partisans in their sample. This belief leads them to play A.

5 Applications and Related Literature

Coordination and social learning naturally interact in the diffusion of technologies and goods with network externalities, and in social movements. In this section, we discuss the assumptions and predictions of our model in these two environments, together with the related literature.

5.1 Social Movements

A social movement can be viewed as a coordination problem among citizens. Supporters of a regime change can gain privileged status if the status quo regime is replaced, but run the risk of being punished if it stays in place. Supporters of the status quo regime face similar but opposite risks. When viewed as a regime-change game, the definition of θ in our model captures the idea that the success of a social movement depends to a large extent on the mobilization rate (DeNardo 1985). The parameter θ is a proxy for the status quo regime's strength and repressive capacity.[5]

5.1.1 Spontaneous Uprisings

Prominent historical cases of revolutions have been sparked by small groups of unorganized protesters. Our model can capture a spontaneous revolution ignited by unorganized activists. Experts are vis-

[5] Revolutions and political regime changes are complex events brought on by many factors. We do not explicitly model the economic causes of regime change and focus instead on the coordination problem among citizens. See, e.g., Acemoglu and Robinson (2000, 2001, 2006), Lizzeri and Persico (2004), and Ellis and Fender (2011) for more on the economic origins of regime transitions.

ible citizens with superior information about the strength or popularity of the status quo regime, and have the opportunity to participate in the initial phase of a social movement. The experts cannot collude, but their spontaneous coordination is facilitated by an external event that reveals the malignant nature of the regime. Followers are the mass population. They decide whether to join the movement after having heard of the initial mobilization. The assumption that each follower only sees a private sample of experts captures the idea that the regime prevents the mass media from reporting the initial mobilization, so the population learns only through interpersonal ties.[6]

The Leipzig demonstrations of 1989 provide a natural case with which to confront the assumptions and predictions of our model with the historical evidence. According to Lohmann (1994), "[t]he first three demonstrations in Leipzig were characterized by relatively low turnout, yet they played a critical role in triggering the protest throughout the GDR." The local newspaper, the Leipziger Volkzeitung, downplayed the first revolts and described the participants as a "mob... with obvious anti-socialist tendency." Moreover, the regime prevented the participants from sending public signals about the movement. Nevertheless, these small-scale, peaceful protests were soon followed by massive demonstrations in which "large numbers of individuals were able to coordinate their participation decisions spontaneously." Lohmann notices that "social embeddedness and personal networks... have influenced individual participation decisions" (see also Opp and Gern 1993) but "political entrepreneurship and organization played a secondary role." These facts are consistent with our finding that a small group of unorganized but visible individuals can influence the likelihood of a massive, spontaneous mobilization (see Corollaries 1 and 2). Lohmann argues that the participants in the first demonstrations were not necessarily representative of the population, but they were not extremists, and their number grew rapidly from several thousand to hundreds of thousands. In contrast, "the organized demonstrations of the third, fourth, and fifth cycles were failures." She attributes this failure to the fact that the organized groups who participated in the latter cycles had more extreme preferences, and their "organized efforts appear to have been discounted by the people." This is consistent with our findings that nonpartisan leaders can be influential even if their preferences do not coincide with the preferences of the followers, but partisans have, if anything, an adverse effect on mass mobilization.

[6] Empirically, knowing someone who is already involved in a social movement is one of the strongest predictors of participation. See, among others, McAdam (1986), McAdam and Paulsen (1993), and Opp and Gern (1993).

The Hungarian revolt of 1956 and the Jasmine revolution in Tunisia also began as small, spontaneous, and unorganized demonstrations (see, e.g., Irving 1981 and Malewski 2011). The participants in the initial phase of the Hungarian revolt were mostly students, journalists, or writers; they were visible and informed citizens, but not extremists.[7] The main actors of the first protests in the Tunisian revolution were young, unemployed graduates (Honwana 2011). This social group was particularly aware of the consequences of the poor economic management of President Ben Ali. The mass media did not report the first demonstrations in Sidi Bouzid, but the population learned about it via private channels and social media.[8] In both cases, the movement rapidly escalated into a massive uprising that led to the fall of the government.

In an influential series of papers, Lohmann (1993, 1995, 2000) analyzes various versions of a signalling model of spontaneous and unorganized political action. Citizens have dispersed information about the status quo, and can signal that information by taking political action. A central feature of Lohmann's model is that activists are motivated by the likelihood of being pivotal over the outcome, while in our model, activists' actions are driven by their belief about the success of the movement, independently of their action. Moreover, since the aggregate level of political participation is publicly observed and the preference distribution is common knowledge, in equilibrium, citizens and policy makers can factor out the preferences of activists. Therefore, in contrast to our model, the outcome cannot be systematically biased by the activists' preferences.

5.1.2 Predicting Uprisings

A common feature of the aforementioned uprisings is that most political observers, social scientists, and activists did not foresee them (Kuran 1991, Gause 2011). Our results (see Corollary 2) can provide an explanation for the seemingly unpredictable nature of revolutions: under some conditions, the likelihood of a revolution is insensitive to the state and the preferences of the population (i.e., the followers) and depends only on the preferences of a negligible share of the population (i.e.,

[7]In fact, many of the writers who participated in the uprising did not show any sign of opposition to the regime before the movement began (Kuran 1991).

[8]The role and importance of the social media the Arab Spring is still a subject of debate among scholars and observers (Beaumont 2011). Nevertheless, it appears that during the initial phase of the Tunisian uprising, the social media were mostly used to post photos and videos of the demonstrations. In our model, this scenario can be modelled by having followers observe a private, noisy signal of the aggregate mobilization among experts, instead of a private sample of experts. We conjecture that our results continue to hold in that case: leaders' actions are still influential, and the likelihood of a mass protest still depends on the experts' preferences.

the experts). This result is consistent with the empirical finding that the degree of mass discontent is a poor predictor of revolutions (see, e.g., Snyder and Tilly 1972, Tilly, Tilly, and Tilly 1975, and Skocpol 1979).

Building on Granovetter's (1978) threshold model, Kuran (1989) develops a dynamic model of revolutionary bandwagons to explain the unpredictability of uprisings. His main finding is that the rate of mobilization is a discontinuous function of the preference distribution of the citizens. A key assumption is that citizens' willingness to mobilize is not driven by the likelihood of success of the social movement as in our model, but rather by the participation rate. Contagion occurs via a direct externality from the extremists (who prefer to mobilize irrespective of the participation rate) on the moderates. In contrast, in our paper, contagion occurs via an informational channel in which extremists have little impact on the outcome, and the mass protest is sparked instead by moderate activists.

5.1.3 Organized Insurgencies

A large strand of literature discusses the role of organized groups in social movements. In a complete information environment, coordination games have multiple equilibria. In that case, the standard account is that an insurgency can change citizens' conjectures about one another and create focal points (Schelling 1960). A more recent literature introduces incomplete information to analyze the informational role of insurgencies. For instance, Bueno De Mesquita (2010) shows that a vanguard can use political violence to reveal information about the (un)popularity of the status quo regime. The author assumes that the level of political violence depends not only on the vanguards' effort but also on the type of the regime. Under this assumption, an uninformed vanguard can affect citizens' willingness to mobilize, even if it has no direct effect on the regime.[9]

This paper differs qualitatively from the literature on the role of organized groups in social movements in that the influence of the activists is totally decentralized. The experts from our model can be viewed as potential recruits to the insurgency. Our assumptions on payoffs require that some experts join the insurgency only if they believe that the regime is sufficiently likely to collapse. Members of the mass population (the followers) choose whether to support the regime or

[9] The status quo regime can also attempt to influence beliefs. For instance, Edmond (2011) shows that when media manipulation is costly, a privately informed regime can use the media as a costly signal jamming technology to convey credible information about its strength, and increase its probability of survival.

not. For simplicity, the regime is not directly affected by the insurgency, and its survival depends only on its popular support. Hence, unlike in Bueno De Mesquita (2010), the insurgency is harmless and uninformative, but it induces some citizens to make a visible and informative choice. The assumption that followers observe a private sample of experts means that the regime prevents the insurgency from communicating publicly about its recruitment, but the decision of an individual to join the insurgency is observed by the members of his community. Our results show that the preferences of the insurgency recruits can have an important impact on the regime, even if their preferences are not representative of the population.

Since Intriligator and Brito (1988), the literature on nonconventional warfare has identified the key role of recruitment in perpetuating a conflict (see, e.g., Faria and Arce 2005). These dynamic models typically assume that the popular support for the insurgency affects its recruitment. Our game theoretic approach shows that its recruitment can also affect its popular support. In terms of counterinsurgency strategy, our results suggest that targeting the incentives of the potential recruits of the insurgency (i.e., changing the bias of independent experts) by creating economic or educational opportunities can be more effective than fighting its current members.

5.2 Opinion Leaders and Viral Marketing

The adoption of a convention, a network technology, or a social good (social media, software, movies, fashion) generates network externalities (Katz and Shapiro 1985): the value of a product to a each consumer increases with the number of people who use it, thereby creating a coordination motive. Social learning helps consumers to acquire information about which product will become the most popular. In this environment, the experts can be thought of as product reviewers who choose to endorse one product and receive some benefit from endorsing a product that ends up being widely adopted. Our results indicate that, even if the pool of experts is negligibly small relative to the population, the outcome is biased toward their preferences.

Alternatively, the experts can be thought of as well informed, early adopters, while the followers are prospective consumers who observe the choices of the early adopters among their social ties. Since the seminal work of Katz and Lazarsfeld (1955) and Arndt (1967), a large literature in marketing and sociology has shown that consumer-to-consumer communication is much more effective than firm-to-consumer communication. Numerous empirical studies have consistently found that the diffusion of information among consumers is overwhelmingly driven by a small number of individuals, initially referred to as "opinion leaders" by Katz and Lazarsfeld (1955).[10] These influentials are typically visible individuals with product-specific or marketplace expertise. Moreover, they are often early adopters, and their influence on other individuals is restricted to their interpersonal ties (see, e.g., Price and Feick 1984 or Watts and Dodds 2007). These patterns are consistent with our assumptions that experts move early, have superior information, and that each follower observes a private sample of experts.

The difficulty of defining and observing causal social influence has led some scholars to question the impact of opinion leaders (Watts 2007, Watts and Dodds 2007, Aral 2011). Our approach contributes to this debate by allowing us to quantify the influence of opinion leaders and to determine the main drivers of their influence. Our results show that leaders can be very influential, and that their influence depends primarily on their expertise and their visibility. More precisely, what matters is not the degree of expertise of the leaders, but their relative expertise compared to the followers (see Proposition 1). Hence, leaders' influence also depends on followers' characteristics (Watts and Dodds 2007). An important question in the marketing literature is whether similarity affects social influence (Price and Feick 1984, Gilly, Graham, Wolfinbarger, and Yale 1998). Our model shows that leaders can be influential even if their preferences differ from that of the followers. However, as leaders become too biased, their influence vanishes because their choices are no longer informative to the followers. Hence, a correlation between similarity and social influence may simply reflect a concern for informational relevance, rather than homophily or ease of communication (Price and Feick 1984). In terms of marketing strategy, our results suggests that it makes sense for a company to adapt its product to the specific tastes of these opinion leaders or to provide discounts, but it can be counterproductive to pay them to adopt its product.

5.3 Leadership and Coordination

This paper is related to a broader literature on leadership and coordination. Dewan and Myatt (2007, 2008, 2012) study leadership in the context of a political party conference: party activists

[10]Depending on the breadth, depth, and origin of their superior information, and their visibility, these agents are variously called opinion leaders (Katz and Lazarsfeld 1955), influentials (Merton 1957), early adopters (Baumgarten 1975), market mavens (Feick and Price 1987, Gladwell 2000), or social hubs (Goldenberg, Han, Lehmann, and Hong 2009).

receive private information about the best policy platform, and the desire for party unity generates a coordination motive. Leaders coordinate followers by publicly communicating their information. In Dewan and Myatt (2008) and (2012), leaders improve coordination among followers, but they do not systematically bias the coordination outcome. Instead, the authors focus on the impact of leaders' information precision and communication skills on their leadership. In (2007), followers care about pivotally influencing the outcome. The authors show, among other things, that even though all players have the same ex-ante preferences, the leader wants to bias the followers' strategies toward the alternative that requires the least coordination.

Corsetti, Dasgupta, Morris, and Shin (2004), Edmond (2011), and Ekmekci (2009) characterize the influence of a large player who can act as a coordination device for followers and who internalizes the impact of her action. In contrast, we focus on the case of many experts with negligible individual influence who cannot act as a coordination device, making the mechanism underlying leaders' influence quite different. Hence, our paper differs from the above contributions in that it provides a theory of decentralized leadership in which leaders' influence is incidental.

6 Discussion

The influence of experts in our model results from a combination of social learning and coordination. To clarify the roles that these two features play, consider the following variant of the model with no coordination motive: instead of the predominant action being determined by followers' choices, suppose that the pivotal state θ^* is exogenously fixed; independent agents prefer to choose A if and only if $\theta < \theta^*$.

Experts' biases affect the coordination outcome only when experts' behavior depends on their own biases, and followers' behavior is influenced by experts' behavior. As in the model with coordination, the action chosen by an independent expert depends on her bias only in a σ_e-neighborhood of θ^* where she is uncertain of the optimal action. When experts have very precise information, such contingencies are rare. Therefore, in the absence of a coordination motive, the ex ante probability that experts' biases affect their own behavior vanishes as the precision of their information increases. Likewise, the action chosen by a follower depends on the experts' behavior only in a σ_f-neighborhood of θ^* where she is uncertain of the optimal action. When followers have very precise information about θ, contingencies in which they are uncertain are rare. Therefore, in the absence of a coordination motive, the ex ante probability that expert behavior influences followers' behavior vanishes as the precision of the followers' information increases.

When θ^* is determined endogenously by the followers' behavior, the effect of experts' biases is multiplied and does not vanish even if experts and followers have precise information. Instead, experts' influence depends on the *relative* precision of the information of experts and followers. Consider the effect of a shift in expert bias in favor of action A. Starting from the original equilibrium value of θ^*, this shift generates more expert choice of A in the small neighborhood of θ^* in which experts are uncertain of the coordination outcome. If experts are better informed than followers, followers do not know when experts cannot predict the coordination outcome, and followers cannot factor out the effect of experts' bias. As a result, the increase in expert choice of A in turn leads to followers choosing A more often in states close to θ^*, thereby increasing the pivotal threshold. Because of the coordination motive, the increase in the threshold leads to further increases in the number of agents choosing A, repeatedly multiplying the effect. No matter how small is the direct social learning effect, the desire to coordinate makes the overall effect non-vanishing.

In our model, followers know only the distribution of expert preferences, not the preferences of any particular expert. If followers have perfect knowledge of each expert's bias, then our results do not hold. In this case, followers who observe conflicting choices from independent experts deduce that the state is close to the pivotal one, and are able to correct for experts' biases. If, however, followers observe only a noisy signal of each expert's preference, then results similar to ours continue to hold. As the experts become increasingly informed, followers again effectively neglect those states in which the experts are uncertain about the coordination outcome, believing that conflicting expert actions are more likely to be the result of the presence of partisan experts in the sample. Consequently, the outcome depends on the experts' biases.

The assumption that experts' choices are privately observed by followers is not essential for our results. If instead all followers observe the actions of the same n experts (drawn at random from the continuum of experts), then the equilibrium is again unique and exhibits the same features as in the private case. Moreover, although the equilibria in the two cases involve different thresholds, they converge to the same limit as n grows large. This strongly suggests that experts can also exert influence over the outcome in intermediate cases where the action of a given expert may be observed by many but not all followers (as is natural for marketing or political campaigns). Note that drawing the observed experts at random from a continuum precludes any signaling motive on the part of the experts. We conjecture that incorporating such a motive would only strengthen the influence of experts.[11]

A Proofs

Lemma 4. *Let u denote the p.d.f. of θ. For all $x \in [\sigma_f, 1-\sigma_f]$, $u\left(\theta | x_f^i = x\right) = g\left(\frac{x-\theta}{\sigma_f}\right)/\sigma_f$ and $\pi_f\left(x, \theta^*, \sigma_f\right) = 1 - G\left(\frac{x-\theta^*}{\sigma_f}\right)$. Likewise, for all $x \in [\sigma_e, 1-\sigma_e]$, $\pi_e\left(x, \theta^*, \sigma_e\right) = 1 - F\left(\frac{x-\theta^*}{\sigma_e}\right)$.*

Proof. Using Bayes' rule, we have

$$\begin{aligned} u\left(\theta | x_f^i = x\right) &= \frac{\frac{1}{\sigma_f} g\left(\frac{x-\theta}{\sigma_f}\right) u(\theta)}{\int_0^1 g\left(\frac{x-\theta}{\sigma_f}\right) u(\theta)\, d\theta} = \frac{\frac{1}{\sigma_f} g\left(\frac{x-\theta}{\sigma_f}\right)}{G\left(\frac{x}{\sigma_f}\right) - G\left(\frac{x-1}{\sigma_f}\right)}, \\ \text{and} \qquad \pi_f\left(x, \theta^*, \sigma_f\right) &= \int_0^{\theta^*} u\left(\theta | x_f^i = x\right) d\theta = \frac{G\left(\frac{x}{\sigma_f}\right) - G\left(\frac{x-\theta^*}{\sigma_f}\right)}{G\left(\frac{x}{\sigma_f}\right) - G\left(\frac{x-1}{\sigma_f}\right)}. \end{aligned}$$

If $x \in [\sigma_f, 1-\sigma_f]$, then $\frac{x}{\sigma_f} \geq 1$ and $\frac{x-1}{\sigma_f} \leq -1$, so $G\left(\frac{x}{\sigma_f}\right) = 1$ and $G\left(\frac{x-1}{\sigma_f}\right) = 0$, which establishes the expressions for $u\left(\theta | x_f^i = x\right)$ and π_f. The proof for π_e is similar. □

Proof of Lemma 1. If $\theta_0^* \in \left[s_f^A, 1 - s_f^B\right]$, then conditional on $\theta = \theta_0^*$ all realized signals x_f^i are in $\left[s_f^A - \sigma_f, 1 - s_f^B + \sigma_f\right]$, which is included in $[\sigma_f, 1-\sigma_f]$ since we assume that $\sigma_e + \sigma_f \leq \frac{1}{2}\min\left(s_f^A, s_f^B\right)$. Using Lemma 4, we have

$$\begin{aligned} \Pr\left(\pi_f\left(x_f^i, \theta_0^*, \sigma_f\right) < p \mid \theta_0^*\right) &= \Pr\left(1 - G\left(\frac{x_f^i - \theta_0^*}{\sigma_f}\right) < p \mid \theta_0^*\right) \\ &= \Pr\left(1 - G\left(\varepsilon_f^i\right) < p\right) \\ &= 1 - G\left(\left(G^{-1}(1-p)\right)\right. \\ &= p, \end{aligned}$$

as needed. □

[11] Proposition 7 of Corsetti, Dasgupta, Morris, and Shin (2004) pertains to a model closely related to a variant of our model with one expert who has a signaling motive. In their setting, the expert exerts a large influence over the outcome.

Proof of Proposition 1. Using Bayes' rule,

$$p_f\left(x_f^i,\lambda^i,\theta_n^*,\sigma_e,\sigma_f\right)=\frac{\int_0^{\theta_n^*}\Pr\left(\lambda^i|\theta'\right)u\left(\theta'|x_f^i\right)d\theta'}{\int_0^1\Pr\left(\lambda^i|\theta'\right)u\left(\theta'|x_f^i\right)d\theta'}.$$

Note that $\theta_n^*\in\left[s_f^A,1-s_f^B\right]$. Hence in state θ_n^*, all realized signals x_f^i are in $\left[s_f^A-\sigma_f,1-s_f^B+\sigma_f\right]$, which is contained in $[\sigma_f,1-\sigma_f]$ since $\sigma_e+\sigma_f\leq\frac{1}{2}\min\left(s_f^A,s_f^B\right)$. Therefore, Lemma 4 implies that $u\left(\theta|x_f^i\right)=\frac{1}{\sigma_f}g\left(\frac{x_f^i-\theta}{\sigma_f}\right)$. Since $x_f^i=\theta_n^*+\sigma_f\varepsilon_f^i$ in the pivotal state, we have that $u\left(\theta|x_f^i\right)=\frac{1}{\sigma_f}g\left(\varepsilon_f^i+\frac{\theta_n^*-\theta}{\sigma_f}\right)$, and

$$p_f\left(\theta_n^*+\sigma_f\varepsilon_f^i,\lambda^i,\theta_n^*,\sigma_e,\sigma_f\right)=\frac{\int_{\theta_n^*-2\sigma_f}^{\theta_n^*}\Pr\left(\lambda^i|\theta'\right)g\left(\varepsilon_f^i+\frac{\theta_n^*-\theta'}{\sigma_f}\right)d\theta'}{\int_{\theta_n^*-2\sigma_f}^{\theta_n^*+2\sigma_f}\Pr\left(\lambda^i|\theta'\right)g\left(\varepsilon_f^i+\frac{\theta_n^*-\theta'}{\sigma_f}\right)d\theta'}$$

since g has support on $[-1,1]$.

When $\theta\in[\theta_n^*-2\sigma_f,\theta_n^*+2\sigma_f]$, all realizations of x_e^i are in $\left[s_f^A-2\sigma_f-\sigma_e,1-s_f^B+2\sigma_f+\sigma_e\right]$, which is contained in $[\sigma_e,1-\sigma_e]$ since $\sigma_e+\sigma_f\leq\frac{1}{2}\min\left(s_f^A,s_f^B\right)$. So Lemma 4 and (3) imply that for all $\theta\in[\theta_n^*-2\sigma_f,\theta_n^*+2\sigma_f]$,

$$\begin{aligned}l\left(\theta,\theta_n^*,\sigma_e\right)&=s_e^A+\left(1-s_e^A-s_e^B\right)\Pr\left(1-F\left(\frac{\theta+\sigma_e\varepsilon_e^j-\theta_n^*}{\sigma_e}\right)>p_e^j\right)\\&=s_e^A+\left(1-s_e^A-s_e^B\right)\Pr\left(\varepsilon_e^j<F^{-1}\left(1-p_e^j\right)+\frac{\theta_n^*-\theta}{\sigma_e}\right),\end{aligned}$$

where j is a randomly chosen independent expert. Therefore, $l\left(\theta,\theta_n^*,\sigma_e\right)$ depends on θ, θ_n, and σ_e only through the value of $\frac{\theta_n^*-\theta}{\sigma_e}$; accordingly, let $\hat{l}$ be such that $\hat{l}\left(\frac{\theta_n^*-\theta}{\sigma_e}\right)\equiv l(\theta,\theta_n^*,\sigma_e)$. Using the transformation $\Delta=\frac{\theta_n^*-\theta'}{\sigma_f}$ and the fact that, conditional on θ, λ^i is distributed according to the Binomial distribution $B\left(n,\hat{l}\left(\frac{\theta_n^*-\theta}{\sigma_e}\right)\right)$, we have that

$$\begin{aligned}p_f\left(\theta_n^*+\sigma_f\varepsilon_f^i,\lambda^i,\theta_n^*,\sigma_e,\sigma_f\right)&=\frac{\int_{\theta_n^*-2\sigma_f}^{\theta_n^*}\binom{n}{\lambda^i}\hat{l}\left(\frac{\theta_n^*-\theta'}{\sigma_e}\right)^{\lambda^i}\left(1-\hat{l}\left(\frac{\theta_n^*-\theta'}{\sigma_e}\right)\right)^{n-\lambda^i}g\left(\varepsilon_f^i+\frac{\theta_n^*-\theta'}{\sigma_f}\right)d\theta'}{\int_{\theta_n^*-2\sigma_f}^{\theta_n^*+2\sigma_f}\binom{n}{\lambda^i}\hat{l}\left(\frac{\theta_n^*-\theta'}{\sigma_e}\right)^{\lambda^i}\left(1-\hat{l}\left(\frac{\theta_n^*-\theta'}{\sigma_e}\right)\right)^{n-\lambda^i}g\left(\varepsilon_f^i+\frac{\theta_n^*-\theta'}{\sigma_f}\right)d\theta'}\\&=\frac{\int_0^2\hat{l}\left(\frac{\sigma_f}{\sigma_e}\Delta\right)^{\lambda^i}\left(1-\hat{l}\left(\frac{\sigma_f}{\sigma_e}\Delta\right)\right)^{n-\lambda^i}g(\varepsilon_f^i+\Delta)d\Delta}{\int_{-2}^2\hat{l}\left(\frac{\sigma_f}{\sigma_e}\Delta\right)^{\lambda^i}\left(1-\hat{l}\left(\frac{\sigma_f}{\sigma_e}\Delta\right)\right)^{n-\lambda^i}g(\varepsilon_f^i+\Delta)d\Delta}.\end{aligned}\tag{6}$$

In particular, the last expression does not depend on θ_n^* and depends on the scaling parameters only through their ratio $\frac{\sigma_f}{\sigma_e}$. Since $v(\theta_n^*, \theta_n^*, \sigma_e, \sigma_f) = \Pr\left(p_f\left(x_f^i, \lambda^i, \theta_n^*, \sigma_e, \sigma_f\right) > p_f^i \mid \theta_n^*\right)$, and the distribution of λ^i conditional on θ_n^* does not depend on θ_n^*, equation (6) implies that $v(\theta_n^*, \theta_n^*, \sigma_e, \sigma_f)$ is independent of θ_n^*, establishing that the equilibrium condition (5) has a unique solution. Moreover, since (6) depends on the scaling parameters only through their ratio and the distribution of λ^i conditional on θ_n^* does not depend on the scaling parameters, the solution of (5) similarly depends on the scaling parameters only through their ratio. □

Proof of Proposition 2. Letting $\frac{\sigma_f}{\sigma_e} \to 0$ in (6), we have that conditional on $\theta = \theta_n^*$,

$$p_f\left(\theta_n^* + \sigma_f \varepsilon_f^i, \lambda^i, \theta_n^*, \sigma_e, \sigma_f\right) \to \frac{\int_0^2 g(\varepsilon_f^i + \Delta)d\Delta}{\int_{-2}^2 g(\varepsilon_f^i + \Delta)d\Delta} = 1 - G\left(\varepsilon_f^i\right) = \pi_f\left(\theta_n^* + \sigma_f \varepsilon_f^i, \theta_0^*, \sigma_f\right)$$

(recalling that $\pi_f\left(x_f^i, \theta_0^*, \sigma_f\right)$ is the probability a follower receiving signal x_f^i and observing no experts assigns to the coordination outcome being A when the pivotal state is θ_0^*). Therefore, the pivotal conditions in the settings with and without experts coincide as $\frac{\sigma_f}{\sigma_e} \to 0$. □

Proof of Corollary 2. Rearranging the expression for the threshold in Proposition 3 yields

$$\theta_n^{**} = s_f^A + \left(1 - s_f^A - s_f^B\right) \Pr\left(\left(\frac{\pi}{1-\pi}\frac{1-p_f^i}{p_f^i}\right)^{1/n} > \left(\frac{s_e^A}{1-s_e^B}\right)^{\lambda/n}\left(\frac{1-s_e^A}{s_e^B}\right)^{1-\lambda/n}\right), \qquad (7)$$

where π, λ, and i are independent random variables with $\pi \sim U[0,1]$, $\lambda \sim B\left(n, s_e^A + \left(1 - s_e^A - s_e^B\right)(1-p_e)\right)$, and i a randomly chosen independent follower. The left-hand side of the inequality in (7) converges in probability to 1. Since λ/n converges in probability to $s_e^A + \left(1 - s_e^A - s_e^B\right)(1-p_e)$ the right-hand side of the inequality converges in probability to

$$\rho_e := \left(\frac{s_e^A}{1-s_e^B}\right)^{s_e^A + \left(1-s_e^A-s_e^B\right)(1-p_e)} \left(\frac{1-s_e^A}{s_e^B}\right)^{s_e^B + \left(1-s_e^A-s_e^B\right)p_e}.$$

Therefore, if $\rho_e > 1$ then $\lim_{n\to\infty} \theta_n^{**} = s_f^A$. If $\rho_e < 1$ then $\lim_{n\to\infty} \theta_n^{**} = 1 - s_f^B$. The result follows since ρ_e is a monotone function of p_e. □

References

Acemoglu, D. and J. Robinson (2000). Why did the west extend the franchise? democracy, inequality, and growth in historical perspective. *Quarterly Journal of Economics 115*(4), 1167–1199.

Acemoglu, D. and J. Robinson (2001). A theory of political transitions. *American Economic Review*, 938–963.

Acemoglu, D. and J. Robinson (2006). *Economic Origins of Dictatorship and Democracy*. Cambridge University Press.

Aral, S. (2011). Identifying social influence: A comment on opinion leadership and social contagion in new product diffusion. *Marketing Science 30*(2), 217–223.

Arndt, J. (1967). Role of product-related conversations in the diffusion of a new product. *Journal of Marketing Research*, 291–295.

Baumgarten, S. (1975). The innovative communicator in the diffusion process. *Journal of Marketing Research*, 12–18.

Beaumont, P. (2011). The truth about twitter, facebook and the uprisings in the arab world. *The Guardian*.

Bueno De Mesquita, E. (2010). Regime change and revolutionary entrepreneurs. *American Political Science Review 104*(3), 446–466.

Corsetti, G., A. Dasgupta, S. Morris, and H. Shin (2004). Does one Soros make a difference? A theory of currency crises with large and small traders. *Review of Economic Studies 71*(1), 87–113.

DeNardo, J. (1985). *Power in Numbers: The Political Strategy of Protest and Rebellion*. Princeton University Press.

Dewan, T. and D. Myatt (2007). Leading the party: Coordination, direction, and communication. *American Political Science Review 101*(4), 827.

Dewan, T. and D. Myatt (2008). The qualities of leadership: Direction, communication, and obfuscation. *American Political Science Review 102*(3), 351–368.

Dewan, T. and D. Myatt (2012). On the rhetorical strategies of leaders: Speaking clearly, standing back, and stepping down. *Journal of Theoretical Politics 24*(4), 431–460.

Edmond, C. (2011). Information manipulation, coordination, and regime change. Technical report, National Bureau of Economic Research.

Ekmekci, M. (2009). Manipulation through political endorsements. *Journal of Economic Theory 144*(3), 1227–1248.

Ellis, C. and J. Fender (2011). Information cascades and revolutionary regime transitions. *Economic Journal 121*(553), 763–792.

Faria, J. and D. G. Arce (2005). Terror support and recruitment. *Defence and Peace Economics 16*(4), 263–273.

Feick, L. and L. Price (1987). The market maven: A diffuser of marketplace information. *Journal of Marketing*, 83–97.

Gause, F. (2011). Why Middle East Studies missed the Arab spring: The myth of authoritarian stability. *Foreign Affairs 90*, 81.

Gilly, M., J. Graham, M. Wolfinbarger, and L. Yale (1998). A dyadic study of interpersonal information search. *Journal of the Academy of Marketing Science 26*(2), 83–100.

Gladwell, M. (2000). *The Tipping Point: How Little Things Can Make a Big Difference*. Little, Brown.

Goldenberg, J., S. Han, D. Lehmann, and J. Hong (2009). The role of hubs in the adoption processes. *Journal of Marketing 73*(2).

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 1420–1443.

Guimaraes, B. and S. Morris (2007). Risk and wealth in a model of self-fulfilling currency attacks. *Journal of Monetary Economics 54*(8), 2205–2230.

Honwana, A. (2011). Youth and the Tunisian revolution. Conflict Prevention and Peace Forum policy paper.

Intriligator, M. and D. Brito (1988). A predator-prey model of guerrilla warfare. *Synthese 76*(2), 235–244.

Irving, D. (1981). *Uprising!: One Nation's Nightmare: Hungary 1956*. Hodder and Stoughton.

Katz, E. and P. Lazarsfeld (1955). *Personal Influence: The Part Played by People in the Flow of Mass Communications*. Free Press.

Katz, M. and C. Shapiro (1985). Network externalities, competition, and compatibility. *American Economic Review 75*(3), 424–440.

Kuran, T. (1989). Sparks and prairie fires: A theory of unanticipated political revolution. *Public Choice 61*(1), 41–74.

Kuran, T. (1991). Now out of never: The element of surprise in the East European revolution of 1989. *World Politics 44*(1), 7–48.

Lizzeri, A. and N. Persico (2004). Why did the elites extend the suffrage? Democracy and the scope of government, with an application to Britain's 'Age of Reform'. *Quarterly Journal of Economics 119*(2), 707–765.

Lohmann, S. (1993). A signaling model of informative and manipulative political action. *American Political Science Review*, 319–333.

Lohmann, S. (1994). The dynamics of informational cascades. *World Politics 47*(1), 42–101.

Lohmann, S. (1995). A signaling model of competitive political pressures. *Economics & Politics 7*(3), 181–206.

Lohmann, S. (2000). Collective action cascades: An informational rationale for the power in numbers. *Journal of Economic Surveys 14*(5), 655–684.

Malewski, J. (2011). An ongoing revolution. *International Viewpoint Online Magazine IV435*.

McAdam, D. (1986). Recruitment to high-risk activism: The case of Freedom Summer. *American Journal of Sociology*, 64–90.

McAdam, D. and R. Paulsen (1993). Specifying the relationship between social ties and activism. *American Journal of Sociology*, 640–667.

Merton, R. (1957). Patterns of influence: Local and cosmopolitan influentials. *Social Theory and Social Structure*, 387–420.

Opp, K. and C. Gern (1993). Dissident groups, personal networks, and spontaneous cooperation: The East German revolution of 1989. *American Sociological Review*, 659–680.

Price, L. and L. Feick (1984). The role of interpersonal sources in external search: An informational perspective. *Advances in Consumer Research 11*(1), 250–255.

Schelling, T. (1960). *The Strategy of Conflict.* Harvard University Press.

Skocpol, T. (1979). *States and Social Revolutions.* Cambridge University Press.

Snyder, D. and C. Tilly (1972). Hardship and collective violence in france, 1830 to 1960. *American Sociological Review*, 520–532.

Steiner, J. (2006). Coordination in a mobile world. *CERGE-EI Working Paper 295.*

Tilly, C., L. Tilly, and R. Tilly (1975). *The Rebellious Century: 1830–1930.* Harvard University Press.

Watts, D. (2007). Challenging the influentials hypothesis. *WOMMA Measuring Word of Mouth 3*(4), 201–211.

Watts, D. and P. Dodds (2007). Influentials, networks, and public opinion formation. *Journal of Consumer Research 34*(4), 441–458.

Weimann, G. (1991). The influentials: Back to the concept of opinion leaders? *Public Opinion Quarterly 55*(2), 267–279.