



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Influential opinion leaders

Citation for published version:

Loeper, A, Steiner, J & Stewart, C 2014, 'Influential opinion leaders', *The Economic Journal*, vol. 124, no. 581, pp. 1147-1167. <https://doi.org/10.1111/eoj.12100>

Digital Object Identifier (DOI):

[10.1111/eoj.12100](https://doi.org/10.1111/eoj.12100)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

The Economic Journal

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Influential Opinion Leaders*

Antoine Loeper[†]

Jakub Steiner[‡]

Universidad Carlos III de Madrid

The University of Edinburgh and CERGE-EI

Colin Stewart[§]

University of Toronto

September 24, 2013

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. Expert influence is fully decentralised in the sense that each individual expert has a negligible impact. The bias in favour of experts results from a social learning effect that is multiplied through a coordination motive. We apply our results to the onset of social movements and to the diffusion of products with network externalities.

When a large group of agents seek to coordinate their behaviour 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 behaviour? We show that the choices of expert early movers can have a large effect on outcomes, biasing the results toward their own preferences. The effect arises even though our model features Bayesian decision-makers who

*We thank David Austen-Smith, Mehmet Ekmekci, Yosh Halberstam, Stephen Morris, Alessandro Pavan, Carolyn Pitchik, and Matt Turner for helpful comments. Regina Tukhbatullina provided excellent research assistance. Loeper acknowledges the financial support from grant ECO2010-19596 from the Spanish "Ministerio de Ciencia e Innovación". Steiner was supported by Purkyne fellowship of the Czech Academy of Sciences and by GACR grant 13-34759S. Stewart is grateful to SSHRC for financial support of this research.

[†]email: aloeper@eco.uc3m.es

[‡]email: jakub.steiner@cerge-ei.cz

[§]email: colin.stewart@utoronto.ca

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. Empirically, knowing an activist involved in a social movement is one of the main determinants of mobilisation (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. Expert influence can arise in models of social learning in which a follower observes expert choices, but the effect is bound to be small if the follower is well informed. In our model, a coordination motive multiplies this social learning effect leading to a sizeable total effect even when the direct social learning effect is vanishingly small. This multiplier effect provides a novel channel through which experts influence mass behaviour.

Agents in our model choose one of two actions, A or B . The coordination outcome is A if A is chosen by sufficiently many agents, and is B otherwise. Each agent has an intrinsic bias for a particular action, which he trades off against his preference for being aligned with the coordination outcome. There are two kinds of agents—experts and followers—with a continuum of each. In the first stage of the game, experts simultaneously choose actions. In the second stage, followers observe a private sample of experts’ actions and simultaneously choose their actions. The coordination outcome is determined only by the choices of the followers, capturing the idea that experts form a small group having negligible direct influence. In equilibrium, however, the coordination outcome depends on a combination of both the experts’ and the followers’ biases, enabling us to quantify the expert influence.

Experts affect the coordination outcome through the information revealed by their actions. Each agent is endowed with private information about the fraction of followers who must choose a given action for it to become the coordination outcome. Because this fraction is unknown, agents face uncertainty about the coordination outcome in equilibrium. Since experts’ actions depend on their private information, observing their actions provides information to followers about what the experts expect the coordination outcome to be. Followers account for the distribution of experts’ biases when evaluating their observations, despite which the interplay between experts and followers can skew the outcome in the experts’ favour. The experts become opinion leaders.

Experts’ influence depends on their biases in a non-monotone way. We distinguish between two degrees of bias. Moderately biased experts, if certain of the coordination outcome, prefer to align their action with the outcome even if it goes against their personal bias. In contrast, the actions of partisan experts are always driven by their personal biases, which are strong enough to overwhelm any desire to coordinate. As the share of experts who are moderately biased in favour of an action increases, followers become more likely to coordinate on that action. However, an increase in the partisan support for an action may have the opposite effect: followers can become more likely to coordinate on the other action. The difference between the effects of the two kinds of experts arises because partisan experts’ actions, unlike those of the moderate experts, do not convey information.

Expert influence arises when experts are well informed relative to followers. In that case, experts are rarely uncertain about the coordination outcome, making it very unlikely that the biases of moderate experts affect their actions. Consequently, followers view the sample of expert action

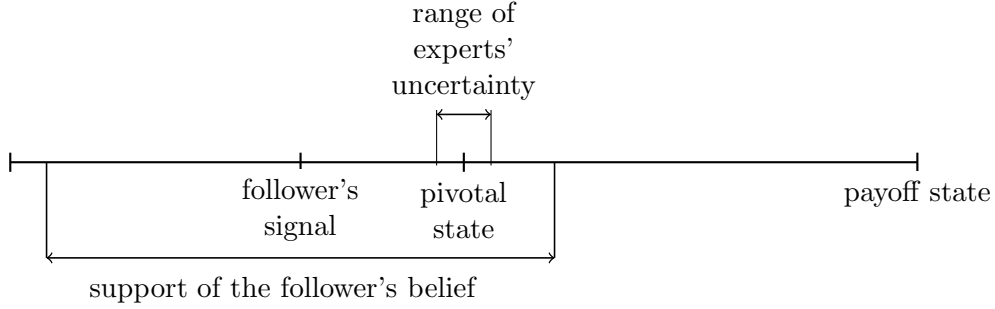


Figure 1: The coordination outcome is A below and B above the pivotal state. When the realised state is close to the pivotal one, (some) experts receive inconclusive signals and cannot predict the outcome with certainty. Each follower receives relatively imprecise signal, and thus his belief about the state has relatively large support. As the range of experts' uncertainty is small, conditional on the follower's information, it is very likely that the expert knows the outcome. If, incidentally, the realised state is the pivotal one, the followers overestimate experts' ability to predict the outcome.

as directly informative about the coordination outcome. Followers believe that each action they observe is either from a moderate expert who is matching the coordination outcome or from a partisan expert; they attribute any disagreement among experts to partisanship rather than expert uncertainty (see Figure 1). As a result, followers take experts' actions at face value even in those rare states in which, because they are uncertain of the coordination outcome, experts' biases do affect their actions.

Given that moderate experts' biases rarely affect their actions, how can their biases have a large impact on outcomes? The reason is that equilibrium behaviour is determined precisely by what happens in those rare states in which experts are uncertain about the outcome. The equilibrium features a pivotal state on one side of which the coordination outcome is A , and on the other side of which it is B . In the pivotal state, the fraction of followers choosing each action must be exactly the one needed for that action to be the coordination outcome. If, for example, all moderate experts have a reasonably strong bias toward action A , then, in the pivotal state, most moderate experts will choose action A ; B will be chosen only by the small fraction whose signals strongly suggest that B will be the coordination outcome. Most followers, then, will observe a preponderance of experts choosing action A , suggesting that A is likely to be the coordination outcome. Hence most followers in turn choose A . Therefore, the pivotal state must be one that strongly favours B in the sense that many followers must choose A for A to be the coordination outcome. Hence, when

experts are biased in favour of A , A is the coordination outcome in every state that does not favour B so strongly.

Our model focuses on the interaction between strategic uncertainty and social learning. The global games literature provides several insights on the role of social learning in coordination. If social learning is public then the observation of early actions correlates the beliefs of late movers. This correlation can lead to equilibrium multiplicity as in Angeletos *et al.* (2007), or to a disproportionate influence of one large opinion leader as in Corsetti *et al.* (2004).¹ In our model, social learning alters the form of strategic uncertainty in way that gives undue influence to a decentralised group.

1 Applications

This section discusses the assumptions and predictions of our model in the contexts of social movements and of the diffusion of goods with network externalities.

1.1 *Spontaneous Regime Changes*

Social movements are prominent examples of large scale coordination problems in which social learning can play an important role. In the context of a popular uprising, our experts can be thought of as visible citizens with superior information about the strength or popularity of the status quo regime. These elite members choose whether to participate in the initial phase of a protest. Followers—the mass population—decide whether to join the protest after observing the initial mobilisation. Protesters can gain privileged status if the regime is replaced, but run the risk of being punished if it stays in place. The fate of the regime depends in turn on the mobilisation rate (DeNardo, 1985), thus creating a coordination motive. Since these gains and risks may not be the same for experts and followers, their incentives may be misaligned. The decentralised interactions in our model fit the nature of revolts. Collusion among experts is difficult while the status quo regime is in power, and the regime may prevent public reporting of the initial mobilisation; followers typically learn about it through interpersonal ties.²

¹See Dasgupta (2007) for the pioneering study of private social learning in global games.

²Empirically, 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 Leipzig demonstrations of 1989, which triggered the East German uprising and eventually led to the fall of the Berlin wall, provide a natural case with which to confront the assumptions and predictions of our model. According to Lohmann, '[t]he first three demonstrations in Leipzig were characterised by relatively low turnout, yet they played a critical role in triggering the protest throughout the GDR' (Lohmann, 1994, p.87). The regime prevented the participants and the press from sending public signals about the first revolts. The participants in the first demonstrations had a 'keen sense of the mood and strength of the regime' (Francisco, 1993, p.667). They were less attached to socialist ideals and less fearful of unification with West Germany than were the rest of the population. Despite this bias, and despite the 'dispersed nature of the information about the GDR's precarious economic and political situation' (Lohmann, 1994, p.43), these small-scale, peaceful protests were soon followed by 'spontaneous, unorganised mass demonstrations' (Lohmann, 1994, p.69).

Expert influence in our model is driven by private social learning that is completely decentralised. Consistent with these features, Lohmann notes that 'the successful Leipzig Monday demonstrations were dominated by unorganised individuals... who made their participation decisions in a decentralised way' (Lohmann, 1994, p.89). In addition, private social learning seems to have been the main channel for the spread of mobilisation: 'social embeddedness and personal networks... influenced individual participation decisions' but 'political entrepreneurship and organisation played a secondary role' (Lohmann, 1994, pp.68–9). In our model, experts' biases influence outcomes because followers believe experts are well informed even when experts are uncertain about the outcome of the movement. The importance of such cases is consistent with the widespread view that 'the revolution came as a surprise even for leading dissidents' (Kuran, 1991, p.8).

The Hungarian revolt of 1956 and the Jasmine revolution in Tunisia also began as small and unorganised demonstrations (see, e.g., Kuran, 1991, 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 organised extremists.³ 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 poor economic management of President Ben Ali. Political organisations

³In 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).

played little role in these events (Malewski, 2011)). The mass media did not report the first demonstrations in Sidi Bouzid, but the population learned about it via private channels and social media.⁴ In both cases, the movement rapidly escalated into a massive uprising that led to the fall of the government.

The literature on social movements has extensively analysed the informational role of groups which are organised, and therefore internalise the impact of information they provide to the broader public (Bueno De Mesquita, 2010; Edmond, 2013). In contrast, in our model, each expert can influence only a negligible fraction of followers, and thus takes the outcome as given, and has no signalling motive. Building on Granovetter’s (1978) threshold model, Kuran (1989) develops a model of revolutionary bandwagons. A key assumption is that citizens’ willingness to mobilise is driven directly by the participation rate, so contagion occurs via a direct externality from the extremists on the moderates. In contrast, in our paper, contagion occurs via an informational channel in which early movers are influential only if their preferences are not too extreme. Most closely related to this paper, Lohmann (1993; 1995; 2000) analyses various versions of a signalling model of unorganised political action. Citizens have dispersed information about the status quo, and can signal that information by taking political action. 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. Because information provision and political participation are public goods, Lohmann’s models predict free-riding, and thus low levels of participation in all stages of the movement.⁵ Hence, Lohmann’s model cannot explain why ‘tiny oppositions mushroomed into crushing majorities’ (Kuran, 1991, p.13).

⁴The 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, 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 mobilisation among experts, instead of a private sample of experts. We conjecture that our results continue to hold in that case: leaders’ actions remain influential, and the likelihood of a mass protest still depends on the experts’ preferences.

⁵Lohmann does not formally characterise the dynamics of mobilisation in the case of a large population. However, it is easy to see that if the number of citizens whose equilibrium action depend on their signal (i.e., nonextremist and nonapathetic citizens) become very large, the probability that each one of them is pivotal vanishes, so their incentive to participate vanishes as well.

1.2 *Opinion Leaders and Viral Marketing*

The adoption of a network technology or a social good (social media, software, movies, fashion) generates network externalities: the value of a product to a each consumer increases with the number of people who use it, thereby creating a coordination motive (Katz and Shapiro, 1985). 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 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 documented numerous examples of product diffusion that are overwhelmingly driven by a small number of individuals, initially referred to as ‘opinion leaders’ by Katz and Lazarsfeld (1955).⁶ Consistently with the assumptions of our model, these influentials have been characterised as visible, well informed individuals, and their influence has been found to be restricted to their interpersonal ties (see, e.g., Price and Feick, 1984, or Watts and Dodds, 2007).

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 model contributes to this debate by quantifying and determining the drivers of opinion leaders’ influence. Our results show that in the case of network goods, leaders can be very influential, and their influence depends primarily on their expertise and visibility. More precisely, what matters is not their degree of expertise, but their relative expertise compared to the followers. 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 *et al.*, 1998). Our model shows that leaders can be influential even if their preferences differ from those of the followers. However, as leaders become too biased, their influence vanishes because their choices are no longer informative to the followers. Hence, the observed correlation between similarity and social influence may simply reflect a concern for

⁶Depending 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 *et al.*, 2009).

	Regime changes	Regime survives
Expert	$1 - p_e$	$-p_e$
Follower	$1 - p_f$	$-p_f$

Table 1: Payoffs for protesting against the current regime.

informational relevance, rather than homophily or ease of communication (Price and Feick, 1984).

2 Example

We begin our analysis with an example consisting of a simple regime change game modelling the onset of a social movement. There are two groups of agents: a continuum of experts and a continuum of followers, where the latter group comprise the great mass of the population. Members of each group choose whether to join or abstain from a protest against the current regime. The outcome of the social movement depends only on the followers' actions. More precisely, regime change occurs if the fraction of followers joining the protest is at least θ , where θ is a stochastic parameter that captures the strength of the regime. Experts have no direct impact on the regime; their influence arises only indirectly through the information their actions reveal to followers.

Each agent's payoff depends on her own action, the fate of the regime, and whether she is an expert or a follower. The payoff for abstaining from protest is normalised to 0; the payoff for protesting is given in Table 1. The payoffs encompass two possibly conflicting goals. While agents prefer to join a successful protest and abstain from a failing one, they also have a bias for or against the current regime that differs across the two groups. Joining the protest is a best response for an expert if and only if she assigns probability at least p_e to the protest succeeding, and similarly for followers with probability p_f . A stronger bias against the regime corresponds to lower values of p_e or p_f .

In order to capture uncertainty about the actions of others, we endow agents with private information. First, the state θ is drawn from a uniform distribution on $[-1, 2]$. Then each expert j receives a private signal $x_e^j = \theta + \sigma_e \varepsilon_e^j$, and each follower i receives a private signal $x_f^i = \theta + \varepsilon_f^i$. The error terms ε_e^j and ε_f^i are drawn from a uniform distribution on $[-1, 1]$, are independent across agents, and are independent of θ . We consider three values of σ_e : 0.01, 0.5, and 100. Contrasting these three cases highlights the role of the precision of experts' information relative to that of

followers.

The timing is as follows. After signals have been observed, first experts simultaneously choose their actions. In addition to his private signal x_f^i , each follower observes the action of a random, uniformly chosen expert, with the draws being independent across followers. Followers then simultaneously choose their own actions.

A monotone equilibrium of this game consists of four thresholds: a pivotal state θ^* together with three threshold signals x_e^* , x_{f0}^* , and x_{f1}^* . The regime change occurs precisely when $\theta \leq \theta^*$. Each agent joins the protest when her private signal is below her threshold signal: experts' threshold is x_e^* , that of followers who observe an expert abstaining is x_{f0}^* , and that of followers who observe an expert protesting is x_{f1}^* .

The equilibrium is determined by four conditions involving the four thresholds. First, an agent receiving her threshold signal must be indifferent between protesting and abstaining; that is,

$$\begin{aligned} \Pr(\theta \leq \theta^* \mid x_e^j = x_e^*) &= p_e, \\ \Pr(\theta \leq \theta^* \mid x_f^i = x_{f0}^* \text{ and } x_e^j > x_e^*) &= p_f, \\ \text{and } \Pr(\theta \leq \theta^* \mid x_f^i = x_{f1}^* \text{ and } x_e^j \leq x_e^*) &= p_f. \end{aligned}$$

Second, in the pivotal state, the fraction of followers who protest must be exactly that needed for regime change to occur. Thus we have the *pivotal condition*

$$\Pr\left(\left(x_f^i \leq x_{f0}^* \text{ and } x_e^j > x_e^*\right) \text{ or } \left(x_f^i \leq x_{f1}^* \text{ and } x_e^j \leq x_e^*\right) \mid \theta^*\right) = \theta^*.$$

Figure 2 depicts the equilibrium pivotal state θ^* as a function of the experts' bias p_e when $p_f = 0.8$. To interpret the figure, notice that the ex ante probability of regime change increases linearly with θ^* . When $\sigma_e = 100$ —meaning that the experts are practically uninformed—the pivotal state is approximately independent of the experts' bias, and is relatively low. In this case, the followers view the experts' actions as essentially uninformative of the regime's fate, they effectively ignore them, and they coordinate on protesting only when the regime is quite weak. When $\sigma_e = 0.5$ —meaning that experts are moderately informed—experts influence the outcome, but the pivotal state is not monotone in the experts' bias. As experts become more biased against




Figure 2: The pivotal state θ^* as a function of experts' bias p_e against action A . Followers have a fairly weak incentive to protest, with $p_f = 0.8$.

the current regime, two countervailing effects occur. On the one hand, more followers observe experts protesting, which, all else equal, makes them more likely to protest. On the other hand, since experts choose to protest more often, observing an expert protesting is less strongly indicative of regime change, making followers less likely to follow the expert's lead. The combination of these two effects generates the non-monotonicity. In addition, in this case, experts' influence on followers is limited in the sense that, irrespective of the experts' bias, the followers coordinate on protesting only when the regime is weak.

Experts exert a strong influence when they are much more informed than are followers. When $\sigma_e = 0.01$, the ex ante probability of regime change increases with the experts' bias against the current regime. In particular, when experts are strongly biased against the current regime, followers coordinate on protesting in the vast majority of states in which protesting could overturn the regime. What leads the followers to coordinate on protesting even though they have a lot to lose and little to gain from doing so? If p_e is close to 0, experts only abstain from protesting if they are virtually certain that the regime will stand. However, from the perspective of a follower who observes an expert protesting, since experts are very well informed, it is much more likely than not that the expert is quite certain that regime change will occur. Most followers therefore mimic the experts' choice to protest, ignoring the experts' bias, which tilts the equilibrium toward regime change.⁷

More generally, the key to understanding the effect lies in the distinction between two conditional

⁷The only followers who do not protest are those whose own signals strongly suggest that regime change will not occur.

probabilities. Consider a shift in the experts' bias. On the one hand, *conditional on the state being pivotal*, this shift affects the probability that any given expert protests. This occurs because, in the pivotal state, experts have dispersed beliefs about the likelihood of regime change—beliefs that they trade off against their bias. Thus some experts will change their actions as the bias shifts. On the other hand, *conditional on a given follower's signal*, the probability that the shift affects the share of experts who protest vanishes as the experts' signals become arbitrarily precise (i.e. as $\sigma_e \rightarrow 0$). This occurs because, as σ_e vanishes, followers become increasingly certain that experts know the fate of the regime. When experts know the outcome, they align their actions with it regardless of their bias. Thus, conditional on any follower's information, the experts' bias has vanishing impact on the experts' actions.

Both conditional probabilities shape equilibrium behaviour. The former indicates that, in the pivotal state, the experts' bias affects followers' behaviour—behaviour that determines the equilibrium value of the pivotal state itself. The latter indicates that the followers' beliefs do not reflect the experts' bias. Followers view experts' actions as precise predictions of the outcome, rendering their own signals virtually irrelevant. Followers' behaviour in the pivotal state is therefore driven solely by the distribution of experts' actions, reflecting the experts' bias.

While the experts' bias has a monotone effect when it is moderate—that is, for $p_e \in (0, 1)$ —the effect is not monotone if we allow for extreme biases. For example, even though experts have a strong influence when p_e is just above 0, their influence vanishes if p_e falls below 0. When experts become partisans who always choose a particular action, their actions cease to be informative, leading the followers to ignore them.

3 Model

The main model generalises the preferences and information structure from the example; among other generalisations, it allows for followers to observe an arbitrary number of experts' actions. As in the example, there are 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 kinds 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.⁸ Although we normalise the measure of both groups 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 actions of the followers.

Preferences of a follower i are described by $\pi_{AA}^i, \pi_{AB}^i, \pi_{BA}^i, \pi_{BB}^i$, where, π_{XY}^i is his payoff when he plays 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, $\pi_{AA}^i - \pi_{BA}^i > \pi_{AB}^i - \pi_{BB}^i$. Under this assumption, follower i 's best response takes the same form as in the example: he plays A if he believes that the coordination outcome is A with a probability at least $p_f^i = (\pi_{BB}^i - \pi_{AB}^i) / (\pi_{AA}^i - \pi_{AB}^i + \pi_{BB}^i - \pi_{BA}^i)$, and he plays B otherwise. We assume that the bias p_f^i of follower i against action A is an integrable function of i . For all i such that $p_f^i < 0$ or $p_f^i > 1$, A or B , respectively, is a dominant action. For $X \in \{A, B\}$, we denote by s_f^X the fraction of followers for whom X is dominant. We refer to these followers as X -partisans and to nonpartisan followers as independents. In the sequel, we assume that $s_f^A > 0$, $s_f^B > 0$, and $s_f^A + s_f^B < 1$.

Experts' payoffs have the same structure. For an expert j , we again denote by p_e^j the critical belief that characterises his best response, and by s_e^A and s_e^B the fraction of A -partisan and 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 agents. Similarly, if $\theta > 1 - s_f^B$, the coordination outcome is B . We focus on the coordination outcome for $\theta \in [s_f^A, 1 - s_f^B]$, where it depends on the independent followers' actions. In this region, the agents face a coordination problem.

The information structure and timing are as follows. The state θ is drawn from a uniform distribution on $[0, 1]$, and experts and followers receive private signals $x_e^j = \theta + \sigma_e \varepsilon_e^j$, and $x_f^i = \theta + \sigma_f \varepsilon_f^i$, respectively.⁹ The errors ε_e^j and ε_f^i are drawn from continuous distributions F and G , with densities f and g , respectively, and each with support $[-1, 1]$. Errors are independent across agents and independent of θ . After the experts simultaneously choose their actions, each follower observes a random sample of n expert actions and makes a choice. The sample is private and taken with

⁸In case of a tie, the coordination outcome may be chosen arbitrarily.

⁹It is important for our results that the support of the prior distribution of states contains regions with dominant actions for independent agents. These dominance regions arise in the main model because of the presence of partisans. In the example, which has no partisans, dominance regions are assumed directly with a larger support of the distribution of states.

uniform probability over all experts. Followers do not observe the biases or signals of the experts in their sample (see Section 6 for a discussion of this assumption). The special case of our model in which $n = 0$ is equivalent to a standard static regime change game played by the followers (see, e.g., Angeletos *et al.* (2007)). To simplify the exposition, we assume that $\sigma_e + \sigma_f \leq 0.5 \min(s_f^A, s_f^B)$; this condition eliminates the need to consider effects arising at the boundary of the support of θ in the equilibrium analysis.

A strategy for an expert maps each signal x_e^j to an action $a_e^j \in \{A, B\}$. Letting $\lambda^i \in \{0, \dots, n\}$ denote the number of actions A in follower i 's sample, a strategy for a 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 he chooses B and below which he 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 $x' \geq x$ 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 weak perfect Bayesian equilibria in monotone strategies.¹⁰ Any such equilibrium features a pivotal state θ^* such that the coordination outcome is A below and B above θ^* . All parameters of the model, including all distributions, are common knowledge.

4 Equilibrium

As in the example, the equilibrium is determined by the pivotal condition on the followers' aggregate behaviour in the pivotal state θ^* . To understand the followers' decisions, we must start with those of the experts. Recall that expert j plays A if and only if his posterior belief that the coordination outcome will be A is at least p_e^j , and let $l(\theta, \theta^*)$ denote the probability that a randomly selected expert chooses A in state θ given the threshold θ^* .

Returning to the followers, let $p_f(x, \lambda, \theta^*) = \Pr(\theta \leq \theta^* \mid x, \lambda)$ denote the posterior probability that a follower assigns to the coordination outcome A after observing a signal x and a number λ of experts choosing A (given the threshold θ^*). By Bayes' rule, this probability is

$$p_f(x, \lambda, \theta^*) = \frac{\int_{\theta \leq \theta^*} g\left(\frac{x-\theta}{\sigma_f}\right) \Pr(\lambda \mid \theta) d\theta}{\int_{\theta} g\left(\frac{x-\theta}{\sigma_f}\right) \Pr(\lambda \mid \theta) d\theta}, \quad (1)$$

¹⁰The restriction to monotone strategies is without loss of generality. A standard infection argument along the lines of Morris and Shin (2003) can be used to show that there is an essentially unique serially undominated strategy profile.

where the distribution of λ conditional on θ is binomial with n draws and probability parameter $l(\theta, \theta^*)$. The pivotal condition is then

$$\theta^* = s_f^A + (1 - s_f^A - s_f^B) \Pr\left(p_f(x_f^i, \lambda^i, \theta^*) \geq p_f^i \mid \theta^*\right). \quad (2)$$

The pivotal condition (2) with beliefs (1) has a unique solution.

Proposition 1. *The game has a unique monotone equilibrium. Moreover, the equilibrium pivotal state θ^* depends on σ_e and σ_f only through their ratio σ_e/σ_f .*

All proofs are in the Appendix.

Equilibrium uniqueness allows us to analyse comparative statics of the likelihood of coordination on a particular action. In the next section, we consider how experts' and followers' preferences influence the coordination outcome.

5 Experts' influence

We now examine the influence of experts as a function of their informational advantage σ_f/σ_e , their degree of partisanship s_e^A and s_e^B , and their visibility n .

Before we get to the main result, we extend the finding from the example that experts lack influence if they are poorly informed relative to followers. In that case, experts' actions are uninformative and followers ignore them when forming their beliefs.

Let p_f denote the average bias in the population of independent followers.

Proposition 2. *As $\sigma_f/\sigma_e \rightarrow 0^+$, the pivotal state θ^* converges to*

$$\theta_0^* = s_f^A + (1 - s_f^A - s_f^B) (1 - p_f),$$

which is the pivotal state of the game without experts (i.e. the case $n = 0$). In particular, the pivotal state is independent of experts' preferences.

In the absence of experts, or when they are poorly informed, the equilibrium outcome aggregates the preferences of followers in a naturally monotone way. In the context of regime change, the ex ante probability of a revolution increases with the population of partisan opponents of the regime

and with the average level of dissatisfaction among nonpartisan citizens, and decreases with the population of partisan supporters of the regime. The channel through which the preferences affect the outcome is best understood through the pivotal condition. The pivotal state is determined by the best responses *in the pivotal state* itself, in which the agents receive inconclusive signals making them unsure about the coordination outcome, thereby suppressing the significance of the coordination motive relative to the bias. Consequently, the aggregate action is monotone in the average bias.

We now return to the case in which experts are well informed. Let θ_n^* denote the limit (given n) of the pivotal states θ^* as $\sigma_f/\sigma_e \rightarrow +\infty$, and let p_e denote the average bias of the independent experts. The following proposition shows that well informed experts influence the outcome, and do so in a very powerful way when they are highly visible.

Proposition 3. *1. The limiting pivotal state θ_n^* is strictly decreasing in the independent experts' average bias p_e .*

2. For each s_e^A and s_e^B , there exists p_e^ , independent of followers' preferences, such that*

$$\lim_{n \rightarrow \infty} \theta_n^* = \begin{cases} s_f^A & \text{if } p_e > p_e^*, \\ 1 - s_f^B & \text{if } p_e < p_e^*. \end{cases}$$

The first part of the proposition states that well informed experts influence the outcome toward their own interest. As independent experts become more biased toward a given outcome, followers become more likely to coordinate on that outcome. In the context of network good adoption, this result implies that the firm can increase the likelihood of widespread adoption by offering discounts to influential early adopters. This strategy is effective as long as the price remains high enough that early adopters' choices still convey information to other prospective buyers.

The second part of the proposition says that when well informed experts are very visible (i.e., when $n \rightarrow \infty$), the followers' preferences cease to impact the outcome. In the limit, expert influence is maximal: the coordination outcome is completely determined by the experts' preferences in every state in which the outcome depends on the choices of independent followers.

Proposition 3 sheds light on repression by regimes facing the threat of a revolution. Although

the literature on the repression of social movements focuses on the effect of repression on large, violent, or organised movements (see, e.g., Cunningham (2003) and the references therein), our results suggest that, to avoid contagion, the status quo regime should also target small groups of well informed, visible activists, even if these groups are loosely organised and not representative of the population as a whole. Millman (2005) documents the active efforts of the British government in fighting civil opposition to World War I. Although ‘dissent, at the beginning, was quite plainly confined to an elite’ (Millman, 2005, p.416), the government feared that ‘the principled objection of the few could become the mass disaffection of the many’ (Millman, 2005, p.414). The government took various measures (including censorship, conscription, and imprisonment) to ‘stop the connection between leadership and the mass’ (Millman, 2005, p.415). Cunningham (2003) analyses covert actions by the FBI against the various groups of the New Left counterculture between 1968 and 1971. His analysis shows that the directorate was ‘acting out of fear of a contagion effect’ (Cunningham, 2003, p.229) and that ‘repression was primarily allocated against visible targets rather than always directly against large, active, or violent groups’ (Cunningham, 2003, p.232).

To explain the main result, we expand on the reasoning used in the example. First consider the experts. In the pivotal state, experts’ posterior beliefs $\pi_e(x_e^j, \theta^*) = \Pr(\theta \leq \theta^* | x_e^j)$ reflect only the noise in their signals rather than useful information about the coordination outcome. As a result, as stated in the following lemma, the beliefs are diffuse in this contingency.

Lemma 1. *Posterior beliefs in the pivotal state are distributed uniformly on $[0, 1]$ regardless of the noise distribution; that is,*

$$\Pr(\pi_e(x_e^j, \theta^*) > p | \theta^*) = 1 - p,$$

for any $p \in [0, 1]$ and any j .

A similar uniform property of posterior beliefs has been used in Steiner (2006) and Guimaraes and Morris (2007), among others.

The lemma implies that, in the pivotal state, the fraction of independent experts who choose action A is $1 - p_e$. Away from the pivotal state, computing the distribution of expert actions is particularly simple in states where every expert’s signal perfectly reveals the coordination outcome, for then every independent expert matches their own action with the outcome. Putting these cases

together gives

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

In particular, (i) in typical states—those outside a $2\sigma_e$ -neighbourhood of θ^* —experts know the outcome, and thus the biases of independent experts do not influence their actions, and (ii) in the pivotal state, experts’ posterior beliefs are random, and thus the share of experts choosing A decreases with their average bias against A .

Both observations are important for the analysis of followers’ behaviour. When σ_e is small relative to σ_f , each follower believes with high probability that the state is typical, and because of (i) effectively neglects the experts’ biases when evaluating their actions. Followers neglect the experts’ biases even in the pivotal state in which, because of (ii), experts’ biases do shape their actions. Although the observed samples of experts typically contain supporters of both actions in the pivotal state, followers do not interpret this as an indication of expert uncertainty; instead, they attribute any disagreement to the presence of partisans. In particular, the followers’ posteriors increase in the number λ of experts’ choices of action A , but do not depend on the biases of the independent experts. Since the equilibrium is determined by the followers’ behaviour 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, leading to the first result in Proposition 3.

The second result in Proposition 3 may be understood as follows. 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 population will coordinate on the action supported by the majority of experts in their sample, and attribute the minority choice of experts to partisans for that action; followers ignore their own private information.¹¹ In this case, $p_e^* = 1/2$, meaning that the outcome favoured by the average independent expert prevails in all states in which followers

¹¹There is a small caveat to this claim. A follower will make use of her private information if she receives a sample in which experts are almost exactly equally divided. However, this event becomes vanishingly unlikely as n grows large.

face 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. If all experts become partisans, their influence vanishes. Formally, as $s_e^A + s_e^B \rightarrow 1$, aggregate expert action $l(\theta, \theta^*) \rightarrow s_e^A$ in all states θ , so the distribution of expert actions is independent of θ . Therefore, the followers' posterior beliefs become independent of the observed expert actions and the pivotal state tends to that of the setting without experts.

Unlike independent experts, partisan experts do not necessarily move the outcome in their preferred direction.

Proposition 4. *The pivotal state θ_n^* tends to $1 - s_f^B$ as $s_e^A \rightarrow 0$ and $n \rightarrow \infty$.*

The proposition implies that as the share of A -partisans vanishes, it becomes more likely that the coordination outcome will be A . 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 conflicting actions in their expert samples. When A -partisans are rare, followers conclude from these conflicting actions 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 .

The qualitative difference between the influence of partisan and nonpartisan experts has consequences for network goods. While Proposition 3 suggests that offering discounts for early adopters of a network good may be an effective marketing strategy even if the late customers understand the biased incentives of the early adopters, the marketing strategy ceases to be effective if the discounts are too large. In that case, once purchasing becomes a dominant action for the early adopters, the strategy may backfire and reduce overall adoption of the good.

Proposition 4 can also explain the fate of the various cycles of the Leipzig demonstrations. As discussed in Lohmann (1994), the participants in the first demonstrations were more dissatisfied with the regime than the rest of the population, but they were not extremists. Their mobilisation successfully triggered a massive protest: participation grew rapidly from several thousand to

hundreds of thousands. In contrast, ‘the organised demonstrations of the third, fourth, and fifth cycles were failures’ (Lohmann, 1994, p.89). Lohmann argues that the participants in the latter cycles had more extreme preferences, and their ‘efforts appear to have been discounted by the people’ (Lohmann, 1994, p.89). These patterns are consistent with our findings that experts can be influential even if their preferences do not coincide with the preferences of the followers, but the actions of partisan experts are not informative, and are thus ignored by the followers.

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 followers only when experts’ behaviour depends on their own biases; that is, only in a $2\sigma_e$ -neighbourhood of θ^* where experts may be uncertain of the optimal action. When experts have very precise information, such contingencies are rare. Likewise, the action chosen by a follower depends on the experts’ behaviour only in a $2\sigma_f$ -neighbourhood of θ^* where the followers may be uncertain of the optimal action. Therefore, in the absence of a coordination motive, the ex ante probability that experts’ biases affect the coordination outcome vanishes as the precision of information of either group increases.

When the pivotal state is determined endogenously by the followers’ behaviour, the effect of expert bias is multiplied and may 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 a shift in expert bias in favour of regime change in the social movement setting. Starting from the original equilibrium value of the pivotal state, this shift generates more expert protests in the small neighbourhood of the pivotal state in which experts are uncertain of the outcome. If experts are relatively well informed, followers cannot factor out the effect of experts’ biases. Thus the increase in expert participation in turn leads to followers protesting more often in states close to the pivotal one, thereby increasing the pivotal state. Because of the coordination

motive, the increase in the pivotal state leads to further increases in the number of protesters, 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 that 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.

Similarly, the assumption that all experts have identical precision of information is not essential for our results. Suppose that the signal precision of some experts diverges but the noise in other experts' signals is non-vanishing. As in our baseline model, the followers overvalue experts' knowledge in the pivotal state and therefore they undervalue the impact of experts' biases when interpreting the observed expert actions. We conjecture that Proposition 3 carries over to this setting. Furthermore, the assumption that the followers know the precision of the experts' signal is also not necessary for our argument; it suffices for the followers to believe there is an upper bound on the noise in experts' signals. Our argument applies to the limit in which this bound vanishes.

The assumption that experts' choices are privately observed by followers can be relaxed. 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 signalling motive on the part of the experts. We conjecture that incorporating such a motive by having finitely many experts would

only strengthen their influence.¹²

A Proofs

Let $\pi_f(x, \theta^*) = \Pr(\theta \leq \theta^* | x_f^i = x)$ and recall that $\pi_e(x, \theta^*) = \Pr(\theta \leq \theta^* | x_e^j = x)$. Note that we suppress from the notation the dependence of π_f on σ_f and of π_e on σ_e .

Lemma 2. *Let u denote the p.d.f. of θ . For all $x \in [\sigma_f, 1 - \sigma_f]$, $u(\theta | x_f^i = x) = g[(x - \theta)/\sigma_f]/\sigma_f$ and $\pi_f(x, \theta^*) = 1 - G[(x - \theta^*)/\sigma_f]$. Likewise, for all $x \in [\sigma_e, 1 - \sigma_e]$, $\pi_e(x, \theta^*) = 1 - F[(x - \theta^*)/\sigma_e]$.*

Proof. Using Bayes' rule, we have

$$u(\theta | x_f^i = x) = \frac{\frac{1}{\sigma_f} g\left(\frac{x-\theta}{\sigma_f}\right) u(\theta)}{\int_0^1 g\left(\frac{x-\theta}{\sigma_f}\right) u(\theta) d\theta} = \frac{\frac{1}{\sigma_f} g\left(\frac{x-\theta}{\sigma_f}\right)}{G\left(\frac{x}{\sigma_f}\right) - G\left(\frac{x-1}{\sigma_f}\right)},$$

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

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

Proof of Proposition 1. Using Bayes' rule, we have

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

Note that $\theta^* \in [s_f^A, 1 - s_f^B]$. Hence in state θ^* , all realised signals x_f^i are in $[s_f^A - \sigma_f, 1 - s_f^B + \sigma_f]$, which is contained in $[\sigma_f, 1 - \sigma_f]$ since $\sigma_e + \sigma_f \leq 0.5 \min(s_f^A, s_f^B)$. Therefore, Lemma 2 implies that $u(\theta | x_f^i) = \frac{1}{\sigma_f} g\left[\frac{(x_f^i - \theta)}{\sigma_f}\right]$. Since $x_f^i = \theta^* + \sigma_f \varepsilon_f^i$ in the pivotal state, we have that $u(\theta | x_f^i) = g\left[\varepsilon_f^i + (\theta^* - \theta)/\sigma_f\right]/\sigma_f$, and

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

¹²Proposition 7 of Corsetti *et al.* (2004) pertains to a model closely related to a variant of our model with one expert who has a signalling motive. In their setting, the expert exerts a large influence over the outcome.

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

When $\theta \in [\theta^* - 2\sigma_f, \theta^* + 2\sigma_f]$, all realisations of x_e^i are in $[s_f^A - 2\sigma_f - \sigma_e, 1 - s_f^B + 2\sigma_f + \sigma_e]$, which is contained in $[\sigma_e, 1 - \sigma_e]$ since $\sigma_e + \sigma_f \leq 0.5 \min(s_f^A, s_f^B)$. So, using Lemma 2, we have for all $\theta \in [\theta^* - 2\sigma_f, \theta^* + 2\sigma_f]$,

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

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

$$\begin{aligned} p_f(\theta^* + \sigma_f \varepsilon_f^i, \lambda^i, \theta^*) &= \frac{\int_{\theta^* - 2\sigma_f}^{\theta^*} \binom{n}{\lambda^i} \hat{l} \left(\frac{\theta^* - \theta'}{\sigma_e} \right)^{\lambda^i} \left[1 - \hat{l} \left(\frac{\theta^* - \theta'}{\sigma_e} \right) \right]^{n - \lambda^i} g \left(\varepsilon_f^i + \frac{\theta^* - \theta'}{\sigma_f} \right) d\theta'}{\int_{\theta^* - 2\sigma_f}^{\theta^* + 2\sigma_f} \binom{n}{\lambda^i} \hat{l} \left(\frac{\theta^* - \theta'}{\sigma_e} \right)^{\lambda^i} \left[1 - \hat{l} \left(\frac{\theta^* - \theta'}{\sigma_e} \right) \right]^{n - \lambda^i} g \left(\varepsilon_f^i + \frac{\theta^* - \theta'}{\sigma_f} \right) d\theta'} \\ &= \frac{\int_0^2 \hat{l} \left(\frac{\sigma_f}{\sigma_e} \Delta \right)^{\lambda^i} \left[1 - \hat{l} \left(\frac{\sigma_f}{\sigma_e} \Delta \right) \right]^{n - \lambda^i} g(\varepsilon_f^i + \Delta) d\Delta}{\int_{-2}^2 \hat{l} \left(\frac{\sigma_f}{\sigma_e} \Delta \right)^{\lambda^i} \left[1 - \hat{l} \left(\frac{\sigma_f}{\sigma_e} \Delta \right) \right]^{n - \lambda^i} g(\varepsilon_f^i + \Delta) d\Delta}. \end{aligned} \quad (3)$$

In particular, the last expression does not depend on θ^* and depends on the scaling parameters only through their ratio σ_f / σ_e . Since the distribution of λ^i conditional on θ^* does not depend on θ^* , equation (3) implies that $\Pr \left[p_f(x_f^i, \lambda^i, \theta^*) > p_f^i \mid \theta^* \right]$ is independent of θ^* , establishing that the equilibrium condition (2) has a unique solution. Moreover, since (3) depends on the scaling parameters σ_e and σ_f only through their ratio, and the distribution of λ^i conditional on θ^* does not depend on the scaling parameters, the solution of (2) similarly depends on the scaling parameters only through their ratio. \square

Proof of Lemma 1. Since $\sigma_e + \sigma_f \leq 0.5 \min(s_f^A, s_f^B)$, we can ignore any boundary effects. Using

Lemma 2, we have

$$\begin{aligned}
\Pr [\pi_e (x_e^j, \theta^*) < p \mid \theta^*] &= \Pr \left[1 - G \left(\frac{x_e^j - \theta^*}{\sigma_e} \right) < p \mid \theta^* \right] \\
&= \Pr [1 - G (\varepsilon_e^j) < p] \\
&= 1 - G [G^{-1}(1 - p)] \\
&= p,
\end{aligned}$$

as needed. □

Proof of Proposition 2. We first solve the game without experts, that is, with $n = 0$. The pivotal condition is

$$\theta_0^* = s_f^A + (1 - s_f^A - s_f^B) \Pr [\pi_f (x_f^i, \theta_0^*) \geq p_f^i \mid \theta_0^*].$$

Using Lemma 1, the condition may be rewritten as

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

We now return to the setting with poorly informed experts. Letting $\sigma_f/\sigma_e \rightarrow 0$ in (3), we have that, conditional on $\theta = \theta^*$, follower i 's posterior belief is

$$p_f (\theta^* + \sigma_f \varepsilon_f^i, \lambda^i, \theta^*) \rightarrow \frac{\int_0^2 g(\varepsilon_f^i + \Delta) d\Delta}{\int_{-2}^2 g(\varepsilon_f^i + \Delta) d\Delta} = 1 - G (\varepsilon_f^i) = \pi_f (\theta^* + \sigma_f \varepsilon_f^i, \theta^*),$$

for every λ^i (recall that $\pi_f (x_f^i, \theta^*)$ is the belief of a follower who has received signal x_f^i and has not observed any expert actions). Therefore, the pivotal conditions in the settings with and without experts coincide as $\sigma_f/\sigma_e \rightarrow 0$. □

The next proposition is used to establish Propositions 3 and 4.

Proposition 5. *As $\sigma_f/\sigma_e \rightarrow +\infty$, the pivotal threshold θ_n^* converges to*

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

where π , λ and p_f^i are independent random variables with $\pi \sim U[0, 1]$, i is a randomly chosen independent follower, and $\lambda \sim B[n, s_e^A + (1 - s_e^A - s_e^B)(1 - p_e)]$.¹³

Proof of Proposition 5. Without loss of generality, we fix σ_f and let $\sigma_e \rightarrow 0^+$. Recall that $\pi_f(x_f^i, \theta^*) = \Pr(\theta < \theta^* | x_f^i)$ denotes the “pre-expert” probability that action A prevails evaluated by follower i conditioning only on his private signal x_f^i (as opposed to the “post-expert” probability $p_f(x_f^i, \lambda^i, \theta^*)$). For $\theta > \theta^*$, $\lim_{\sigma_e \rightarrow 0^+} l(\theta, \theta^*) = s_e^A$ and for $\theta < \theta^*$, $\lim_{\sigma_e \rightarrow 0^+} l(\theta, \theta^*) = 1 - s_e^B$. Hence, using Bayes’ rule, $p_f(x, \lambda, \theta^*)$ converges to

$$\frac{\pi_f(x, \theta^*) \binom{n}{\lambda} (1 - s_e^B)^\lambda (s_e^B)^{n-\lambda}}{\pi_f(x, \theta^*) \binom{n}{\lambda} (1 - s_e^B)^\lambda (s_e^B)^{n-\lambda} + [1 - \pi_f(x, \theta^*)] \binom{n}{\lambda} (s_e^A)^\lambda (1 - s_e^A)^{n-\lambda}}.$$

Straightforward algebraic manipulation gives that, for every x , λ , and θ^* ,

$$\lim_{\frac{\sigma_f}{\sigma_e} \rightarrow +\infty} p_f(x, \lambda, \theta^*) = \frac{\pi_f(x, \theta^*)}{\pi_f(x, \theta^*) + [1 - \pi_f(x, \theta^*)] \left(\frac{s_e^A}{1 - s_e^B}\right)^\lambda \left(\frac{1 - s_e^A}{s_e^B}\right)^{n-\lambda}}. \quad (4)$$

In the limit, 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.

The result follows from (4) together with the pivotal condition (2). \square

Proof of Proposition 3. Proposition 3 is a corollary of Proposition 5. The first statement in Proposition 3 follows from the fact that the distribution of λ is decreasing in p_e (in the sense of first-order stochastic dominance), and the posterior belief

$$\frac{\pi}{\pi + (1 - \pi) \left(\frac{s_e^A}{1 - s_e^B}\right)^\lambda \left(\frac{1 - s_e^A}{s_e^B}\right)^{n-\lambda}}$$

is increasing in λ .

The second statement of Proposition 3 is established as follows. Rearranging the expression for the threshold in Proposition 5 yields

$$\theta_n^* = s_f^A + (1 - s_f^A - s_f^B) \Pr \left[\left(\frac{\pi}{1 - \pi} \frac{1 - p_f^i}{p_f^i} \right)^{1/n} > \left(\frac{s_e^A}{1 - s_e^B} \right)^{\lambda/n} \left(\frac{1 - s_e^A}{s_e^B} \right)^{1 - \lambda/n} \right], \quad (5)$$

¹³Here $B(n, p)$ denotes the binomial distribution for n draws with probability p .

where π , λ , and i are independent random variables with $\pi \sim U[0, 1]$, $\lambda \sim B[n, s_e^A + (1 - s_e^A - s_e^B)(1 - p_e)]$, and i a randomly chosen independent follower. The left-hand side of the inequality in (5) converges in probability to 1. Since λ/n converges in probability to $s_e^A + (1 - s_e^A - s_e^B)(1 - p_e)$ the right-hand side of the inequality converges in probability to

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

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

Proof of Proposition 4. The result follows immediately from Proposition 5. □

References

- Angeletos, G., Hellwig, C. and Pavan, A. (2007). ‘Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks’, *Econometrica*, vol. 75(3), pp. 711–756.
- Aral, S. (2011). ‘Identifying social influence: A comment on opinion leadership and social contagion in new product diffusion’, *Marketing Science*, vol. 30(2), pp. 217–223.
- Arndt, J. (1967). ‘Role of product-related conversations in the diffusion of a new product’, *Journal of Marketing Research*, vol. 4(3), pp. 291–295.
- Baumgarten, S. (1975). ‘The innovative communicator in the diffusion process’, *Journal of Marketing Research*, vol. 12(1), pp. 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*, vol. 104(3), pp. 446–466.
- Corsetti, G., Dasgupta, A., Morris, S. and Shin, H. (2004). ‘Does one Soros make a difference? A theory of currency crises with large and small traders’, *Review of Economic Studies*, vol. 71(1), pp. 87–113.

- Cunningham, D. (2003). 'The patterning of repression: FBI counterintelligence and the New Left', *Social Forces*, vol. 82(1), pp. 209–240.
- Dasgupta, A. (2007). 'Coordination and delay in global games', *Journal of Economic Theory*, vol. 134(1), pp. 195–225.
- DeNardo, J. (1985). *Power in Numbers: The Political Strategy of Protest and Rebellion*, Princeton, NJ: Princeton University Press.
- Edmond, C. (2013). 'Information manipulation, coordination, and regime change', *Review of Economic Studies*, forthcoming.
- Feick, L. and Price, L. (1987). 'The market maven: A diffuser of marketplace information', *Journal of Marketing*, pp. 83–97.
- Francisco, R.A. (1993). 'Theories of protest and the revolutions of 1989', *American Journal of Political Science*, vol. 37(3), pp. 663–680.
- Gilly, M., Graham, J., Wolfinbarger, M. and Yale, L. (1998). 'A dyadic study of interpersonal information search', *Journal of the Academy of Marketing Science*, vol. 26(2), pp. 83–100.
- Gladwell, M. (2000). *The Tipping Point: How Little Things Can Make a Big Difference*, Boston, MA: Little, Brown.
- Goldenberg, J., Han, S., Lehmann, D. and Hong, J. (2009). 'The role of hubs in the adoption process', *Journal of Marketing*, vol. 73(2), pp. 1–13.
- Granovetter, M. (1978). 'Threshold models of collective behavior', *American Journal of Sociology*, vol. 83(6), pp. 1420–1443.
- Guimaraes, B. and Morris, S. (2007). 'Risk and wealth in a model of self-fulfilling currency attacks', *Journal of Monetary Economics*, vol. 54(8), pp. 2205–2230.
- Honwana, A. (2011). 'Youth and the Tunisian revolution', Conflict Prevention and Peace Forum Policy Paper, Social Science Research Council.
- Katz, E. and Lazarsfeld, P. (1955). *Personal Influence: The Part Played by People in the Flow of Mass Communications*, Glencoe, IL: Free Press.

- Katz, M. and Shapiro, C. (1985). 'Network externalities, competition, and compatibility', *American Economic Review*, vol. 75(3), pp. 424–440.
- Kuran, T. (1989). 'Sparks and prairie fires: A theory of unanticipated political revolution', *Public Choice*, vol. 61(1), pp. 41–74.
- Kuran, T. (1991). 'Now out of never: The element of surprise in the East European revolution of 1989', *World Politics*, vol. 44(1), pp. 7–48.
- Lohmann, S. (1993). 'A signaling model of informative and manipulative political action', *American Political Science Review*, vol. 87(2), pp. 319–333.
- Lohmann, S. (1994). 'The dynamics of informational cascades', *World Politics*, vol. 47(1), pp. 42–101.
- Lohmann, S. (1995). 'A signaling model of competitive political pressures', *Economics & Politics*, vol. 7(3), pp. 181–206.
- Lohmann, S. (2000). 'Collective action cascades: An informational rationale for the power in numbers', *Journal of Economic Surveys*, vol. 14(5), pp. 655–684.
- Malewski, J. (2011). 'An ongoing revolution', *International Viewpoint Online Magazine*, vol. IV435.
- McAdam, D. (1986). 'Recruitment to high-risk activism: The case of Freedom Summer', *American Journal of Sociology*, vol. 92(1), pp. 64–90.
- McAdam, D. and Paulsen, R. (1993). 'Specifying the relationship between social ties and activism', *American Journal of Sociology*, vol. 99(3), pp. 640–667.
- Merton, R. (1957). *Social Theory and Social Structure*, Glencoe, IL: Free Press.
- Millman, B. (2005). 'HMG and the war against dissent, 1914-18', *Journal of Contemporary History*, vol. 40(3), pp. 413–440.
- Morris, S. and Shin, H.S. (2003). 'Global games: Theory and applications', in (M. Dewatripont, L. P. Hansen and S. J. Turnovsky, eds.), *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, Cambridge: Cambridge University Press.

- Opp, K. and Gern, C. (1993). 'Dissident groups, personal networks, and spontaneous cooperation: The East German revolution of 1989', *American Sociological Review*, vol. 58(5), pp. 659–680.
- Price, L. and Feick, L. (1984). 'The role of interpersonal sources in external search: An informational perspective', *Advances in Consumer Research*, vol. 11(1), pp. 250–255.
- Steiner, J. (2006). 'Coordination in a mobile world', Working Paper, CERGE-EI.
- Watts, D. (2007). 'Challenging the influentials hypothesis', *WOMMA Measuring Word of Mouth*, vol. 3(4), pp. 201–211.
- Watts, D. and Dodds, P. (2007). 'Influentials, networks, and public opinion formation', *Journal of Consumer Research*, vol. 34(4), pp. 441–458.
- Weimann, G. (1991). 'The influentials: Back to the concept of opinion leaders?', *Public Opinion Quarterly*, vol. 55(2), pp. 267–279.