Keeping the Listener Engaged: A Dynamic Model of Bayesian Persuasion

Yeon-Koo Che

Columbia University

Kyungmin Kim

Emory University

Konrad Mierendorff

University College London

We consider a dynamic model of Bayesian persuasion in which information takes time and is costly for the sender to generate and for the receiver to process, and neither player can commit to their future actions. Persuasion may totally collapse in a Markov perfect equilibrium of this game. However, for persuasion costs sufficiently small, a version of a folk theorem holds: outcomes that approximate Kamenica and Gentzkow's sender-optimal persuasion as well as full revelation and everything in between are obtained in Markov perfect equilibrium as the cost vanishes.

I. Introduction

Persuasion is a quintessential form of communication in which one individual (the sender) pitches an idea, a product, a political candidate, a point of view, or a course of action to another individual (the receiver). Whether the

We thank Emir Kamenica and four anonymous referees for many insightful and constructive suggestions. We are also grateful to Martin Cripps, Jeff Ely, Faruk Gul, Stephan Lauermann,

Electronically published June 2, 2023

Journal of Political Economy, volume 131, number 7, July 2023.

https://doi.org/10.1086/722985

receiver ultimately accepts that pitch—or is persuaded—depends on the underlying truth (the state of the world) but, importantly, also on the information the sender manages to communicate. In remarkable elegance and generality, Kamenica and Gentzkow (2011; henceforth KG) show how the sender should communicate information in such a setting, when she can perform any (Blackwell) experiment instantaneously, without any cost incurred by her or by the receiver. This frictionlessness gives full commitment power to the sender, as she can publicly choose any experiment and reveal its outcome, all before the receiver can act.

In practice, however, persuasion is rarely frictionless. Imagine a salesperson pitching a product to a potential buyer. The buyer may have an interest in buying the product but requires some evidence that it matches his needs. To convince the buyer, the salesperson might demonstrate certain features of the product or marshal customer testimonies and sales records, any of which takes real time and effort. Likewise, to process information, the buyer must pay attention, which is costly. Clearly, these features are present in other persuasion contexts, such as a prosecutor seeking to convince juries or a politician trying to persuade voters.

In this paper, we study the implications of these realistic frictions. Importantly, if information takes time to generate but the receiver can act at any time, the sender no longer automatically enjoys full commitment power. Specifically, she cannot promise to the receiver what experiments she will perform in the future, effectively reducing her commitment power to a current flow experiment. Given the lack of commitment by the sender, the receiver may stop listening and take an action if he does not believe that the sender's future experiments are worth waiting for. The buyer in the example above may walk away at any time when he becomes sufficiently pessimistic about the product or about the prospect of the salesperson eventually persuading him. We will examine to what extent and in what manner the sender can persuade the receiver in this environment with limited commitment. As we will demonstrate, the key challenge facing the sender is to instill the belief that she is worth listening to, namely, to keep the receiver engaged.

We develop a dynamic version of the canonical persuasion model: the state is binary, L or R, and the receiver can take a binary action, ℓ or r. The receiver prefers to match the state by taking action ℓ in state L and r in state R, while the sender prefers action r regardless of the state. Time is

George Mailath, Meg Meyer, Sven Rady, Nikita Roketskiy, Hamid Sabourian, Larry Samuelson, Sara Shahanaghi, and audiences in various seminars and conferences for helpful comments and discussions. Che is supported by the National Science Foundation (SES-1851821). Che and Kim are supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2020S1A5A2A03043516). This paper was edited by Emir Kamenica. Our dear friend and coauthor, Konrad Mierendorff, passed away in August 2021. All the ideas and results in this paper are collaborative work by the three authors, but Che and Kim are responsible for any remaining errors.

continuous and the horizon is infinite. At each point in time, unless the game has ended, the sender may perform some flow experiment. In response, the receiver either takes an action and ends the game or simply waits and continues the game. Both the sender's choice of experiment and its outcome are publicly observable. Therefore, the two players always share a common belief about the state.

The sender has a rich class of Poisson experiments at her disposal. Specifically, we assume that at each instant, the sender can generate a collection of Poisson signals. The possible signals are flexible in their directionalities: a signal can be either good news (inducing a posterior above the current belief) or bad news (inducing a posterior below the current belief). In addition, the news can be of arbitrary accuracy: the sender can choose any target posterior, although more accurate signals (with targets closer to 0 or 1) arrive at a lower rate. Our model generalizes the existing Poisson models in the literature, which considered either a good news or a bad news Poisson experiment of given accuracy (e.g., Keller, Rady, and Cripps 2005; Keller and Rady 2015; Che and Mierendorff 2019).

Any experiment, regardless of its accuracy, requires a flow cost $c > 0$ (per unit of time) for the sender to perform and for the receiver to process. That the cost is the same for both players is a convenient normalization, with no material consequence (see n. 8). Our model of information allows for the flexibility and richness of Kamenica and Gentzkow (2011) but adds the friction that information takes time to generate. This serves to isolate the effects of the friction.

We may interpret the model in the canonical communication context, such as a salesperson pitching a product to a buyer. The former is trying to persuade the latter that the product fits his needs, an event denoted by R. Once inside the store, the buyer is deciding whether to listen to the pitch (wait), leave the store (action ℓ), or purchase the product (action r). We interpret the series of pitches made by the salesperson as experiments. A salesperson's pitches may include the types of product features demonstrated as well as her manner, tone, and body language with which her messages are delivered. Hence, the pitches can reveal a lot about what she is intending to say, not just what she is saying, consistent with public observability of experiments assumed in our model. Meanwhile, whether the pitches succeed depends on the buyer's idiosyncratic needs and is uncertain from the salesperson's perspective. It is also reasonable that an experienced salesperson could get feedback on her pitch directly or indirectly from the buyer's reactions, which would make the outcome of the experiment public. As in our model, the key issue is whether the buyer believes the salesperson's pitches to be worth listening to. Our analysis will focus on this issue.

We study *Markov perfect equilibria* (MPEs) of this game, that is, subgame perfect equilibrium strategy profiles that prescribe the sender's flow experiment and the receiver's action (ℓ, r, or wait) at each belief p, the probability that the

state is R. We are particularly interested in the equilibrium outcomes when the frictions are sufficiently small (i.e., in the limit as the flow cost c converges to zero). In addition, we investigate the *persuasion dynamics*, or the type of pitch the sender uses to persuade the receiver in equilibria of this game.

Is persuasion possible? If so, to what extent?—Whether the sender can persuade the receiver depends on whether the receiver finds her worth listening to—or, more precisely, on his belief that the sender will provide enough information to justify his listening costs. This belief depends on the sender's future experimentation strategy, which in turn rests on what the receiver will do if the sender betrays her trust and reneges on her information provision. The multitude of ways in which the players can coordinate on these choices yields a folk theorem–like result. There is an MPE in which no persuasion occurs. When the cost c becomes arbitrarily small, however, we also obtain a set of persuasion equilibria that ranges from ones that approximate Kamenica and Gentzkow's (2011) sender-optimal persuasion to ones that approximate full revelation; we show that any sender (receiver) payoff between these two extremes is attainable in the limit as c tends to zero.

In the *no-persuasion* equilibrium, the receiver is pessimistic about the sender generating sufficient information, so he simply takes an action without waiting for information. Facing this pessimism, the sender becomes desperate and maximizes her chance of once-and-for-all persuasion involving minimal information, which turns out to be the sort of strategy that the receiver would not find worth waiting for, justifying his pessimism.

In a persuasion equilibrium, by contrast, the receiver expects the sender to deliver sufficient information to compensate his listening costs. This optimism in turn motivates the sender to deliver on her promise of informative experimentation; if she reneges on her experimentation, the ever optimistic receiver would simply wait for experimentation to resume an instant later instead of taking the action that the sender would like him to take. In short, the receiver's optimism fosters the sender's generosity in information provision, which in turn justifies this optimism. As we will show, equilibria with this virtuous cycle of beliefs can support a wide range of outcomes from Kamenica and Gentzkow's (2011) optimal persuasion to full revelation, as the flow cost c tends to zero.[1]

Persuasion dynamics.—Our model informs us of what kind of pitch the sender should make at each point in time; how long it takes for the sender to persuade the receiver, if ever; and how long the receiver listens

[1] The mechanism using a virtuous cycle of beliefs to support cooperative behavior in a dynamic environment has been utilized in other economic contexts. Among others, Che and Sákovics (2004) show how this mechanism can be used to overcome the holdup problem. In fact, the main tension in our dynamic persuasion problem can be interpreted as a holdup problem: the receiver wants to avoid incurring listening costs if the sender will behave opportunistically and not provide sufficient information. However, the current paper differs in other crucial aspects; in particular, the rich choice of information structures is unique here and has no analog in Che and Sákovics (2004).

to the sender before taking an action. The dynamics of the persuasion strategy adopted in equilibrium involves rich behavioral implications that are absent in the static persuasion model.

In our MPEs, the sender optimally makes use of the following three strategies: (1) confidence building, (2) confidence spending, and (3) confidence preserving. The *confidence-building* strategy involves a bad news Poisson experiment that induces the receiver's belief (that the state is R) to either drift upward or jump to zero. Under this strategy, the belief moves upward for sure when the state is R and quite likely even when the state is *L*; in fact, this strategy minimizes the probability of bad news by insisting that the news be conclusive. The sender finds it optimal to use this strategy when the receiver's belief is already close to the persuasion target (i.e., the belief that will trigger him to choose r).

The *confidence-spending* strategy involves a good news Poisson experiment that generates an upward jump to some target belief, either one inducing the receiver to choose r or at least one inducing him to listen to the sender. Such a jump arises rarely, however, and absent this jump, the receiver's belief drifts downward. In this sense, this strategy is a risky one that spends the receiver's confidence over time. This strategy is used when the receiver is already quite pessimistic about R, so that either the confidence-building strategy would take too long, or the receiver would simply not listen. In particular, it is used as a "last ditch" effort, when the sender is close to giving up on persuasion or when the receiver is about to choose ℓ.

The *confidence-preserving* strategy combines the above two strategies—namely, a good news Poisson experiment inducing the belief to jump to a persuasion target and a bad news Poisson experiment inducing the belief to jump to zero. This strategy is effective if the receiver is sufficiently skeptical relative to the persuasion target so that the confidence-building strategy will take too long. Confidence spending could expedite persuasion for a range of beliefs but would run down the receiver's confidence in the process. Hence, at some point, the sender finds it optimal to switch to the confidence-preserving strategy, which prevents the receiver's belief from deteriorating further. The belief where the sender switches to this strategy constitutes an absorbing point of the belief dynamics; from then on, the belief does not move, unless either a sudden persuasion breakthrough or breakdown occurs.

The equilibrium strategy of the sender combines these three strategies in different ways under different economic conditions, thereby exhibiting rich and novel persuasion dynamics. Our characterization in section V describes precisely how the sender uses them in different equilibria.

Related literature.—This paper primarily contributes to the Bayesian persuasion literature that began with Kamenica and Gentzkow (2011) by studying the problem in a dynamic environment. Several recent papers also consider dynamic models (e.g., Brocas and Carrillo 2007; Kremer,

Mansour, and Perry 2014; Au 2015; Ely 2017; Renault, Solan, and Vieille 2017; Che and Hörner 2018; Henry and Ottaviani 2019; Ely and Szydlowski 2020; Orlov, Skrzypacz, and Zryumov 2020; Marinovic and Szydlowski 2020; Bizzotto, Rüdiger, and Vigier 2021). Our focus is different from most of these papers since we consider gradual production of information and assume that there is no commitment.[2]

Two papers closest to ours in this regard are Brocas and Carrillo (2007) and Henry and Ottaviani (2019), who restrict the set of feasible experiments so that information arrives gradually. The former considers a binary signal in a discrete-time setting, and the latter employs a drift-diffusion model in a continuous-time setting.[3] Unlike our model, the receiver in their models cannot stop listening and take an action at any time: he can move only after the sender stops experimenting (Brocas and Carrillo 2007) or applies for approval (Henry and Ottaviani 2019). This modeling difference reflects interests in different economic problems/contexts; for example, Henry and Ottaviani (2019) focus on regulatory approval, while we study persuasive communication. However, the difference leads to very different persuasion outcomes: in their models, complete persuasion failure never occurs, and there exists a unique equilibrium.[4] Another important difference is that the sender in their models does not enjoy the richness and control of information structures: in both papers, the sender

[2] Orlov, Skrzypacz, and Zryumov (2020) characterize an equilibrium that resembles some aspects of our equilibrium in a model where the sender (agent) faces no constraint in the release of information. In particular, they show that the sender may pipet information—release information gradually—in a way that resembles our confidence-building (R-drifting) strategy. The resemblance is more apparent than fundamental, however. In their main model, the sender intrinsically prefers the receiver to delay exercise of a real option; i.e., the delay of the receiver's action per se is desired by the sender. She can fully reveal the state instantaneously but chooses to delay release of information in order to incentivize the receiver to wait longer. In our model, the sender has no intrinsic preferences for delay and provides information only to persuade the receiver to take a particular final action.

[3] McClellan (2022) and Escudé and Sinander (2023) also study dynamic persuasion in drift-diffusion models. McClellan (2022) characterizes the optimal dynamic approval mechanism under full commitment. Escudé and Sinander (2023) consider a sender dynamically optimizing against a receiver who chooses a series of actions myopically.

[4] Henry and Ottaviani (2019) consider three regimes that differ in the players' commitment power. Their informer authority regime corresponds to the sender-optimal dynamic outcome, in that the sender stops as soon as the belief reaches the minimal point at which the receiver is willing to take action r (approves the project). It is easy to show that in this case, if the receiver could reject/accept the project unilaterally at any time and discounted his future payoff or incurred a flow cost as in our model, he would take an action immediately without listening, and persuasion would fail completely. Their no-commitment regime is similar to our model but with the crucial difference that the sender does not have the option to pass, i.e., to stop experimenting without abandoning the project. This feature allows the receiver (e.g., a drug approver) to force the sender to keep experimenting, resulting in the receiver-optimal persuasion as the unique equilibrium outcome. If passing were an available option, as we assume in our model, multiple equilibria supported by virtuous cycles of beliefs would arise even in their drift-diffusion model, producing a range of persuasion outcomes and ultimately leading to the same kind of result as our theorem 2 (see n. 18). Finally, their evaluator authority case is obtained when the receiver can commit to an acceptance threshold.

decides simply whether to continue and has no influence over the type of information generated.

The receiver's problem in our paper involves a stopping problem, which has been studied extensively in the single agent context, beginning with Wald (1947) and Arrow, Blackwell, and Girshick (1949). In particular, Nikandrova and Pancs (2018), Che and Mierendorff (2019), and Mayskaya (2020) study an agent's stopping problem when she acquires information through Poisson experiments.[5] Che and Mierendorff (2019) introduced the general class of Poisson experiments adopted in this paper. However, the generality is irrelevant in their model, because unlike here, the decision maker optimally chooses only between two conclusive experiments (i.e., never chooses a nonconclusive experiment).

The paper is organized as follows. Section II introduces the model. Section III illustrates the main ideas of our equilibria. Sections IV and V characterize our MPE strategies and study their payoff implications. Section VI concludes.

II. Model

We consider a game in which a sender (she) wishes to persuade a receiver (he). There is an unknown state ω that can be either L (left) or R (right). The receiver ultimately takes a binary action ℓ or r, which yields the following payoffs:

	Payoffs for the Sender and Receiver	
States/Actions	ℓ	r
L	$(0,\ u_\ell^{\mathrm{L}})$	$(v,\ u_{\mathrm{r}}^{\mathrm{L}})$
R	$(0,\ u_\ell^{\mathrm{R}})$	$(v,\ u_{\mathrm{r}}^{\mathrm{R}})$

The receiver gets u_a^ω if he takes action $a \in \{\ell, \mathrm{r}\}$ when the state is $\omega \in \{\mathrm{L}, \mathrm{R}\}$. The sender's payoff depends only on the receiver's action: she gets v if the receiver takes r and zero otherwise. We assume $u_\ell^{\mathrm{L}} > \max\{u_{\mathrm{r}}^{\mathrm{L}}, 0\}$ and $u_{\mathrm{r}}^{\mathrm{R}} > \max\{u_\ell^{\mathrm{R}}, 0\}$, so that the receiver prefers to match the action with the state, and also $v > 0$, so that the sender prefers action r to action ℓ. Both players begin with a common prior p_0 that the state is R and use Bayes's rule to update their beliefs.

KG benchmark.—By now, it is well understood how the sender optimally persuades the receiver if she can commit to an experiment without any restrictions. For each $a \in \{\ell, \mathrm{r}\}$, let $U_a(p)$ denote the receiver's expected

[5] The Wald stopping problem has also been studied with drift-diffusion learning (e.g., Moscarini and Smith 2001; Fudenberg, Strack, and Strzalecki 2018; Ke and Villas-Boas 2019) and in a model that allows for general endogenous experimentation (see Zhong 2022).

payoff when he takes action a with belief p. In addition, let $\hat{p}$ denote the belief at which the receiver is indifferent between actions ℓ and r, that is, $U_\ell(\hat{p}) = U_r(\hat{p})$.[6]

If the sender provides no information, then the receiver takes action r when $p_0 \geq \hat{p}$. Therefore, persuasion is necessary only when $p_0 < \hat{p}$. In this case, the KG solution prescribes an experiment that induces only two posteriors, $q_- = 0$ and $q_+ = \hat{p}$. The former leads to action ℓ, while the latter results in action r. This experiment is optimal for the sender, because $\hat{p}$ is the minimum belief necessary to trigger action r, and setting $q_- = 0$ maximizes the probability of generating $\hat{p}$ and thus action r. The resulting payoff for the sender is $p_0 v/\hat{p}$, as given by the dashed line in the left panel of figure 1. The flip side is that the receiver enjoys no rents from persuasion; his payoff is $\mathcal{U}(p_0) := \max\{U_\ell(p_0), U_r(p_0)\}$, the same as if no information were provided, as depicted in the right panel of figure 1.

Dynamic model.—We consider a dynamic version of the above Bayesian persuasion problem. Time flows continuously starting at 0. Unless the game has ended, at each point in time $t \geq 0$, the sender may perform an experiment at a constant flow cost c from a feasible set, which will be described precisely below, or pass—not running any experiment and not incurring the flow cost c.[7] Just as it is costly for the sender to produce information, it is also costly for the receiver to process it. Specifically, if the sender experiments, then the receiver also pays the same flow cost and observes the experiment and its outcome. After that, he decides whether to take an irreversible action (ℓ or r) or to wait and listen to the information provided by the sender in the next instant. The former ends the game, while the latter lets the game continue.

There are two notable modeling assumptions. First, the receiver can stop listening to the sender and take a game-ending action at any point in time. This is the fundamental difference from Kamenica and Gentzkow (2011), wherein the receiver is allowed to take an action only after the sender finishes her information provision. Second, the players' flow costs are assumed

[6] Specifically, for each $p \in [0, 1]$, $U_\ell(p) := pu_\ell^{\mathrm{R}} + (1-p)u_\ell^{\mathrm{L}}$ and $U_r(p) := pu_r^{\mathrm{R}} + (1-p)u_r^{\mathrm{L}}$. Therefore, $\hat{p} = (u_\ell^{\mathrm{L}} - u_r^{\mathrm{L}})/(u_r^{\mathrm{R}} - u_\ell^{\mathrm{R}} + u_\ell^{\mathrm{L}} - u_r^{\mathrm{L}})$, which is well defined in (0, 1) under our assumptions on the receiver's payoffs.

[7] Passing enables the sender to stop experimenting at no cost. As will be seen, the experimentation always costs $c > 0$ even at low intensity (informativeness). While this involves a form of discontinuity, it is largely for analytic convenience. Our results remain unchanged even if the cost is proportional to the intensity of the experiment (see Che, Kim, and Mierendorff 2021). One may also wonder what would happen if passing incurs the same cost c as experimentation—a natural assumption if c is interpreted as the waiting cost rather than the experimentation cost. Our main results would still go through under this assumption, except for some details of the equilibrium characterization. Without the sender being able to freely stop experimenting, she would never give up on persuading, so the lower boundary of the experimentation region, denoted by p_* later, is always determined by the receiver's incentives, as in proposition 2.

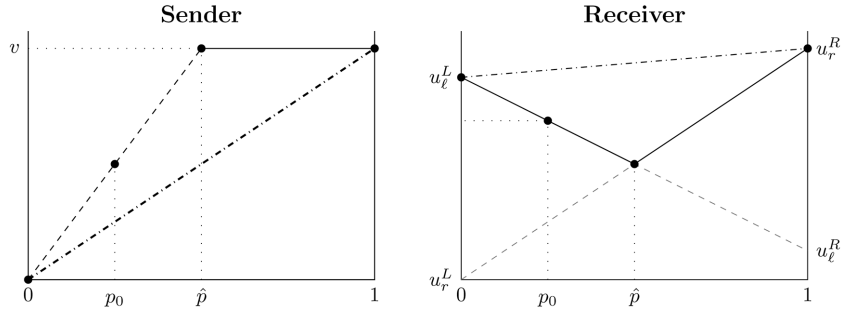

FIG. 1.—Payoffs from static persuasion. *Solid lines*, payoffs without persuasion (information). *Dashed line*, sender's expected payoff in KG solution. *Dash-dotted lines*, payoffs under a fully revealing experiment.

to be the same. This is, however, just a normalization that allows us to directly compare the players' payoffs, and all subsequent results can be reinterpreted as relative to each player's individual flow cost.[8]

Feasible experiments.—We consider a general class of experiments whose informativeness per unit time is bounded in a proper way. Formally, we let p_t denote the belief that $\omega = \mathrm{R}$ at time t and represent an experiment by a regular martingale process $\langle p_t \rangle$—that is, a càdlàg martingale process over $X \coloneqq [0, 1]$ that is progressively measurable with respect to its natural filtration $\{\mathcal{F}_t\}$—with countably many discontinuities and a deterministic continuous path at each point in history. Its martingale property follows from the law of iterated expectations (or Bayes plausibility). We let $\mathcal{P}$ denote the set of all regular martingale processes.[9]

For any $q \neq p_{t-} \coloneqq \lim_{t' \uparrow t} p_{t'}$, let $\lambda^\omega(q, p_{t-}) \coloneqq \lim_{dt \to 0} \mathbb{P}[p_t = q | p_{t-dt}, \omega]/dt$ denote the rate at which the belief changes from p_{t-} to q in state ω. The set of feasible experiments is then defined as

$$\mathcal{P}^* \coloneqq \left\{ \langle p_t \rangle \in \mathcal{P} : \sum_{q \neq p_{t-}} |\lambda^{\mathrm{R}}(q, p_{t-}) - \lambda^{\mathrm{L}}(q, p_{t-})| \leq \lambda \text{ for all } t \text{ and } p_{t-} \right\}.$$

[8] Suppose that the sender's cost is given by c_s, while that of the receiver is c_r. Such a model is equivalent to our normalized one in which $c_r' = c_s' = c_r$ and $v' = v(c_r/c_s)$. When solving the model for a fixed set of parameters $(u_a^\omega, v, c, \lambda)$, this normalization does not affect the results. If we let c tend to zero, we are implicitly assuming that the sender's and receiver's (unnormalized) costs, c_s and c_r, converge to zero at the same rate. See n. 26 for a relevant discussion.

[9] The requirement of a deterministic continuous path means that $\mathcal{P}$ does not include diffusion processes such as Brownian motion. But the class $\mathcal{P}$ encompasses a large class of jump (Poisson) processes. The implications of relaxing this requirement for our results (i.e., whether the sender would prefer a belief process failing this requirement to Poisson processes we allow in our model) remain an open question.

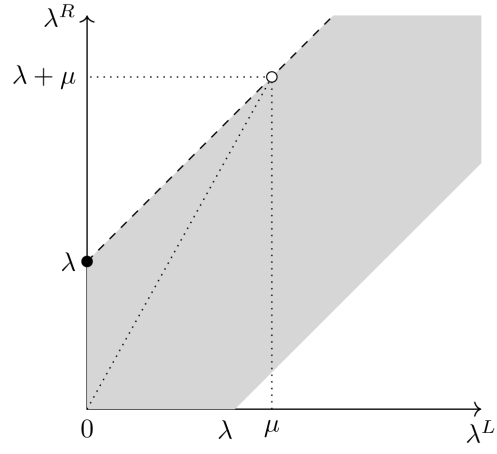

FIG. 2.—Arrival rates of feasible Poisson experiments.

The set $\mathcal{P}^*$ includes all Poisson processes whose state-contingent jump rates (λ^L, λ^R) satisfy $|\lambda^L - \lambda^R| \leq \lambda$ at each point in history; the feasible arrival rates are depicted by the shaded area in figure 2. It also includes all mixtures of those Poisson experiments.

In fact, any information in our class can be generated by diluting a conclusive Poisson signal arriving at rate λ. Consider a conclusive signal that arrives in state R at rate λ (black circle in fig. 2). One can then add a white noise arriving in both states at some rate μ to this conclusive signal. The resulting signal (white circle in fig. 2) then arrives more frequently at rates $(\mu, \lambda + \mu)$ but is less precise, moving the belief only to posterior $q = p(\lambda + \mu)/[(1 - p)\mu + p(\lambda + \mu)](< 1)$. The constant bound for the arrival rate differences means that the constraint on flow information is independent of a prior, or *experimental*, as defined by Denti, Marinacci, and Rustichini (2022); this stands in contrast to other models, such as rational inattention, which assumes (prior-dependent) Shannon information cost or capacity.

LEMMA 1. An experiment $\langle p_t \rangle$ is feasible (i.e., $\langle p_t \rangle \in \mathcal{P}^*$) if and only if the following property holds at each point in history: there exists $\alpha : [0, 1] \to [0, 1]$ such that $\Sigma_{q \neq p} \alpha(q) \leq 1$;

a. for any $q \neq p$, the arrival rate of posterior belief q given $p_{t^-} = p$ is equal to

$$\alpha(q) \frac{\lambda p(1 - p)}{|q - p|};$$

and

b. conditional on no jump, the belief drifts according to

$$\dot{p} = -\left(\sum_{q>p}\alpha(q) - \sum_{q<p}\alpha(q)\right)\lambda p(1-p).$$

Proof.—See appendix A.

Lemma 1 shows that a feasible flow experiment can be represented by the shares α of a unit capacity allocated to Poisson experiments that trigger jumps to alternative posterior beliefs q at rates $\alpha(q)\lambda p(1-p)/|q-p|$. The jump rate in part a simplifies to an expression familiar from the existing literature when the sender triggers a single jump with $\alpha(q) = 1$ to conclusive news with either $q = 0$ or $q = 1$. For instance, conclusive R-evidence ($q = 1$) is obtained at the rate of λp, as is assumed in good news models (see, e.g., Keller, Rady, and Cripps 2005). Likewise, conclusive L-evidence ($q = 0$) is obtained at the rate of $\lambda(1-p)$, as is assumed in bad news models (see, e.g., Keller and Rady 2015). Our model allows for such conclusive news, but it also allows for arbitrary nonconclusive news with $q \in (0, 1)$ as well as any arbitrary mixture among such experiments. Further, our information constraint captures the intuitive idea that more accurate information takes longer to generate. For example, if we assume $q > p$, the arrival rate increases as the news becomes less precise (q falls), and it approaches infinity as the news becomes totally uninformative (i.e., as q tends to p). Last, limited arrival rates capture an important feature of our model that any meaningful persuasion takes time and requires delay.

Part b describes the law of motion governing the drift of beliefs when no jump occurs. Strikingly, the drift rate depends only on the difference between the fractions of the capacity allocated to right versus left Poisson signals. That is, the rate does not depend on the precision q of the news in the individual experiments. The reason is that the precision of news and its arrival rate offset each other, leaving the drift rate unaffected.[10] This feature makes the analysis tractable while at the same time generalizing conclusive Poisson models in an intuitive way.

Among many feasible experiments, the following three (visualized in fig. 3) will prove particularly relevant for our purposes. They formalize the three modes of persuasion discussed in section I.

- R-drifting experiment (confidence building): $\alpha(0) = 1$. The sender devotes all her capacity to a Poisson experiment with (posterior) jump target $q = 0$. In the absence of a jump, the posterior drifts to the right at rate $\dot{p} = \lambda p(1-p)$.

[10] Suppose $q > p$. This means that the sender has chosen a rate λ for the informative signal and $\mu \geq 0$ for the noise. It is clear that μ does not affect the updating of the state since the noise arrives at the same rate in both states.

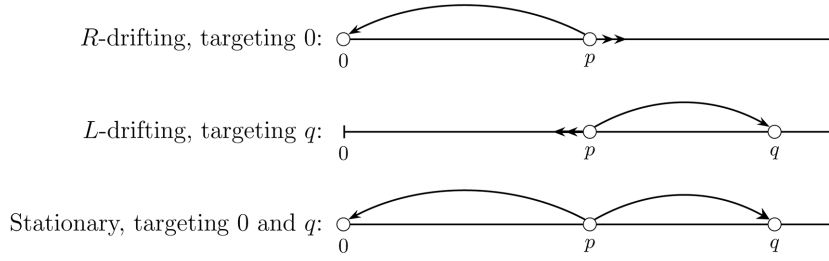

FIG. 3.—Three prominent feasible experiments.

- L-drifting experiment (confidence spending): $\alpha(q) = 1$ for some $q > p$. The sender devotes all her capacity to a Poisson experiment with jumps targeting some posterior $q > p$. The precise jump target q will be specified in our equilibrium construction. In the absence of a jump, the posterior drifts to the left at rate $\dot{p} = -\lambda p(1 - p)$.
- Stationary experiment (confidence preserving): $\alpha(0) = \alpha(q) = 1/2$ for some $q > p$. The sender assigns an equal share of her capacity to an experiment targeting 0 and one targeting q. Absent jumps, the posterior remains unchanged.

Solution concept.—We study (pure-strategy) MPEs of this dynamic game in which both players' strategies depend only on the current belief p.[11] Formally, a profile of Markov strategies specifies, for each $p \in [0, 1]$, a flow experiment $\sigma^{S}(p) = (\alpha(q; p))_{q\in[0,1]}$ chosen by the sender and an action $\sigma^{R}(p) \in \{\ell, \mathrm{r}, \mathrm{wait}\}$ chosen by the receiver. Given $\sigma = (\sigma^{S}, \sigma^{R})$ and prior belief p_0, let p_t denote the belief at time t induced by the strategy profile and τ denote the stopping time at which the receiver takes action ℓ or r. Then, the sender's expected payoff is given by

$$V^{\sigma}(p_0) = v\,\mathbb{P}[\sigma^{R}(p_\tau) = \mathrm{r}|p_0] - c\mathbb{E}\left[\int_0^\tau \mathbf{1}_{\left\{\sum_{q\neq p}\alpha(q;p_t) > 0\right\}}\,dt|p_0\right],$$

while the receiver's expected payoff is given by

$$U^{\sigma}(p_0) = \mathbb{E}\left[U_{\sigma^{R}(p_\tau)}(p_\tau)|p_0\right] - c\mathbb{E}\left[\int_0^\tau \mathbf{1}_{\left\{\sum_{q\neq p}\alpha(q;p_t) > 0\right\}}\,dt|p_0\right].$$

[11] Naturally, this solution concept limits the use of (punishment) strategies depending on the payoff-irrelevant part of the histories and serves to discipline strategies off the equilibrium path. For non-Markov equilibria, see our discussion in sec. VI.

A strategy profile $\sigma = (\sigma^S, \sigma^R)$ is *admissible* if the law of motion governing the belief evolution is well defined (see app. B for details) and the stopping time τ is also well defined. Let Σ denote the set of all admissible strategy profiles.

DEFINITION 1 (Markov perfect equilibrium). A strategy profile $\sigma = (\sigma^S, \sigma^R) \in \Sigma$ is an MPE if

i. $V^{\sigma}(p) \geq V^{\hat{\sigma}}(p)$ for all $p \in [0, 1]$ and $\hat{\sigma} = (\hat{\sigma}^S, \sigma^R) \in \Sigma$;
ii. $U^{\sigma}(p) \geq U^{\hat{\sigma}}(p)$ for all $p \in [0, 1]$ and $\hat{\sigma} = (\sigma^S, \hat{\sigma}^R) \in \Sigma$; and
iii. for any p such that the receiver stops (i.e., $\sigma^R(p) \in \{\ell, r\}$):

$$\sigma^S(p) \in \arg\max_{\alpha(\cdot;p)} \sum_q \alpha(q, p) \frac{\lambda p(1-p)}{|q-p|} \left(V^{\sigma}(q) - \mathbf{1}_{\{\sigma^R(p)=r\}} v \right) - \mathbf{1}_{\{\sum \alpha(q;p_t)>0\}} c.$$

Whereas properties i and ii are obvious equilibrium requirements, property iii imposes a restriction that captures the spirit of perfection in our continuous-time framework. To see its role clearly, suppose that the receiver would choose action ℓ unless the sender changes the belief significantly by running a flow experiment. In discrete time, the sender would simply choose a flow experiment that maximizes her expected payoff. In continuous time, however, the sender's strategy at such a point is inconsequential for her payoff; with probability 1, the game would end with the receiver taking action ℓ. With no further restriction on the sender's strategy, this continuous time peculiarity leads to severe but uninteresting equilibrium multiplicity (see n. 15). Property iii enables us to avoid the problem by requiring the sender to choose a strategy that maximizes her instantaneous payoff normalized by dt in the stopping region; it can be seen as selecting an MPE that is robust to a discrete-time approximation.

III. Illustration: Persuading the Receiver to Listen

We begin by illustrating the key issue facing the sender: persuading the receiver to listen. To this end, consider any prior $p_0 < \hat{p}$ so that persuasion is not trivial and suppose that the sender repeatedly chooses R-drifting experiments with jumps targeting $q = 0$ until the posterior either jumps to zero or drifts to $\hat{p}$, as depicted on the horizontal axis in figure 4. This strategy exactly replicates the KG solution (in the sense that it yields the same probabilities of reaching the two posteriors, 0 and $\hat{p}$), provided that the receiver listens to the sender for a sufficiently long time.

But will the receiver wait until the belief reaches 0 or $\hat{p}$? The answer is no. The KG experiment leaves no rents for the receiver without listening costs, and thus with listening costs the receiver will be strictly worse off than if he picks ℓ immediately. In figure 4, the receiver's expected gross

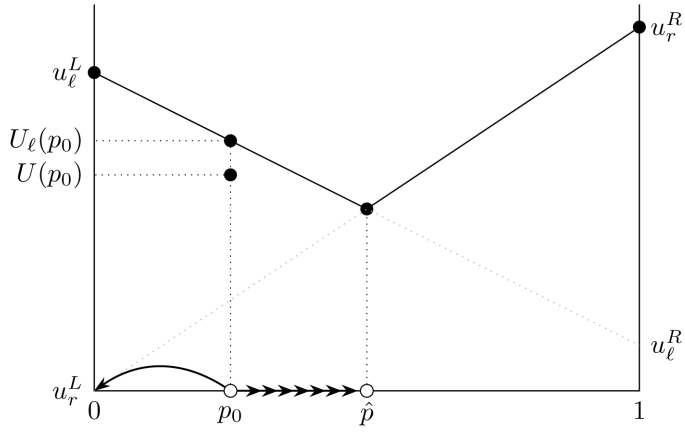

FIG. 4.—Replicating KG outcome through R-drifting experiments.

payoff from the static KG experiment is $U_\ell(p_0)$. Because of the listening costs, the receiver's expected payoff under the dynamic KG strategy, denoted here by $U(p_0)$, is strictly smaller than $U_\ell(p_0)$. In other words, the dynamic strategy implementing the KG solution cannot persuade the receiver to wait and listen, so it does not permit any persuasion.[12] Indeed, this problem leads to the existence of a no-persuasion MPE, regardless of the listening cost.

THEOREM 1 (Persuasion failure). For any $c > 0$, there exists an MPE in which no persuasion occurs; that is, for any p_0, the receiver immediately takes either action ℓ or r.

Proof. Consider the following strategy profile: the receiver chooses ℓ for $p < \hat{p}$ and r for $p \geq \hat{p}$, and the sender chooses the L-drifting experiment with jump target $\hat{p}$ for all $p \in [\hat{\pi}_{\ell L}, \hat{p})$ and passes for all $p \notin [\hat{\pi}_{\ell L}, \hat{p})$, where the cutoff $\hat{\pi}_{\ell L}$ is the belief at which the sender is indifferent between the L-drifting experiment and stopping (followed by ℓ).[13]

[12] The KG outcome can also be replicated by other dynamic strategies. For instance, the sender could repeatedly choose a stationary strategy with jumps targeting 0 and $\hat{p}$ until either jump occurs. However, this (and, in fact, any other) strategy would not incentivize the receiver to listen, for the same reason as in the case of repeating R-drifting experiments.

[13] Specifically, $\hat{\pi}_{\ell L}$ equates the sender's flow cost c to the instantaneous benefit from the L-drifting experiment:

$$c = \frac{\lambda \hat{\pi}_{\ell L}(1 - \hat{\pi}_{\ell L})}{\hat{p} - \hat{\pi}_{\ell L}} v,$$

where the right-hand side is the sender's benefit v from persuasion multiplied by the rate at which the rightward jump to $\hat{p}$ occurs (under the L-drifting experiment) at belief $\hat{\pi}_{\ell L}$. Solving the equation yields

In order to show that this strategy profile is indeed an equilibrium, first consider the receiver's incentives given the sender's strategy. If $p \notin [\hat{\pi}_{\ell L}, \hat{p})$, then the sender never provides information, so the receiver has no incentive to wait and will take an action immediately. If $p \in [\hat{\pi}_{\ell L}, \hat{p})$, then the sender never moves the belief into the region where the receiver strictly prefers to take action r (i.e., strictly above $\hat{p}$). This implies that the receiver's expected payoff is equal to $U_\ell(p_0)$ minus any listening cost she may incur. Therefore, again, it is optimal for the receiver to take an action immediately.

Now consider the sender's incentives given the receiver's strategy. If $p \geq \hat{p}$, then it is trivially optimal for the sender to pass. Now suppose that $p < \hat{p}$. Our refinement (property iii in definition 1) requires that the sender choose a flow experiment that maximizes her instantaneous payoff, which is given by[14]

$$\max_{\alpha(\cdot;p)} \sum_{q \neq p} \alpha(q; p) \lambda \frac{p(1-p)}{|q-p|} \mathbf{1}_{\{q \geq \hat{p}\}} v - \mathbf{1}_{\left\{\sum_{q \neq p} \alpha(q;p) > 0\right\}} c \text{ subject to } \sum_{q \neq p} \alpha(q; p) \leq 1.$$

If the sender chooses any nontrivial experiment, its jump target must be $q = \hat{p}$. Hence the sender's best response is either to maximize the jump rate to $\hat{p}$ (i.e., $\alpha(\hat{p}; p) = 1$) or to pass. The former is optimal if and only if $\lambda p(1-p)/(\hat{p}-p)v \geq c$, or equivalently, $p \geq \hat{\pi}_{\ell L}$.[15] QED

The no-persuasion equilibrium constructed in the proof showcases a total collapse of trust between the two players. The receiver does not trust the sender to convey valuable information (i.e., to choose an experiment targeting $q > \hat{p}$), so he refuses to listen to her. This attitude makes the sender desperate for a quick breakthrough; she tries to achieve persuasion by targeting just $\hat{p}$, which is indeed not enough for the receiver to be willing to wait.

Can trust be restored? In other words, can the sender ever persuade the receiver to listen to her? She certainly can, if she can commit to a dynamic strategy, that is, if she can credibly promise to provide more information in the future. Consider the following modification of the dynamic KG strategy discussed above: the sender repeatedly chooses R-drifting experiments with jumps targeting zero until either the jump occurs or the

$$\hat{\pi}_{\ell L} = \frac{1}{2} + \frac{c}{2\lambda v} - \sqrt{\left(\frac{1}{2} + \frac{c}{2\lambda v}\right)^2 - \frac{c\hat{p}}{\lambda v}}.$$

[14] The objective function follows from the fact that under the given strategy profile, the sender's value function is $V(p) = v$ if $p \geq \hat{p}$ and $V(p) = 0$ otherwise; and when the target posterior is q, a Poisson jump occurs at rate $\lambda p(1-p)/|q-p|$.

[15] Absent point iii in definition 1, there are many additional equilibria in which, in the stopping region, the sender may simply refuse to experiment or adopt an arbitrary Poisson experiment with jumps targeting beliefs other than $\hat{p}$ within the same stopping region. None of these alternative equilibria survive in the corresponding discrete-time setting. Our refinement allows us to select the continuous-time limit of the unique discrete-time no-persuasion equilibrium, and theorem 1 holds despite this refinement.

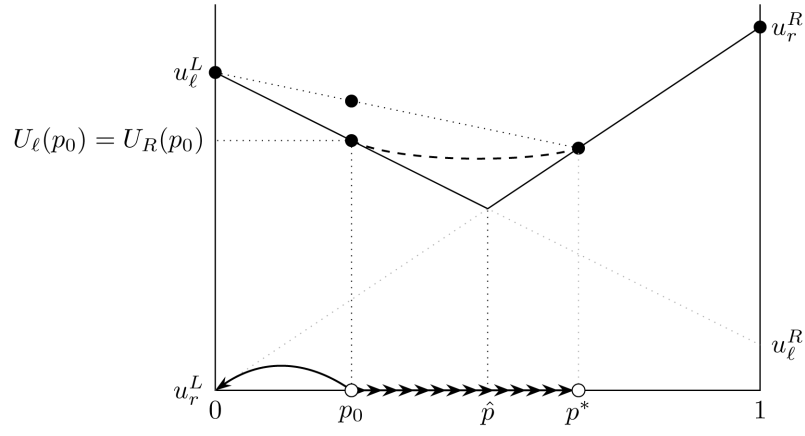

FIG. 5.—Persuasive R-drifting experiments.

belief reaches $p^* > \hat{p}$. If the receiver waits until the belief either jumps to zero or reaches p^*, then her expected payoff is equal to[16]

$$U_{\mathrm{R}}(p) = \frac{p^* - p}{p^*} u_\ell^{\mathrm{L}} + \frac{p}{p^*} U_{\mathrm{r}}(p^*) - \left(p \log\left(\frac{p^*}{1 - p^*} \frac{1 - p}{p} \right) + 1 - \frac{p}{p^*} \right) \frac{c}{\lambda}.$$

Importantly, if p^* is sufficiently large relative to c, then $U_{\mathrm{R}}(p)$ (dashed line in fig. 5) stays above $\max\{U_\ell(p), U_{\mathrm{r}}(p)\}$ (solid kinked line) while p drifts toward p^*, so the receiver prefers to wait. Intuitively, unlike in the KG solution, this "more generous" persuasion scheme promises the receiver enough rents that make it worth listening to.

If c is sufficiently small, the required belief target p^* need not exceed $\hat{p}$ by much. In fact, p^* can be chosen to converge to $\hat{p}$ as $c \to 0$. In this fashion, a dynamic persuasion strategy can be constructed to approximate the KG solution when c is sufficiently small.

[16] To understand this explicit solution, first notice that under the prescribed strategy profile, the receiver takes action ℓ when p jumps to zero, which occurs with probability $(p^* - p)/p^*$, and action r when p reaches p^*, which occurs with probability p/p^*. The last term captures the total expected listening cost. The length of time τ it takes for p to reach p^* absent jumps is derived as follows:

$$p^* = \frac{p}{p + (1 - p)e^{-\lambda\tau}} \Leftrightarrow \tau = \frac{1}{\lambda} \log\left(\frac{p^*}{1 - p^*} \frac{1 - p}{p} \right).$$

Hence, the total listening cost is equal to

$$(1 - p) \int_0^\tau ctd\left(1 - e^{-\lambda t}\right) + \left(p + (1 - p)e^{-\lambda\tau}\right) c\tau = \left(p \log\left(\frac{p^*}{1 - p^*} \frac{1 - p}{p} \right) + 1 - \frac{p}{p^*} \right) \frac{c}{\lambda}.$$

At first glance, this strategy seems unlikely to work without the sender's commitment power. How can she credibly continue her experiment even after the posterior has risen past $\hat{p}$? Why not simply stop at the posterior $\hat{p}$—the belief that should have convinced the receiver to choose r? Surprisingly, however, the strategy works even without commitment. This is because the equilibrium beliefs generated by the Markov strategies themselves can provide a sufficient incentive for the sender to continue beyond $\hat{p}$. We already argued that with a suitably chosen $p^* > \hat{p}$, the receiver is incentivized to wait past $\hat{p}$ because of the optimistic equilibrium belief that the sender will continue to experiment until a much higher belief p^* is reached. Crucially, this optimism in turn incentivizes the sender to carry out her strategy:[17] were she to deviate and, say, pass at $q = \hat{p}$, the receiver would simply wait (instead of choosing r), believing that the sender will shortly resume her R-drifting experiments after the unexpected pause. Given this response, the sender cannot gain from deviating: she cannot convince the receiver to prematurely choose r. To summarize, the sender's strategy instills optimism in the receiver that makes him wait and listen, and this optimism—or the *power of beliefs*—in turn incentivizes the sender to carry out the strategy.

The power of beliefs logic extends beyond the Poisson model we employ here,[18] but it does depend on subtle details of the model. For example, consider a variation of the model in which the sender becomes unable to provide further information at some (Poisson distributed) random time. If the event is also observable to the receiver, then the above logic applies unchanged. If it is unobservable to the receiver, however, the logic no longer holds: no matter how unlikely the event is, the sender will stop providing information as soon as the belief rises above $\hat{p}$, unraveling any persuasion equilibrium. Likewise, with a deadline at which the receiver should take an action, the power of beliefs logic survives if the arrival of the deadline is stochastic but fails if the deadline is deterministic. See section VI for discussions on a few other relevant features.

[17] We will show in sec. V.B that under certain conditions, using R-drifting experiments is not just better than passing but also the optimal strategy (best response), given that the receiver waits. Here, we illustrate the possibility of persuasion for this case. The logic extends to other cases where the sender optimally uses different experiments to persuade the receiver.

[18] Consider Henry and Ottaviani's (2019) model in which the belief, as expressed by the log likelihood ratio $s = \ln(p/(1-p))$, follows a Brownian motion with a drift given by the state. In keeping with our model, suppose at each point in time that the sender either experiments or passes and the receiver chooses ℓ, r, or wait, with the flow cost c incurred on both sides if the sender experiments and the receiver waits. As noted in n. 4, this model is similar to Henry and Ottaviani's (2019) no-commitment regime, except that our sender has the option to pass without ending the game and the receiver incurs a flow cost. An MPE is then characterized by two stopping bounds, $s_* \leq \hat{s} := \ln(\hat{p}/(1-\hat{p}))$ and $s^* \geq \hat{s}$, such that the sender experiments and the receiver waits if and only if $s \in (s_*, s^*)$. Our power of beliefs argument would imply that a range of persuasion targets s^* are supported as MPE for $c > 0$ sufficiently low, and that range would span the entire $(\hat{s}, \infty)$ as $c \to 0$.

IV. Persuasion Equilibria

The equilibrium logic outlined in section III applies not just to strategy profiles that approximate the KG solution but also to other strategy profiles with a persuasion target $p^* \in (\hat{p}, 1)$. Building upon this observation, we establish a folk theorem–like result: any sender (receiver) payoff between the KG solution and full revelation can be supported as an MPE payoff in the limit as c tends to zero.

THEOREM 2. Fix any prior $p_0 \in (0, 1)$.

a. For any sender payoff $V \in (p_0 v, \min\{p_0/\hat{p}, 1\}v)$, if c is sufficiently small, there exists an MPE in which the sender obtains V.
b. For any receiver payoff $U \in (\mathcal{U}(p_0), p_0 u_r^R + (1 - p_0) u_\ell^L)$, if c is sufficiently small, there exists an MPE in which the receiver achieves U.

The proof of theorem 2 follows from the equilibrium constructions of propositions 2 and 3 in section V.B. The main argument for the proof is outlined below.

Figure 6 depicts how the set of implementable payoffs for each player varies according to p_0 in the limit as c tends to zero. Theorem 2 states that any payoffs in the shaded areas can be implemented in an MPE, provided that c is sufficiently small. In the left panel, the upper bound for the sender's payoff is given by the KG-optimal payoff $\min\{p_0/\hat{p}, 1\}v$, and the lower bound is given by the sender's payoff from full revelation $p_0 v$. For the receiver, by contrast, full revelation defines the upper bound $p_0 u_r^R + (1 - p_0) u_\ell^L$, whereas the KG payoff, which leaves no rent for the receiver, is given by $\mathcal{U}(p_0)$.

Note that theorem 2 is silent about payoffs in the dotted region. In the static KG environment, these payoffs can be achieved by the (sender-pessimal) experiment that splits the prior p into two posteriors, 1 and $q \in [0, \hat{p}]$. The following theorem shows that the sender's payoffs in this region cannot be supported as an MPE payoff for a sufficiently small $c > 0$ (even without invoking our refinement).

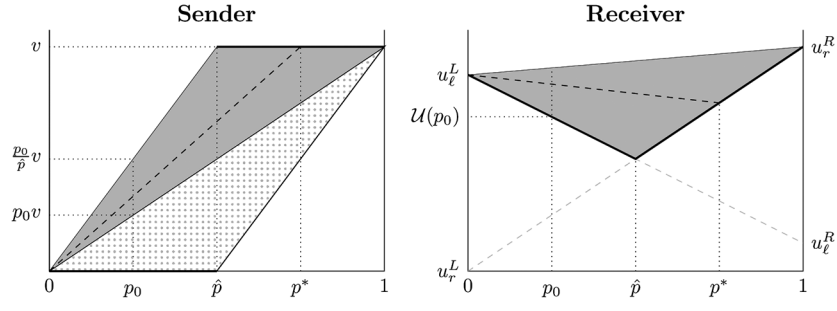

FIG. 6.—Implementable payoff set for each player at each p_0.

THEOREM 3. If $p_0 \leq \hat{p}$, then the sender's payoff in any MPE is either equal to 0 or at least $p_0 v - 2c/\lambda$. If $p_0 > \hat{p}$, then the sender's payoff in any MPE is at least $p_0 v - 2c/\lambda$.

Proof. Fix $p_0 \leq \hat{p}$, and consider any MPE. If the receiver's strategy is to wait at p_0, then the sender can always adopt the stationary strategy with jump targets 0 and 1, which will guarantee her a payoff of $p_0 v - 2c/\lambda$.[19] If the receiver's strategy is to stop at p_0, then the receiver takes action ℓ immediately, in which case the sender's payoff is equal to 0. Therefore, the sender's expected payoff is either equal to 0 or above $p_0 v - 2c/\lambda$.

Now suppose $p_0 > \hat{p}$, and consider any MPE. As above, if p_0 belongs to the waiting region, then the sender's payoff must be at least $p_0 v - 2c/\lambda$. If p belongs to the stopping region, then the sender's payoff is equal to v. In either case, the sender's payoff is at least $p_0 v - 2c/\lambda$. QED

We prove theorem 2 by constructing MPEs with a particularly simple structure:

DEFINITION 2. An MPE is a *simple MPE* (SMPE) if there exist $p_* \in (0, \hat{p})$ and $p^* \in (\hat{p}, 1)$ such that the receiver chooses action ℓ if $p < p_*$, waits if $p \in (p_*, p^*)$, and chooses action r if $p \geq p^*$.[20]

In other words, in an SMPE, the receiver waits for more information if $p \in W$ and takes an action ℓ or r otherwise, where $W = (p_*, p^*)$ or $W = [p_*, p^*)$ denotes the *waiting region*:

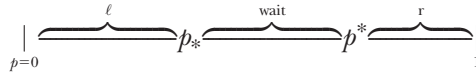

While this is the most natural equilibrium structure, we do not exclude possible MPEs that violate this structure. Whether such MPEs exist is irrelevant for our results. While we construct SMPEs to establish theorem 2, theorem 3 is valid for all MPEs. Finally, we continue to require our refinement with SMPEs.

To prove theorem 2, we begin by fixing $p^* \in (\hat{p}, 1)$. Then, for each c sufficiently small, we identify a unique value of p_* for which an SMPE can be constructed. We then show that as $c \to 0$, p_* approaches zero as well (see propositions 2 and 3 in sec. V.B). This implies that given p^*, the limit SMPE spans the sender's payoffs on the line segment that connects $(0, 0)$ and (p^*, v)—the dashed line in the left panel of figure 6—and the receiver's payoffs

[19] In order to understand this payoff, notice that the strategy fully reveals the state, and thus the sender gets v only in state R. In addition, in each state, a Poisson jump occurs at rate $\lambda/2$, and thus the expected waiting time equals $2/\lambda$, which is multiplied by c to obtain the expected cost.

[20] We do not restrict the receiver's decision at the lower bound p_*, so that the waiting region can be either (p_*, p^*) or $[p_*, p^*)$. Requiring $W = (p_*, p^*)$ can lead to nonexistence of an SMPE in proposition 2. Requiring $W = [p_*, p^*)$ can lead to nonadmissibility of the sender's best response in proposition 3.

on the line segment that connects $(0, u_\ell^{\mathrm{L}})$ and $(p^*, U_{\mathrm{r}}(p^*))$ in the right panel. By varying p^* from $\hat{p}$ to 1, we can cover the entire shaded areas in figure 6. Note that with this construction and the uniqueness claims in propositions 2 and 3, we also obtain a characterization of feasible *payoff vectors* (V, U) for the sender and receiver that can arise in an SMPE in the limit as c tends to zero. We state this in the following corollary.

COROLLARY 1. For any prior $p_0 \in [0, 1]$, in the limit as c tends to zero, the set of SMPE payoff vectors (V, U) is given by

$$\left\{(V, U) \mid \exists\, p^* \in [\max\{p_0, \hat{p}\}, 1] : V = \frac{p_0}{p^*} v,\ U = \frac{p_0}{p^*} U_{\mathrm{r}}(p^*) + \frac{p^* - p_0}{p^*} u_\ell^{\mathrm{L}}\right\},$$

with the addition of the no-persuasion payoff vector $(0, U(p_0))$ for $p_0 < \hat{p}$.

V. Persuasion Dynamics

In this section, we provide a full description of SMPE strategy profiles and illustrate the resulting equilibrium persuasion dynamics. We first explain why the sender optimally uses the three modes of persuasion discussed in sections I and II. Then, using them as building blocks, we construct full SMPE strategy profiles.

A. *Modes of Persuasion*

Fix an SMPE with two threshold beliefs p_* and p^*, where $p_* < \hat{p} < p^*$. We investigate the sender's optimal persuasion/experimentation behavior at any belief $p \in (0, 1)$ in that equilibrium.

Suppose that the sender runs a flow experiment that targets $q \neq p$ when the current belief is p. Then, by lemma 1, the belief jumps to q at rate $\lambda p(1 - p)/|q - p|$ and, absent jumps, moves continuously according to $\dot{p} = -\mathrm{sgn}(q - p)\lambda p(1 - p)$, where $\mathrm{sgn}(x)$ denotes the signum function. Therefore, her flow benefit is given by

$$v(p; q) := \lambda \frac{p(1 - p)}{|q - p|}(V(q) - V(p)) - \mathrm{sgn}(q - p)\lambda p(1 - p)V'(p),$$

where $V(\cdot)$ is the sender's value of playing the candidate equilibrium strategy.[21] Specifically, for $q > p$, the flow benefit consists of the value increase from a breakthrough that arises at rate $\lambda p(1 - p)/|q - p|$ (the first term) and the decay of value in its absence (the second term). For $q < p$, the first term captures the value decrease from a breakdown, while the second term represents the gradual appreciation in its absence.

[21] Note that the sender's value function may not be everywhere differentiable. We ignore this here to give a simplified argument illustrating the properties of the optimal strategy for the sender. The formal proofs can be found in app. C.

At each point in time, the sender can choose any countable mixture over experiments. Therefore, at each p, her flow benefit from optimal persuasion is equal to

$$v(p) \coloneqq \max_{\alpha(\cdot;p)} \sum_q \alpha(q;p) v(p;q) \text{ subject to } \sum_q \alpha(q;p) \leq 1. \tag{1}$$

The function $v(p)$ represents the gross flow value from experimentation. It plays an important role in characterizing the sender's strategy in the stopping region as well as in the waiting region. If $p \geq p^*$, then the receiver takes action r immediately, and thus $V(p) = v$ for all $p \geq p^*$. It follows that $v(p) = 0 < c$, so it is optimal for the sender to pass, which is intuitive. If $p < p_*$, then the sender has only one instant to persuade the receiver, and therefore she experiments only when $v(p) \geq c$: if $v(p) < c$, persuasion is so unlikely that she prefers to pass or, more intuitively, gives up on persuasion.

In the waiting region $p \in (p_*, p^*)$, the sender must have an incentive to experiment, which suggests that $v(p) \geq c$.[22] In particular, when the sender's equilibrium strategy involves experimentation, her value function is characterized by the Hamilton-Jacobi-Bellman (HJB) equation, which means that $V(p)$ is adjusted so that $v(p) = c$ holds.

The following proposition simplifies the potentially daunting task of characterizing the sender's optimal experiment at each belief in (1) to searching among a small subset of feasible experiments.

PROPOSITION 1. Consider an SMPE where the receiver's strategy is given by $p_* < \hat{p} < p^*$.

a. For all $p \in (0, 1)$, there exists a best response that involves at most two distinct Poisson jumps, one to $q_1 (> p)$ at rate $\alpha_1 \coloneqq \alpha(q_1; p)$ and the other to $q_2 (< p)$ at rate $\alpha_2 \coloneqq \alpha(q_2; p)$.
b. Suppose that $V(\cdot)$ is nonnegative, increasing, and strictly convex over $(p_*, p^*]$ and $V(p_*)/p_* \leq V'(p_*)$. Then, the best response in part a has
   i. for $p \in (p_*, p^*)$, $\alpha_1 + \alpha_2 = 1$ with $q_1 = p^*$ and $q_2 = 0$;
   ii. for $p < p_*$, either the sender passes or $\alpha_1 = 1$ and $q_1 = p_*$ or $q_1 = p^*$; and
   iii. for $p > p^*$, the sender passes.

For part a of proposition 1, notice that the right-hand side in equation (1) is linear in each $\alpha(q; p)$ and the constraint $\Sigma_q \alpha(q; p) \leq 1$ is also linear. Therefore, by the standard linear programming logic, there exists a

[22] Suppose that $v(p) < c$. Then, the sender strictly prefers passing forever to conducting any experiment at p followed by the optimal continuation. This implies that the value function must be $V(p) = 0$, the value of passing forever. Hence, we must have $v(p) \geq c$ whenever $V(p) > 0$, which holds if $p \in W$.

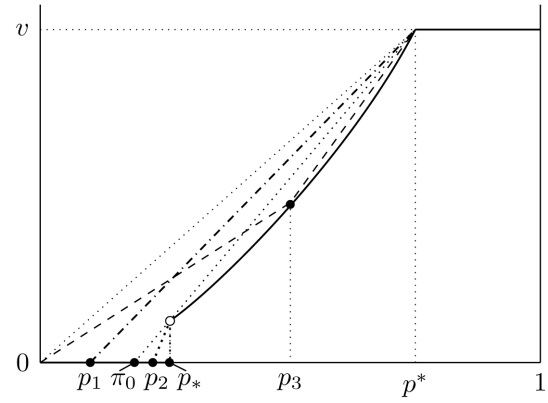

FIG. 7.—Optimal Poisson jump targets for different values of p. The solid curve represents the sender's value function in an SMPE with p_* and p^*.

solution that makes use of at most two experiments, one below p and the other above p.[23] This result implies that

$$v(p) = \max_{(\alpha_1,q_1),(\alpha_2,q_2)} \lambda p(1-p)\left[\alpha_1 \frac{V(q_1) - V(p)}{q_1 - p} - \alpha_2 \frac{V(p) - V(q_2)}{p - q_2} - (\alpha_1 - \alpha_2)V'(p)\right], \quad (2)$$

subject to $\alpha_1 + \alpha_2 \leq 1$ and $q_2 < p < q_1$.

Part b of proposition 1 states that if $V(\cdot)$ satisfies the stated properties, which will be shown to hold in equilibrium later, then there are only three candidates for optimal Poisson jump targets—0, p_*, and p^*—regardless of $p \in (0, p^*)$. As illustrated in figure 7, the right-hand side of (2) boils down to choosing $q_1 > p$ to maximize the slope of V between q_1 and p (i.e., the first fraction) or choosing $q_2 < p$ to minimize the slope of V between q_2 and p (i.e., the second fraction). In the waiting region, the former strategy leads to $q_1 = p^*$, whereas the latter strategy leads to $q_2 = 0$ (see p_3 and the dashed lines in fig. 7).[24] Similarly, if $p < p_*$, then $q_2 = 0$ is optimal and q_1 is either p_* (see p_2 and the dotted line) or p^* (see p_1 and the dash-dotted line).

Proposition 1 implies that the sender makes use of the following three modes of persuasion at each $p < p^*$.

[23] One may wonder why we allow for two experiments. In fact, linearity implies that there exists a maximizer that puts all weight on a single experiment. But to obtain an admissible Markov strategy, using two experiments is sometimes necessary. For example, if p is an absorbing belief, then admissibility requires that the stationary strategy be used at that belief, requiring two experiments. See app. B for details.

[24] Note that $q_1 > p^*$ yields a lower slope than $q_1 = p^*$; intuitively, the sender would be wasting her persuasion rate if she targets above p^*. Meanwhile, when $p \in (p_*, p^*)$, $q_2 = p_*$ yields a higher slope than $q_2 = 0$, given $V(p_*)/p_* \leq V'(p_*)$.

- Confidence building: R-drifting experiment with jump target 0.
- Confidence spending: L-drifting experiment with jump target $q_1 = p^*$ or possibly $q_1 = p_*$ if $p < p_*$.
- Confidence preserving: stationary experiment with jump targets $q_1 = p^*$ and $q_2 = 0$.

Two aspects determine the sender's choice over these experiments in her optimal strategy. First, strategies may differ in the distributions over final posteriors they induce. In particular, they may differ in the probability of persuasion (i.e., of the belief reaching p^*). Second, and more interestingly, they may differ in the time it takes for the sender to conclude persuasion. While the former feature has been studied extensively by the static persuasion models, the latter feature is novel here and is crucial for shaping the precise persuasion dynamics.

To be concrete, compare the confidence-building strategy that uses the R-drifting experiment (with jump target 0) until the belief reaches p^* with the confidence-preserving strategy that uses the stationary experiment (with jump targets $q_1 = p^*$ and $q_2 = 0$) until a jump occurs. Starting from any belief $p \in (p_*, p^*)$, both strategies eventually lead to a posterior of 0 or p^*, with identical probabilities. Hence they yield the same outcome for the two players, except for the time it takes for the persuasion process to conclude. Clearly, the sender wishes to minimize that time, which explains her choice between the two modes of persuasion. Intuitively, if the current belief is close to the persuasion target p^*, then confidence building (i.e., R-drifting) takes less time on average than confidence preserving (i.e., stationary), since the former concludes persuasion within a short period of time, whereas the latter may take a long time and thus proves costly.[25] The opposite is true, however, if the current belief is significantly away from the persuasion target p^*. Intuitively, seeking persuasion by an immediate success is more useful than slowly building up the receiver's confidence in that case.

The confidence-spending strategy (which uses the L-drifting experiment with jump target p^*) offers a similar trade-off as confidence preserving vis-à-vis confidence building. If the current belief is far away from the persuasion target p^*, confidence spending involves less time than

[25] The expected persuasion costs associated with R-drifting and stationary strategies, which can be computed as illustrated in n. 16 and 19, are given by, respectively,

$$C_+(p; p^*) = \frac{c}{\lambda}\left(p \log\left(\frac{p^*}{1-p^*}\frac{1-p}{p}\right) + 1 - \frac{p}{p^*}\right) \text{ and } C_S(p) = \frac{2c(p^* - p)}{\lambda p^*(1-p)}.$$

It can be shown that $C_+(p^*; p^*) = C_S(p^*)$ and $2C'_+(p^*; p^*) = C'_S(p^*) < 0$; i.e., as p tends to p^*, the expected persuasion cost converges to zero faster under R-drifting than under stationary strategy.

confidence building. However, there is another difference. If a success does not arise before the belief falls to p^*, persuasion stops and the receiver chooses ℓ before the belief reaches zero. By the familiar logic from (static) Bayesian persuasion, this leads to a suboptimal distribution over posteriors. To avoid this, the sender may in some cases prefer the confidence-building strategy or, in other cases, switch from the L-drifting experiment to the confidence-preserving strategy before reaching p_*. As will be seen, the confidence-spending strategy is also used in the stopping region $p < p_*$ as a Hail Mary pitch when the receiver is about to choose ℓ an instant later.

B. *Equilibrium Characterization*

We now explain how the sender's equilibrium strategy deploys the three modes of persuasion introduced in section V.A and provide a full description of the unique SMPE strategy profile for each set of parameter values and persuasion target p^*.

The structure of SMPE depends on two conditions. The first condition concerns how demanding the persuasion target p^* is:

$$p^* \leq \eta \approx 0.943. \tag{Cond1}$$

This condition determines whether the sender always prefers the R-drifting strategy to the stationary strategy. The constant η is the largest value of p^* such that the sender prefers the former strategy to the latter for all $p < p^*$ (see app. C1 for a formal definition). Notice that this condition holds for p^* not too large relative to $\hat{p}$; for instance, this is the case when the sender's equilibrium strategy approximates the KG solution (as long as $\hat{p} \leq \eta$).

The structure of the sender's equilibrium strategy also depends on the following condition:

$$v > U_{\mathrm{r}}(p^*) - U_{\ell}(p^*). \tag{Cond2}$$

The left-hand side quantifies the sender's gains when she successfully persuades the receiver and induces action r, while the right-hand side represents the corresponding gains for the receiver.[26] If (Cond2) holds, then the sender has a stronger incentive to experiment than the receiver

[26] As explained in sec. II (see n. 8), the payoffs of the two players are directly comparable, because their flow cost c is normalized to be the same. With different flow costs, (Cond2) has to be stated using each player's payoff relative to their flow cost. In the extreme case when the sender's cost is zero but the receiver's is not, (Cond2) necessarily holds, and the equilibria characterized in proposition 2 below always exist. However, the sender is indifferent over all strategies that yield the same (ex post) distribution of posteriors. Therefore, the claim of uniqueness in proposition 2 no longer holds.

has to listen, so the belief p_* below which some player wishes to stop is determined by the receiver's incentives. Conversely, if (Cond2) fails, then the sender is less eager to experiment, and thus p_* is determined by the sender's incentives.

We first provide an equilibrium characterization for the case where (Cond2) is satisfied.

PROPOSITION 2. Fix $p^* \in (\hat{p}, 1)$ and suppose that $v > U_r(p^*) - U_\ell(p^*)$. For each $c > 0$ sufficiently small, there exists a unique SMPE such that the waiting region has upper bound p^*. The waiting region is $W = [p_*, p^*)$ for some $p_* < \hat{p}$, and the sender's equilibrium strategy is as follows:[27]

a. Suppose that the belief is in the waiting region with $p \in [p_*, p^*)$.

   i. If $p^* \in (\hat{p}, \eta)$, then the sender plays the R-drifting strategy with left-jumps to zero for all $p \in [p_*, p^*)$.
   ii. If $p^* \in (\eta, 1)$,[28] then there exist cutoffs $p_* < \xi < \bar{\pi}_{LR} < p^*$ such that for $p \in [p_*, \xi) \cup (\bar{\pi}_{LR}, p^*)$, the sender plays the R-drifting strategy with left-jumps to zero; for $p = \xi$, she uses the stationary strategy with jumps to zero and p^*; and for $p \in (\xi, \bar{\pi}_{LR}]$, she adopts the L-drifting strategy with right-jumps to p^*.

b. Suppose that the belief is outside the waiting region with $p < p_*$. There exist cutoffs $0 < \pi_{\ell L} < \pi_0 < p_*$ such that for $p \le \pi_{\ell L}$, the sender passes; for $p \in (\pi_{\ell L}, \pi_0)$, she uses the L-drifting strategy with jumps to $q = p^*$; and for $p \in [\pi_0, p_*)$, she uses the L-drifting strategy with jumps to $q = p_*$.

The lower bound p_* of the waiting region converges to zero as $c \to 0$.

Figure 8 summarizes the sender's SMPE strategy in proposition 2, depending on whether $p^* < \eta$. If $p^* \in (\hat{p}, \eta)$, then the sender uses only R-drifting experiments in the waiting region $[p_*, p^*)$, as depicted in the top panel of figure 8. If $p^* > \eta$, then the sender employs other strategies as well, as described in the bottom panel of figure 8. For low beliefs close to p_*, she starts with R-drifting (confindence-building) experiments but switches to the stationary experiment when the belief reaches ξ. For beliefs

[27] We set $W = [p_*, p^*)$ to be a half-open interval, since for beliefs $p < p_*$ close to p_*, the sender's best response is to target $q = p_*$. Hence, existence of the best response requires $p_* \in W$.

[28] Notice that in the knife-edge case when $p^* = \eta$, there are two SMPEs, one as in 2a(i) and another as in 2a(ii). In the latter, however, $\bar{\pi}_{LR} = \xi$ and the L-drifting strategy is not used in the waiting region. The two equilibria are payoff-equivalent but exhibit very different dynamic behavior when $p_0 \in [p_*, \xi]$.

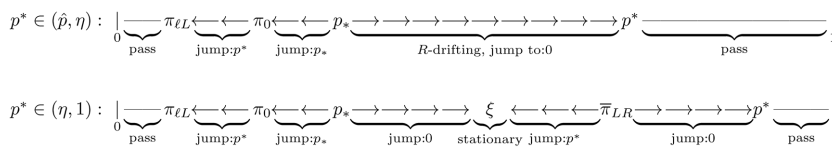

FIG. 8.—Sender's SMPE strategies in proposition 2, that is, when $v > U_r(p^*) - U_\ell(p^*)$.

above ξ but below $\bar{\pi}_{\text{LR}}$, she employs L-drifting (confidence-spending) experiments and also switches to the stationary experiment when the belief reaches ξ.

To understand these different patterns, recall from section V.A that the R-drifting experiment is particularly useful if it does not take too long to build the receiver's confidence and move the belief to p^*. This explains the use of R-drifting experiments when p is rather close to p^* for $p \in [\bar{\pi}_{\text{LR}}, p^*)$ if $p^* \geq \eta$ and for all p in the waiting region if $p^* < \eta$. If p^* is above η, then for p below $\bar{\pi}_{\text{LR}}$, other experiments become optimal. For $p < \xi$, the sender starts by building confidence, but instead of continuing with this strategy until p^* is reached, she cuts it short and switches to the stationary strategy when ξ is reached. At ξ, the arrival rate of a jump to p^* in the stationary experiment is sufficiently high to yield a faster persuasion (on average) than it would take to gradually build confidence to p^* using the R-drifting strategy. For beliefs $p \in (\xi, \bar{\pi}_{\text{LR}})$, a jump to p^* arrives at a higher rate, so that it becomes optimal to spend confidence and use only the L-drifting experiment rather than preserving confidence with the stationary experiment.

For an economic intuition, consider a salesperson courting a potentially interested buyer. If the buyer needs only a bit more reassurance to buy the product, then the salesperson should carefully build up the buyer's confidence until the belief reaches p^*. The salesperson may still slip off and lose the buyer (i.e., p jumps down to zero). But most likely, the salesperson weathers that risk and moves the buyer over the last hurdle (i.e., $q = p^*$ is reached). This is exactly what our equilibrium persuasion dynamics describes when p_0 is close to p^*. When the buyer does not require a high degree of confidence to be persuaded ($p^* \leq \eta$), building up confidence is the optimal strategy for the salesperson whenever the buyer is initially willing to listen (i.e., p_0 is in the waiting region). By contrast, when $p^* > \eta$, the buyer requires a lot of convincing and there are beliefs where the buyer is rather uninterested (as in a cold call). Then, the salesperson's optimal strategy depends on how skeptical the buyer is initially. If $p_0 \in [\bar{\pi}_{\text{LR}}, p^*)$, then it is still an optimal strategy for the salesperson to build up the buyer's confidence until p^*. If $p_0 \in (p_*, \xi)$, the salesperson first tries to build confidence. If the buyer is still listening when the belief reaches ξ, the seller becomes more

convinced that the buyer can be persuaded, and she starts using a big pitch that would move the belief to p^*. For higher beliefs, she is even more convinced that the buyer can be persuaded quickly, so she spends confidence and concentrates all her efforts on quickly persuading the receiver.

(Cond2) means that the lower bound p_* of the waiting region is determined by the receiver's incentive: p_* is the point at which the receiver is indifferent between taking action ℓ immediately and waiting (i.e., $U_\ell(p_*) = U(p_*)$, where $U(p)$ is the receiver's payoff from experimentation). Intuitively, (Cond2) suggests that the receiver gains less from experimentation—and is thus less willing to continue—than the sender. Therefore, at the lower bound p_*, the receiver wants to stop, even though the sender wants to continue persuading the receiver (i.e., $V(p_*) > 0$).

When $p < p_*$, the sender plays only L-drifting experiments unless she prefers to pass (i.e., when $p < \pi_{\ell L}$). This is intuitive, because the receiver takes action ℓ immediately unless the sender generates an instantaneous jump, forcing the sender to effectively make a Hail Mary pitch. It is intriguing, though, that the sender's target posterior can be either p_* or p^*, depending on how close p is to p_*: in the sales context used above, if the buyer is fairly skeptical, then the salesperson needs to use a big pitch. But, depending on how skeptical the buyer is, she may try to get enough attention only for the buyer to stay engaged (targeting $q = p_*$) or use an even bigger pitch to convince the buyer to buy outright (targeting $q = p^*$). If p is just below p_* (see p_2 in fig. 7), then the sender can jump into the waiting region at a high rate: recall that the arrival rate of a jump to p_* grows to infinity as p tends to p_*. In this case, it is optimal to target p_*, thereby maximizing the arrival rate of Poisson jumps: the salesperson is sufficiently optimistic about her chance of grabbing the buyer's attention, so she aims only to make the buyer stay. If p is rather far away from p_* (below π_0, such as p_1 in fig. 7), then the sender does not enjoy a high arrival rate. In this case, it is optimal to maximize the sender's payoff conditional on Poisson jumps, which she gets by targeting p^*: the salesperson tries to sell her product right away, and if it does not succeed, then she just lets it go.

Next, we provide an equilibrium characterization for the case when (Cond2) is violated.

PROPOSITION 3. Fix $p^* \in (\hat{p}, 1)$ and assume that $v \leq U_r(p^*) - U_\ell(p^*)$. For each $c > 0$ sufficiently small, there exists a unique SMPE such that the waiting region has upper bound p^*. The waiting region is $W = (p_*, p^*)$ for some $p_* < \hat{p}$, and the sender's equilibrium strategy is as follows:[29]

[29] We set $W = (p_*, p^*)$ to be an open interval, since the sender uses the L-drifting strategy for beliefs close to p_*. Including p_* would not lead to a well-defined stopping time and therefore violates admissibility.

a. Suppose that the belief is in the waiting region with $p \in (p_*, p^*)$.

   i. If $p^* \in (\hat{p}, \eta)$, then there exists a cutoff $\underline{\pi}_{\mathrm{LR}} \in W$ such that for $p \in (\underline{\pi}_{\mathrm{LR}}, p^*)$, the sender uses the R-drifting strategy with left-jumps to zero, and for $p \in (p_*, \underline{\pi}_{\mathrm{LR}})$, she uses the L-drifting strategy with right-jumps to p^*.
   ii. If $p^* \in (\eta, 1)$, then there exist cutoffs $p_* < \underline{\pi}_{\mathrm{LR}} < \xi < \bar{\pi}_{\mathrm{LR}} < p^*$ such that for $p \in [\underline{\pi}_{\mathrm{LR}}, \xi) \cup [\bar{\pi}_{\mathrm{LR}}, p^*)$, the sender plays the R-drifting strategy with left-jumps to zero; for $p = \xi$, she adopts the stationary strategy with jumps to zero or p^*; and for $p \in (p_*, \underline{\pi}_{\mathrm{LR}}) \cup (\xi, \bar{\pi}_{\mathrm{LR}})$, she uses the L-drifting strategy with right-jumps to p^*.

b. If the belief is outside the waiting region, the sender passes.

The lower bound of the waiting region p_* converges to zero as c tends to zero.

Figure 9 describes the persuasion dynamics in proposition 3. There are two main differences from proposition 2. First, if $p < p_*$, then the sender simply passes, whereas in proposition 2, the sender uses L-drifting experiments when $p \in (\pi_{\ell \mathrm{L}}, p_*)$. Second, when p is just above p_*, the sender adopts L-drifting experiments, and thus the game may stop at p_*. By contrast, in proposition 2, the sender always plays R-drifting experiments just above p_*, and the game never ends with the belief reaching p_*. Both of these differences are precisely due to the failure of (Cond2): if $v \leq U_{\mathrm{r}}(p^*) - U_{\ell}(p^*)$, then the sender is less willing to continue than the receiver, and thus p_* is determined by the sender's participation constraint (i.e., $V(p_*) = 0$). Therefore, the sender has no incentive to experiment once p falls below p_*.

When p is just above p_*, the sender goes for a big pitch by targeting p^* with L-drifting experiments. The sender does not mind losing the buyer's confidence in the process, since the violation of (Cond2) means that as the belief nears p_*, she has very little motivation left for persuading the receiver even though the latter remains willing to listen. By contrast, when (Cond2) holds (as in proposition 2), as the belief nears p_*, the receiver loses interest in listening, but the sender still sees a significant value

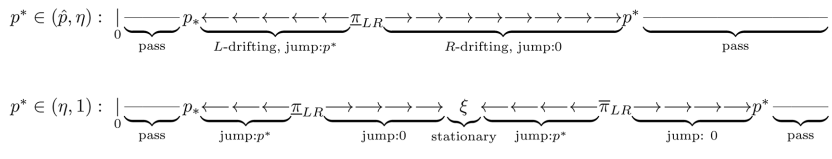

FIG. 9.—Sender's SMPE strategy in proposition 3, that is, when $v \leq U_{\mathrm{r}}(p^*) - U_{\ell}(p^*)$.

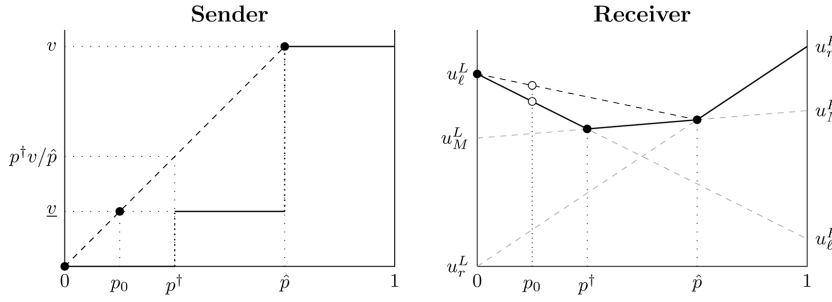

FIG. 10.—Payoffs from static persuasion when there is a middle action, *M*. *Solid lines*, payoffs without persuasion (information). *Dashed line*, sender's expected payoff in KG solution.

in staying in the game. Hence, the sender tries to build—instead of running down—the receiver's confidence in that case.

VI. Concluding Discussions

We conclude by discussing how our results depend on several modeling assumptions and suggesting a few directions for future research.

Binary actions and states.—We have considered the canonical Bayesian persuasion problem with two states and two actions. Some of our results clearly depend on specific features of the problem. However, our main economic insights hold more generally. In fact, it is often straightforward to modify our technical analysis for other persuasion problems.

To be concrete, consider an extension in which the receiver has one additional action M and the players' payoffs are as depicted in figure 10. Specifically, M is the receiver's optimal action when the belief belongs to the intermediate range $[p^\dagger, \hat{p}]$, and the sender earns $\underline{v}$ if the receiver takes action M. Assume $\underline{v} < p^\dagger v/\hat{p}$, so that the KG solution still induces two posteriors, 0 and $\hat{p}$, whenever $p_0 < \hat{p}$.

An important change from our baseline environment is that the receiver does enjoy rents from the KG solution; observe that in the right panel of figure 10, the dashed line strictly exceeds the solid line whenever $p \in (0, \hat{p})$. In this case, it is possible to construct an SMPE exactly implementing the KG solution with $p^* = \hat{p}$ for c sufficiently small. More importantly, theorem 1 no longer holds: if c is sufficiently small, then the receiver prefers to wait when the sender plays an L-drifting experiment targeting $\hat{p}$ at $p_0 (< \hat{p})$.[30] Meanwhile, theorem 2 remains valid: for c sufficiently small

[30] A tempting conjecture may be that the no-persuasion equilibrium is unsustainable if the KG solution offers strictly positive rents to the receiver. This need not be the case. If p_0 is slightly below $p^\dagger$, then it becomes credible (in the sense of satisfying our refinement) that the receiver stops immediately and the sender uses the L-drifting experiment with jump target $p^\dagger$ (and not p^*).

and p^* exceeding $\hat{p}$, the SMPE we construct for our baseline environment in section V continues to be an SMPE in this extended problem. Therefore, our arguments in section IV apply unchanged.

Relaxing the binary state assumption raises a few significant challenges—such as defining the set of feasible experiments and analyzing a system of partial differential equations—which prevents us from providing a tight and comprehensive equilibrium characterization. Nevertheless, at least conceptually, it is not hard to see how our main insights would extend to the environment with more than two states. The no-persuasion equilibrium in theorem 1 would exist if and only if the receiver never earns strictly positive (instantaneous) rents from the sender's optimal flow experiment. By contrast, if there is a (lump-sum) Blackwell experiment that strictly benefits both players, then the resulting outcome would be approximated by equilibria of the dynamic Persuasion model.

Other features of the model.—We have restricted attention to MPEs, which by definition do not rely on incentives provided by off-path punishments. Certainly, other (non-Markov) equilibria could be used so as to enlarge the set of sustainable payoffs.[31] Then, it seems plausible that as the players' persuasion costs vanish, one could implement all individually rational payoffs, including the dotted region in figure 6. For prior beliefs $p_0 < \hat{p}$, this is indeed the case, because the no-persuasion equilibrium in theorem 1 can be used to most effectively control the sender's incentives. For prior beliefs $p_0 > \hat{p}$, however, no clear punishment equilibrium is available; note that for $p_0 > \hat{p}$, the no-persuasion equilibrium maximizes the sender's payoff. This suggests that our construction of MPEs cannot be replaced by arguably more complex constructions that rely on off-path punishments. Indeed, we conjecture that for $p_0 > \hat{p}$, theorem 2 and corollary 1 characterize the full set of equilibrium payoffs.

Our model assumes flow persuasion costs rather than discounting. This assumption simplifies the analysis mainly by additively separating persuasion benefits from persuasion costs. Still, it has no qualitative impact on our main results. Specifically, if we include both flow costs and discounting in the analysis, then the resulting SMPEs would converge to those of our current model as discounting becomes negligible. If we consider only discounting (without flow costs), then the persuasion dynamics needs some modification. Among other things, the sender has no reason to voluntarily stop experimentation, and thus the persuasion dynamics will be similar to that of proposition 2 (as opposed to that of proposition 3).[32] Still, our main economic lessons will continue to apply:

[31] As is well known, it is technically challenging to define a game in continuous time without Markov restrictions (see, e.g., Simon and Stinchcombe 1989). Our subsequent discussion should be understood as referring to the limit of discrete-time equilibria.

[32] Specifically, the lower bound p_* of the waiting region will be determined by the receiver's incentives. In addition, at the lower bound p_*, so as to stay within the waiting region, the sender

all three theorems in section IV would continue to hold.[33] Furthermore, the relative advantages of the three main modes of persuasion remain unchanged, so the persuasion dynamics are in many cases similar to those described in section V.

Our continuous-time game has a straightforward discrete-time analog and can be interpreted as its limit. In a discrete-time model, however, it becomes important whether the receiver's per-period listening cost is independent of the amount of information the sender generates or proportional to it. In the former (independent) case, our power of beliefs logic no longer holds: if the current belief p is just below the persuasion target p^*, then the receiver's gains from waiting one more period are close to zero, in which case she would prefer to stop at $p < p^*$. Thus, any equilibrium with persuasion target $p^* > \hat{p}$ would unravel, leaving $p^* = \hat{p}$ as the only feasible persuasion target. In our baseline model, this renders the no-persuasion equilibrium the unique SMPE. However, if the KG solution provides positive rents for the receiver, as exemplified in the case with three actions depicted in figure 10, any persuasion target $p^* > \hat{p}$ can still be supported as SMPE as c tends to zero. Meanwhile, if the receiver's listening cost is proportional to the amount of new information, he would still be willing to wait, no matter how close p is to p^*. Then, all our analysis and results continue to hold in the discrete-time analog even in the baseline model.[34]

Our model focuses on generalized Poisson experiments to accommodate rich and flexible information choice. By contrast, an alternative such as the drift-diffusion model does not allow for such richness. For example, in Henry and Ottaviani (2019), the sender samples from a fixed exogenous process without choosing the type of experiment. Nevertheless, the logic that gives rise to our theorem 2—namely, the incentivizing power of equilibrium beliefs—applies equally well to such models (see n. 18).

Directions for future research.—The key features of our model are that real information takes time to generate and that neither the sender nor the receiver has commitment power over future actions. There are several avenues along which one could vary these features. For example, one may

---

will play either R-drifting experiments or the stationary strategy. This latter fact implies that if the game starts from $p_0 \in [p_*, p^*)$, then it will end only when the belief reaches either 0 or p^*, and thus the persuasion probability will always be equal to p_0/p^*.

[33] The proofs of theorems 1 and 3 can be readily modified. For theorem 2, it is easy to show that the main economic logic behind it (namely, the power of beliefs explained at the end of sec. III) holds unchanged with discounting.

[34] The same logic applies when there is discounting in terms of the period length Δ. If Δ is independent of the amount of new information, then all persuasive SMPEs with $p^* > \hat{p}$ unravel. However, if Δ is proportional to the amount of information—a sensible assumption if Δ describes information processing time—then such unraveling does not occur, and our analysis goes through unchanged.

consider a model in which the sender faces the same flow information constraint as in our model but has full commitment power over her dynamic strategy: given our discussion in section III, it is straightforward that the sender can approximately implement the KG outcome. However, it is nontrivial to characterize the sender's optimal dynamic strategy. Alternatively, one could further relax the commitment power by allowing the receiver to observe only the outcome of the flow experiment but not the experiment itself.

More broadly, the rich persuasion dynamics found in our model owe a great deal to the general class of Poisson experiments we allow for. At first glance, allowing for the information to be chosen from such a rich class of experiments at each point in time might appear extremely complex to analyze, and a clear analysis might seem unlikely. Yet the model produced a remarkably precise characterization of the sender's optimal choice of information—namely, not just when to stop providing information but, more importantly, what type of information to generate. This modeling innovation may fruitfully apply to other dynamic settings.

Appendix A

Further Characterization on Feasible Experiments

This appendix formally proves lemma 1 and also provides an alternative belief-based characterization for the set $\mathcal{P}^*$ of feasible experiments.

Proof of lemma 1. Fix $\langle p_t \rangle \in \mathcal{P}^*$ and any $t \in \mathbb{R}_+$. For each $q \neq p$, let $\gamma(q, p)$ denote the unconditional arrival rate of posterior belief q given $p_{t-} = p$. For these values to be well defined, it is necessary and sufficient that the associated conditional likelihoods $(\lambda^L(q, p), \lambda^R(q, p))$ satisfy

$$q = \frac{p\lambda^R(q, p)}{(1 - p)\lambda^L(q, p) + p\lambda^R(q, p)} \text{ and } \gamma(q, p) = p\lambda^R(q, p) + (1 - p)\lambda^L(q, p).$$

Solving this system of equations, we obtain

$$\lambda^R(q, p) = \gamma(q, p)\frac{q}{p} \text{ and } \lambda^L(q, p) = \gamma(q, p)\frac{1 - q}{1 - p}.$$

Then, our information constraint can be written as

$$\sum_{q \neq p} |\lambda^R(q, p) - \lambda^L(q, p)| = \sum_{q \neq p} \gamma(q, p)\left|\frac{q}{p} - \frac{1 - q}{1 - p}\right| = \sum_{q \neq p} \gamma(q, p)\frac{|q - p|}{p(1 - p)} \leq \lambda. \quad \text{(A1)}$$

For each q, define $\alpha(q) := \gamma(q, p)|q - p|/[\lambda p(1 - p)]$. Then, the above constraint can be equivalently written as $\Sigma_{q \neq p}\alpha(q) \leq 1$, and the arrival rate of posterior q given p is given by $\gamma(q, p) = \alpha(q)\lambda p(1 - p)/|q - p|$.

Let $\dot{p}$ denote the instantaneous change of $\langle p_t \rangle$ conditional on no jump. Since $\langle p_t \rangle$ is a martingale, $\Sigma_{q \neq p}\gamma(q, p)(q - p) + \dot{p} = 0$, so $\dot{p}$ satisfies

$$\begin{aligned}\dot{p} &= -\sum_{q\neq p}\gamma(q,p)(q-p) = -\sum_{q>p}\gamma(q,p)(q-p) - \sum_{q<p}\gamma(q,p)(q-p)\\ &= -\sum_{q>p}\gamma(q,p)|q-p| + \sum_{q<p}\gamma(q,p)|q-p|\\ &= -\sum_{q>p}\alpha(q)\lambda p(1-p) + \sum_{q<p}\alpha(q)\lambda p(1-p)\\ &= -\left(\sum_{q>p}\alpha(q) - \sum_{q<p}\alpha(q)\right)\lambda p(1-p).\end{aligned}$$

QED

Next, we provide an additional characterization of $\mathcal{P}^*$ based on a measure of information. Let $\langle p_t\rangle$ denote a regular martingale process in $\mathcal{P}$. For each t, $p_{t-} \coloneqq \lim_{t'\uparrow t} p_{t'}$, and $q \neq p_{t-}$, let $\gamma(q, p_{t-})$ denote the rate at which the belief jumps from p_{t-} to q; formally,

$$\gamma(q, p_{t-}) \coloneqq \lim_{dt\to 0}\frac{\mathbb{P}[p_t = q|p_{t-dt}]}{dt}.$$

We measure the amount of flow information of $\langle p_t\rangle$ at each point in history by

$$\mathcal{I}(p_{t-}) \coloneqq \sum_{q\neq p_{t-}}\gamma(q, p_{t-})|p_t - p_{t-}|.$$

In other words, our information measure $\mathcal{I}(p_{t-})$ quantifies the total absolute change of the belief process at each point in time.

By (A1) in the proof of lemma 1, our information constraint can be written as

$$\sum_{q\neq p}|\lambda^{\mathrm{R}}(q,p) - \lambda^{\mathrm{L}}(q,p)| = \sum_{q\neq p}\gamma(q,p)\frac{|q-p|}{p(1-p)} \leq \lambda \Leftrightarrow \mathcal{I}(p) \leq \lambda p(1-p).$$

This implies that the set $\mathcal{P}^*$ of feasible experiments can be equivalently defined as

$$\mathcal{P}^* \coloneqq \{\langle p_t\rangle \in \mathcal{P} : \mathcal{I}(p_{t-}) \leq \lambda p_{t-}(1-p_{t-}) \text{ for all } t \text{ and } p_{t-}\}.$$

In other words, we consider belief processes whose aggregate change at each point in history is bounded by $\lambda p_{t-}(1-p_{t-})$; note that $p(1-p)$ is equal to the variance of the Bernoulli random variable p. The bound's dependence on p_{t-} is natural, given that $p_{t-} \in [0,1]$ and $p_{t-} = 0, 1$ represents perfect information from which no belief change should be feasible; more generally, it captures an intuitive idea that the sender can move the receiver's belief more, the more uncertain the state is.

Appendix B

Admissible Strategies

This appendix completes the definition of our continuous-time game by defining admissible strategies for the sender. We note that this appendix is similar to

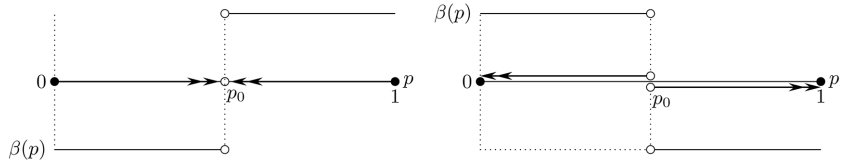

FIG. 11.—*Left*, case where integral equation (B1) does not have a solution. *Right*, case where (B1) has multiple solutions.

appendix B.1 of Klein and Rady (2011): the two models have the same underlying technical issues and natural resolutions to them.[35]

Recall that the sender's strategy is a measurable function σ^S that assigns a flow experiment $\sigma^S(p) = (\alpha(q; p))_{q\in[0,1]}$ to each belief $p \in [0, 1]$. As noted, the strategy induces a belief process satisfying

$$\dot{p}_t = -\beta(p_t)\lambda p_t(1 - p_t), \tag{B1}$$

where

$$\beta(p) \coloneqq \sum_{q>p}\alpha(q; p) - \sum_{q<p}\alpha(q; p).$$

Note that p_t moves leftward if $\beta(p_t) > 0$ and rightward if $\beta(p_t) < 0$.

DEFINITION 3. A measurable function σ^S is an admissible strategy for the sender if for all $p_0 \in [0, 1]$, there exists a solution to (B1).

To see the role of definition 3, first observe that for σ^S with a relatively simple structure, we can find an explicit solution to (B1). For example, if the sender plays only the R-drifting experiment, then $\beta(p) = -1$ for all p, in which case $p_t = p_0 e^{\lambda t}/(p_0 e^{\lambda t} + 1 - p_0)$. If the sender plays only the stationary experiment, then $\beta(p) = 0$ for all p, in which case $p_t = p_0$. Of course, the differential equation (B1) cannot be solved explicitly in general. One may utilize a sufficient condition on $\beta(\cdot)$: for example, it suffices that $\beta(\cdot)$ is continuous or satisfies Carathéodory conditions (see Goodman 1970). For our purpose, however, imposing such a sufficient condition is unnecessarily restrictive. Therefore, we require only that there is a solution to (B1). More precisely, we shall require a couple of conditions, one of which is necessary and the other is of no material consequence. This approach is valid, since the equilibrium with these weaker conditions will ensure that (B1) is well defined for all $p \in [0, 1]$.

To explain the necessary condition that is relevant for our context, consider, for example, a strategy such that the sender plays the R-drifting experiment targeting 0 (so $\beta(p) = -1$) whenever $p \le p_0$ and the L-drifting experiment targeting 1 (so $\beta(p) = 1$) whenever $p > p_0$. As depicted in the left panel of figure 11, the belief moves toward p whether it is below or above p, so $\beta(p_0) = -1$ results

[35] The difference is that the technical problems arise in their model because the evolution of beliefs is jointly controlled by two players, while in our model, it is because the sender can choose from a large set of Poisson experiments.

in (B1) being ill defined at p_0. In fact, p_t should stay constant if starting from p_0. Hence, admissibility requires $\sigma^S(p_0)$ to satisfy $\beta(p_0) = 0$.

We next consider a condition that is not necessary for (B1) to be well defined but is sensible as a selection rule when (B1) admits multiple solutions. Consider, for example, a strategy such that the sender plays the L-drifting experiment targeting 1 (so $\beta(p) = 1$) whenever $p \leq p_0$ and the R-drifting experiment targeting 0 (so $\beta(p) = -1$) whenever $p > p_0$ (see the right panel of fig. 11). Since the former case includes p_0, it is natural that starting from p_0, the belief moves leftward according to

$$p_t = \frac{pe^{-\lambda t}}{pe^{-\lambda t} + (1 - p)}.$$

However, since $p_t = p_0$ only when $t = 0$, the following is also a solution to (B1):

$$p_t = \frac{pe^{\lambda t}}{pe^{\lambda t} + (1 - p)}.$$

Whenever this multiplicity arises, we select the most natural one that would be obtained from the discrete-time approximation. This selection, however, is inconsequential for our equilibrium characterization, because at a point where this selection issue arises (such as $\underline{\pi}_{LR}$ or $\bar{\pi}_{LR}$ in propositions 2 and 3), we can arbitrarily specify the sender's strategy; the selection forces us to adopt a particular belief path but does not restrict the sender's strategy in any way.

Appendix C

Proofs of Propositions 2 and 3

The proofs are presented in several sections. Throughout, we take $p^* \in (\hat{p}, 1)$ as given and construct the corresponding equilibria. Section C1 constructs the value functions that correspond to the equilibrium strategies in propositions 2 and 3. Sections C2 and C3 verify the sender's and the receiver's incentives, respectively. Uniqueness of SMPE is proven in the supplemental material. A brief sketch is provided in section C4.

C1. Constructing Equilibrium Value Functions

We first compute the players' value functions under alternative persuasion strategies; they will be used to compute the players' equilibrium payoffs. In what follows, we take it for granted that the receiver takes an action immediately if the belief reaches either 0 or p^*.

We also assume that the receiver waits while the sender plays each persuasion strategy in this section.

Ordinary differential equations (ODEs) for R-drifting and L-drifting.—For any $p \in (0, p^*)$, let $N_\varepsilon(p)$ denote a small open neighborhood of p. Suppose that for any belief in $N_\varepsilon(p)$, the sender plays the R-drifting experiment with jump target 0. Then, the sender's value function $V_+(p)$ and the receiver's value function $U_+(p)$ satisfy the following ODEs:[36]

[36] The ODEs can be obtained heuristically in the same way as the Hamilton-Jacobi-Bellman equation. The subscripts + and − represent the direction of belief drifting in the absence of Poisson jumps.

$$c = \lambda p(1-p)\left(\frac{-V_+(p)}{p} + V_+'(p)\right) \quad \text{and} \quad c = \lambda p(1-p)\left(\frac{u_\ell^{\mathrm{L}} - U_+(p)}{p} + U_+'(p)\right). \tag{C1}$$

Similarly, suppose that for any belief in $N_\varepsilon(p)$, the sender plays the L-drifting experiment with jump target p^*. Then, the players' value functions, $V_-(p)$ and $U_-(p)$, satisfy

$$c = \lambda p(1-p)\left(\frac{v - V_-(p)}{p^* - p} - V_-'(p)\right) \text{ and } c = \lambda p(1-p)\left(\frac{U_{\mathrm{r}}(p^*) - U_-(p)}{p^* - p} - U_-'(p)\right). \tag{C2}$$

R-drifting strategy.—Suppose that the sender plays R-drifting experiments until the belief reaches p^*. In this case, the players' payoffs are obtained as the solutions to (C1) with boundary conditions $V_+(p^*) = v$ and $U_+(p^*) = U_{\mathrm{r}}(p^*)$, respectively. We obtain

$$V_{\mathrm{R}}(p) = \frac{p}{p^*}v - C_+(p; p^*) \text{ and } U_{\mathrm{R}}(p) = \frac{p^* - p}{p^*}u_\ell^{\mathrm{L}} + \frac{p}{p^*}U_{\mathrm{r}}(p^*) - C_+(p; p^*),$$

where $C_+(p; q) \coloneqq [p\log([q/(1-q)][(1-p)/p]) + 1 - (p/q)](c/\lambda)$ represents the expected cost of using R-drifting experiments until the belief moves from p to either 0 or q.

Stationary strategy.—Suppose that the sender uses the stationary experiment with jump targets 0 and p^* at p. Then, the players' value functions, $V_{\mathrm{S}}(p)$ and $U_{\mathrm{S}}(p)$, are respectively given by

$$V_{\mathrm{S}}(p) = \frac{p}{p^*}v - C_{\mathrm{S}}(p) \text{ and } U_{\mathrm{S}}(p) = \frac{p^* - p}{p^*}u_\ell^{\mathrm{L}} + \frac{p}{p^*}U_{\mathrm{r}}(p^*) - C_{\mathrm{S}}(p), \tag{C3}$$

where $C_{\mathrm{S}}(p) \coloneqq 2c(p^* - p)/[\lambda p^*(1-p)]$ represents the expected cost of playing the stationary strategy.[37]

RS strategy (R-drifting followed by stationary).—Suppose that the sender plays the R-drifting strategy until $q(> p)$ and then switches to the stationary strategy. Then, the players' value functions solve (C1) with boundary conditions $V_+(q) = V_{\mathrm{S}}(q)$ and $U_+(q) = U_{\mathrm{S}}(q)$, yielding

$$V_{\mathrm{RS}}(p; q) = \frac{p}{p^*}v - C_+(p; q) - \frac{p}{q}C_{\mathrm{S}}(q) \text{ and}$$

$$U_{\mathrm{RS}}(p; q) = \frac{p^* - p}{p^*}u_\ell^{\mathrm{L}} + \frac{p}{p^*}U_{\mathrm{r}}(p^*) - C_+(p; q) - \frac{p}{q}C_{\mathrm{S}}(q).$$

Note that p/q is the probability that the belief moves from p to q (whereupon the sender switches to the stationary strategy).

LS strategy (L-drifting followed by stationary).—Suppose that the sender plays the L-drifting strategy until $q(< p)$ and then switches to the stationary strategy. Then, the players' value functions solve (C2) with boundary conditions $V_-(q) = V_{\mathrm{S}}(q)$ and $U_-(q) = U_{\mathrm{S}}(q)$, resulting in

[37] Under the stationary strategy, the total arrival rate of Poisson jumps is equal to $\lambda_{\mathrm{S}}(p) = (\lambda/2)(1-p) + (\lambda/2)p(1-p)/(p^* - p) = (\lambda/2)p^*(1-p)/(p^* - p)$. $C_{\mathrm{S}}(p)$ is equal to c times the expected arrival time $1/\lambda_{\mathrm{S}}(p)$.

$$V_{\mathrm{LS}}(p;q) = \frac{p}{p^*}v - C_-(p;q) - \frac{p^*-p}{p^*-q}C_{\mathrm{S}}(q) \text{ and}$$

$$U_{\mathrm{LS}}(p;q) = \frac{p^*-p}{p^*}u_{\mathrm{r}}^{\mathrm{L}} + \frac{p}{p^*}U_{\mathrm{r}}(p^*) - C_-(p;q) - \frac{p^*-p}{p^*-q}C_{\mathrm{S}}(q),$$

where

$$C_-(p;q) := -\frac{p^*-p}{p^*(1-p^*)}\left(p^*\log\frac{1-q}{1-p} + (1-p^*)\log\frac{q}{p} - \log\frac{p^*-q}{p^*-p}\right)\frac{c}{\lambda}$$

denotes the expected cost of playing L-drifting experiments until the belief drifts down from p to $q(<p)$.

Crossing lemma.—The following lemma provides potential crossing patterns among the value functions and plays a crucial role in the subsequent analysis.

LEMMA 2 (Crossing lemma). Let $V_+(p)$ and $V_-(p)$ be solutions to (C1) and (C2), respectively.

a. Let $p^* < 8/9$. For all $p < p^*$, if $V_+(p) = V_{\mathrm{S}}(p)$, then $V_+'(p) < V_{\mathrm{S}}'(p)$. Similarly, if $V_-(p) = V_{\mathrm{S}}(p)$, then $V_-'(p) < V_{\mathrm{S}}'(p)$.

b. Let $p^* \geq 8/9$, and define $\xi_1 := 3p^*/4 - \sqrt{(3p^*/4)^2 - (p^*/2)}$ and $\xi_2 := 3p^*/4 + \sqrt{(3p^*/4)^2 - (p^*/2)}$.

   i. For all $p < p^*$, if $V_+(p) = V_{\mathrm{S}}(p)$, then $V_+'(p) = V_{\mathrm{S}}'(p)$ if and only if $p \in \{\xi_1, \xi_2\}$, and $V_+'(p) > V_{\mathrm{S}}'(p)$ if and only if $p \in (\xi_1, \xi_2)$.

   ii. For all $p < p^*$, if $V_-(p) = V_{\mathrm{S}}(p)$, then $V_-'(p) = V_{\mathrm{S}}'(p)$ if and only if $p \in \{\xi_1, \xi_2\}$, and $V_-'(p) > V_{\mathrm{S}}'(p)$ if and only if $p \in (\xi_1, \xi_2)$.

c. For all $p < p^*$, if $V_+(p) = V_-(p)$, then $\operatorname{sign}(V_+'(p) - V_-'(p)) = \operatorname{sign}(V_-(p) - V_{\mathrm{S}}(p))$.

All parts also hold for the receiver's value functions $U_+(\cdot)$, $U_-(\cdot)$, and $U_{\mathrm{S}}(\cdot)$.

Proof. We focus on the sender's value functions, as the same proofs apply to the receiver. From (C1)–(C3), we can obtain expressions for $V_+'(p)$, $V_-'(p)$, and $V_{\mathrm{S}}'(p)$. Combining these with $V_+(p) = V_{\mathrm{S}}(p)$ and $V_-(p) = V_{\mathrm{S}}(p)$, we obtain

$$V_+'(p) - V_{\mathrm{S}}'(p) = V_-'(p) - V_{\mathrm{S}}'(p) = -\frac{c(2p^2 - 3p^*p + p^*)}{\lambda p^* p(1-p)^2} \gtreqless 0$$
$$\Leftrightarrow -2p^2 + 3p^*p - p^* \gtreqless 0.$$

For $p^* < 8/9$, the quadratic expression in the last inequality is always negative, which proves part a. For $p^* \geq 8/9$, the quadratic expression has two real roots, ξ_1 and ξ_2, and is positive if and only if $p \in (\xi_1, \xi_2)$. This proves part b.

Similarly, using $V_+(p) = V_-(p)$, we have

$$V_+'(p) - V_-'(p) = \frac{p^*}{p(p^*-p)}(V_-(p) - V_{\mathrm{S}}(p)),$$

which leads to part c. QED

Construction of ξ.—While ξ is part of the equilibrium only for $p^* > \eta$, we define it generally. For $p^* \geq 8/9$, we set $\xi := \xi_1$, and for $p^* < 8/9$, we set $\xi := p^*$. We define it in this way to ensure that $V_{\mathrm{RS}}(p;\xi)$ meets $V_{\mathrm{S}}(p)$ from above at $p = \xi$ (as p rises toward ξ). In particular, together with the crossing lemma 2b, this means

that for any $p < \xi$, $V_{RS}(p; \xi)$ is above $V_S(p)$, and for $p^* \geq 8/9$, these two functions have the same slope at $p = \xi$. This will play a crucial role later.

Construction of η.—The parameter η is the value of $p^* \geq 8/9$ such that $V_R(\xi(p^*)) = V_S(\xi(p^*))$.[38] (We make the dependence of ξ on p^* explicit here and also note that the functions $V_R(\cdot)$ and $V_S(\cdot)$ depend on p^* directly.) Solving this equation yields $p^* = \eta \approx 0.943$. We make the following observations for a later purpose:

LEMMA 3.

a. If $p^* < \eta$, then $V_R(p) > V_S(p)$ for all $p \in (0, p^*)$.
b. If $p^* = \eta$, then $V_R(p) \geq V_S(p)$ for all $p \in (0, p^*)$, with equality only when $p = \xi$.
c. If $p^* > \eta$, then $V_R(\xi) < V_S(\xi)$.

The same results hold for $U_R(\cdot)$ and $U_S(\cdot)$.

Proof. We focus on the sender's value functions, as the same proofs apply to the receiver. Using the explicit solutions of $V_R(p)$ and $V_S(p)$, we can see that $V_S(0) < V_R(0)$, $V_S(p^*) = V_R(p^*)$, and $V_S'(p^*) > V_R'(p^*)$. Therefore, either $V_S(p)$ stays weakly below $V_R(p)$ for all $p < p^*$ or $V_S(p)$ crosses $V_R(p)$ at least twice (from below and then from above). By lemma 2b, the latter occurs only if $V_S(p)$ crosses $V_R(p)$ from below at some $p < \xi$ and then second time from above at some $p' \in (\xi, \xi_2)$, which is equivalent to $V_R(\xi) < V_S(\xi)$. The desired result follows since $V_R(\xi(p^*)) - V_S(\xi(p^*))$ changes the sign only once at $p^* = \eta$ (see n. 38). QED

Pasted strategies.—Given ξ, we combine alternative strategies as follows. For any $p \leq p^*$, we define

$$\hat{V}(p) := \begin{cases} V_{RS}(p; \xi) & \text{if } p < \xi, \\ V_S(\xi) & \text{if } p = \xi, \\ V_{LS}(p; \xi) & \text{if } p \in [\xi, p^*], \end{cases} \quad \text{and} \quad \hat{U}(p) := \begin{cases} U_{RS}(p; \xi) & \text{if } p < \xi, \\ U_S(\xi) & \text{if } p = \xi, \\ U_{LS}(p; \xi) & \text{if } p \in [\xi, p^*]. \end{cases}$$

We next define $\tilde{V}(p) := \max\{V_R(p), \hat{V}(p)\}$ and $\tilde{U}(p) := \max\{U_R(p), \hat{U}(p)\}$. We make several useful observations in the following lemma.

LEMMA 4.

a. Both $\tilde{V}(p)$ and $\tilde{U}(p)$ are strictly convex in p over $[0, p^*]$.
b. If $p^* \leq \eta$, then $\tilde{V}(p) = V_R(p)$ and $\tilde{U}(p) = U_R(p)$ for all $p \in [0, p^*]$.
c. If $p^* > \eta$, then there exists $\bar{\pi}_{LR} \in (\xi, p^*)$ such that $\tilde{V}(p) = \hat{V}(p)$ and $\tilde{U}(p) = \hat{U}(p)$ for $p \leq \bar{\pi}_{LR}$ and $\tilde{V}(p) = V_R(p)$ and $\tilde{U}(p) = U_R(p)$ for $p \in [\bar{\pi}_{LR}, p^*]$.
d. $\tilde{V}(p) \geq V_S(p)$ for all $p < p^*$, and the inequality is strict for $p \neq \xi$.

Proof. The same proof applies to both players, so we focus on the sender's value functions. Recall that for $p^* < 8/9$, we have $\xi = p^*$ so that $\hat{V}(p) = V_{RS}(p; p^*) = V_R(p)$, which implies part b. Since $V_R(p)$ is strictly convex, part a holds as well. In what follows, we consider $p^* \geq 8/9$, in which case $\hat{V}(p) \neq V_R(p)$.

[38] To show that η is well defined, we can define a function $g:(8/9, 1) \to \mathbb{R}$ by $g(p^*) := V_R(\xi(p^*)) - V_S(\xi(p^*))$ so that $g(\eta) = 0$. It can be verified that $g'(p^*) > 0$ for all $p^* \in (8/9, 1)$.

a. Since $\bar{V}(p)$ is the upper envelope of two functions and $V_R(p)$ is strictly convex over $[0, p^*]$, it suffices to prove that $\hat{V}(p)$ is also strictly convex over $[0, p^*]$. Both $V_{RS}(p; \xi)$ and $V_{LS}(p; \xi)$ are strictly convex over their respective supports, and $\hat{V}(p)$ is continuously differentiable at the pasting point ξ. The latter holds because $V_{RS}(\xi; \xi) = V_{LS}(\xi; \xi) = V_S(\xi)$ implies $V'_{RS}(\xi; \xi) = V'_{LS}(\xi; \xi)$ by lemma 2c.

b. If $p^* < \eta$, $V_{RS}(\xi; \xi) = V_S(\xi) < V_R(\xi)$ by lemma 3a. Together with the fact that both $V_{RS}(p; \xi)$ and $V_R(p)$ satisfy the ODE (C1), this implies that $\hat{V}(p) = V_{RS}(p; \xi) < V_R(p)$ for all $p \leq \xi$.[39] For $p \in (\xi, p^*]$, observe that $V_{LS}(\xi; \xi) = V_S(\xi) < V_R(\xi)$ (lemma 3a), $V_{LS}(p^*; \xi) = V_R(p^*)$, and $V'_{LS}(p^*; \xi) > V'_R(p^*)$. Therefore, either $\hat{V}(p) = V_{LS}(p; \xi) < V_R(p)$ for all $p \in (\xi, p^*)$ or $V_{LS}(\cdot; \xi)$ crosses $V_R(\cdot)$ from below at least once at some $p \in (\xi, p^*)$. In the latter case, we must have $V'_R(p) = V'_+(p) < V'_-(p) = V'_{LS}(p; \xi)$. Then, by lemma 2c, $V_R(p) = V_{LS}(p; \xi) = V_-(p) < V_S(p)$, contradicting lemma 3a.

The result for $p^* = \eta$ follows from a continuity argument: both $\hat{V}(p)$ and $V_R(\cdot)$ change continuously in p^*. Since $\hat{V}(p) < V_R(p)$ for all $p < p^*$ whenever $p^* < \eta$, it must be that $\hat{V}(p) \leq V_R(p)$ for all $p < p^*$ when $p^* = \eta$. This concludes the proof for part b.

For parts c and d, the following claim is useful:

Claim 1. Suppose $p^* \geq 8/9$.

i. $\hat{V}(p) \geq V_S(p)$ for all $p \in (0, \xi_2]$, with strict inequality for $p \neq \xi$.
ii. $V_R(p) > V_S(p)$ for all $p \in [\xi_2, p^*)$.

Proof. i. Consider first $p < \xi (= \xi_1)$. We have to show that $\hat{V}(p) = V_{RS}(p; \xi) > V_S(p)$. To see this, pick $q < \xi$. Then by lemma 2b(i), $V_{RS}(p; q)$ stays above $V_S(p)$ for $p < q$ and $V_{RS}(p; \xi) > V_{RS}(p; q)$ for all $q < \xi$. The same logic applies to $V_{LS}(p; \xi)$ for $p \in (\xi, \xi_2]$. For part ii, we check that $V_R(p^*) = V_S(p^*)$ and $V'_R(p^*) < V'_S(p^*)$. Lemma 2b(i) then implies that $V_R(p)$ and $V_S(p)$ cannot intersect at $p \geq \xi_2$. QED

For part c, we first show that $V_{RS}(p; \xi) > V_R(p)$ for $p \leq \xi$. If $p^* > \eta$, then $V_{RS}(\xi; \xi) = V_S(\xi) > V_R(\xi)$ (lemma 3c), which immediately implies that $\hat{V}(p) = V_{RS}(p; \xi) > V_R(p)$ for all $p \leq \xi$. Next, for $p \in (\xi, p^*]$, observe that $V_{LS}(\xi; \xi) = V_S(\xi) > V_R(\xi)$; and $V_{LS}(p; \xi) < V_R(p)$ for $p = p^* - \varepsilon$, since $V_{LS}(p^*; \xi) = V_R(p^*)$, and $V'_{LS}(p^*; \xi) > V'_R(p^*)$. This means that $V_{LS}(\cdot; \xi)$ crosses $V_R(\cdot)$ at least once in (ξ, p^*). To show that there is a unique crossing point $\bar{\pi}_{LR}$, note that claim 1 implies that at any crossing point $p \in (\xi, p^*)$, $V_{LS}(p; \xi) = \hat{V}(p) = V_R(p) > V_S(p)$, and hence by lemma 2c, $V_{LS}(p; \xi)$ can cross $V_R(p)$ only from above. Therefore, there is a unique crossing point.

d. If $p^* \leq \eta$, then the result is immediate from lemmas 3a and 4b. If $p^* > \eta$ the result is immediate from lemma 4c and claim 1. QED

C1.1. Equilibrium Payoffs and Construction of p_* in Proposition 2

When (Cond2) holds, we define p_* as the belief $\phi_{\ell R}$ at which the receiver is indifferent between waiting and stopping with action ℓ; that is, we set $p_* := \phi_{\ell R}$, where $\phi_{\ell R}$ is defined by[40]

[39] It is easy to see that (C1) satisfies the Lipschitz condition for uniqueness on $(0, p^*)$.

[40] To see that $\phi_{\ell R}$ is well defined, observe that, whether $p^* \leq \eta$ or $p^* > \eta$, $\lim_{p \to 0} \tilde{U}(p) = u^1_\ell - c/\lambda < u^1_\ell = U_\ell(0)$, while $\tilde{U}(p^*) = U_r(p^*) > U_\ell(p^*)$ (because $p^* > \hat{p}$). In addition, $\tilde{U}(p)$ is strictly convex over $[0, p^*]$ (lemma 4a), while $U_\ell(p)$ is linear. Therefore, $\tilde{U}(p)$ crosses $U_\ell(p)$ from below only once.

$$U_\ell(\phi_{\ell R}) = \tilde{U}(\phi_{\ell R}). \tag{C4}$$

We focus on the case in which c is sufficiently small. In the limit as $c \to 0$, $\tilde{U}(p) = [(p^* - p)/p^*]u_\ell^L + (p/p^*)U_r(p^*) > U_r(p)$ for all p. Therefore, there exists $c_1 > 0$ such that $p_* = \phi_{\ell R} < \hat{p}$ for all $c \le c_1$. We assume that $c \le c_1$ in the sequel. The following lemma shows that the sender's payoff is positive at p_* if (Cond2) holds.

LEMMA 5. $\tilde{V}(\phi_{\ell R}) > 0$ if and only if (Cond2) holds.

Proof. By (C4), we have

$$\tilde{V}(\phi_{\ell R}) = \frac{\phi_{\ell R}}{p^*} v + \tilde{U}(\phi_{\ell R}) - \left(\frac{p^* - \phi_{\ell R}}{p^*} u_\ell^L + \frac{\phi_{\ell R}}{p^*} U_r(p^*)\right) \overset{(C4)}{=} \frac{\phi_{\ell R}}{p^*}(v - (U_r(p^*) - U_\ell(p^*)),$$

where the first equality holds because both players incur the same costs, so that $\tilde{V}(p) - (p/p^*)v = \tilde{U}(p) - \{[(p^* - p)/p^*]u_\ell^L + (p/p^*)U_r(p^*)\}$ whenever $p \in (0, p^*]$. The last expression is positive if and only if (Cond2) holds. QED

We set the players' value functions as follows:

$$V(p) := \begin{cases} 0 & \text{if } p \in [0, p_*), \\ \tilde{V}(p) & \text{if } p \in [p_*, p^*), \\ v & \text{if } p \ge p^*, \end{cases} \quad \text{and} \quad U(p) := \begin{cases} U_\ell(p) & \text{if } p \in [0, p_*), \\ \tilde{U}(p) & \text{if } p \in [p_*, p^*), \\ U_r(p) & \text{if } p \ge p^*. \end{cases}$$

LEMMA 6. When (Cond2) holds, $V(p)$ is nonnegative and nondecreasing for all $p \in [0, 1]$.

Proof. Since $\tilde{V}(\cdot)$ is convex on $[0, p^*]$, $\tilde{V}(0) = -c/\lambda$, and $\tilde{V}(p_*) \ge 0$ by lemma 5, $\tilde{V}(\cdot)$ is increasing on $[p_*, p^*]$. Hence $V(\cdot)$ is nondecreasing on $[0, 1]$ and nonnegative since $V(0) = 0$. QED

C1.2. Equilibrium Payoffs and Construction of p_* in Proposition 3

When (Cond2) fails, the same construction as above does not work; for example, $\tilde{V}(p_*) < 0$ by lemma 5. The right construction requires us to consider another L-drifting strategy.

L0 strategy (L-drifting followed by passing).—Suppose that the sender continues to play the L-drifting experiment until the belief reaches $q (< p)$ and then she stops experimenting altogether (passes). The resulting value functions are the solutions to (C2) with boundary conditions $V_-(q) = 0$ and $U_-(q) = U_\ell(q)$, which yield

$$V_{L0}(p; q) := \frac{p - q}{p^* - q} v - C_-(p; q) \text{ and}$$

$$U_{L0}(p; q) := \frac{p^* - p}{p^* - q} U_\ell(q) + \frac{p - q}{p^* - q} U_r(p^*) - C_-(p; q).$$

Note that this strategy leads to q with probability $(p^* - p)/(p^* - q)$ and p^* with probability $(p - q)/(p^* - q)$.

Construction of p_*.—Let $\pi_{\ell L}$ denote the lowest value of $q \in (0, \hat{p})$ such that

$$V'_{L0}(q; q) \geq 0 \Leftrightarrow \frac{\lambda q(1-q)}{p^* - q} v \geq c \Leftrightarrow q \geq \pi_{\ell L} := \frac{1}{2} + \frac{c}{2\lambda v} - \sqrt{\left(\frac{1}{2} + \frac{c}{2\lambda v}\right)^2 - \frac{cp^*}{\lambda v}}. \quad \text{(C5)}$$

In words, $\pi_{\ell L}$ is the lowest belief at which the sender is willing to play the L0 strategy even for an instance. When (Cond2) fails, we set $p_* := \pi_{\ell L}$. Clearly, $\lim_{c \to 0} p_* = 0$. We set $c_2 > 0$ such that $p_* = \pi_{\ell L} < \hat{p}$ for all $c \leq c_2$ and assume $c \leq c_2$ hereafter.

LEMMA 7. Suppose that (Cond2) fails and $p_* = \pi_{\ell L}$. There exists $c_3 > 0$ such that for all $c \leq c_3$:

a. $\tilde{V}(p_*) < 0$;
b. there exists $\underline{\pi}_{LR} \in (p_*, \min\{\hat{p}, \xi\})$ such that $V_{L0}(p; p_*) \geq \tilde{V}(p)$ if and only if $p \leq \underline{\pi}_{LR}$.

Proof. For each p^*, there exists $c_3^1 > 0$ such that $p_* < \xi$ and $V_S(\xi) > 0$ for all $c \leq c_3^1$. In the sequel, we assume that $c < c_3 := \min\{c_3^1, c_3^2\}$, where c_3^2 is defined in the proof for part b.

a. Suppose $p^* \leq \eta$ so that $\tilde{V}(p) = V_R(p)$ for all $p \leq p^*$. Since (C5) holds with equality at $q = \pi_{\ell L} = p_*$, we can substitute $\lambda v/c$ in the explicit solution for $V_R(p_*)$ and get

$$\tilde{V}(p_*) = V_R(p_*) < 0 \Leftrightarrow \log\left(\frac{p^*}{1-p^*}\frac{1-p_*}{p_*}\right) > \frac{p^* - p_*}{p^*(1-p_*)}.$$

Define $f_1(p) := \log([p^*/(1-p^*)][(1-p)/p]) - (p^* - p)/[p^*(1-p)]$. The above inequality holds since $f_1(p^*) = 0$ and $f_1'(p) < 0$ for all $p < p^*$. If $p^* > \eta$, then $\tilde{V}(p) = V_{RS}(p; \xi)$ for all $p \leq \xi$. In this case,

$$\tilde{V}(p_*) < 0 \Leftrightarrow \frac{2p_*(p^* - \xi)}{p^*\xi(1-\xi)} + p_* \log\left(\frac{\xi}{1-\xi}\frac{1-p_*}{p_*}\right) + 1 - \frac{p_*}{\xi} > \frac{p^* - p_*}{p^*(1-p_*)}.$$

Define

$$f_2(p) := \frac{2p(p^* - \xi)}{p^*\xi(1-\xi)} + p \log\left(\frac{\xi}{1-\xi}\frac{1-p}{p}\right) + 1 - \frac{p}{\xi} - \frac{p^* - p}{p^*(1-p)}.$$

The desired result ($f_2(p_*) > 0$) holds, because $f_2(0) = 0$, $f_2(\xi) > 0$ and f_2 is concave over $p \in (0, \xi]$.

b. We begin by showing that there exists $c_3^2 > 0$ such that for $c < c_3^2$, $V_{L0}(x; p_*) < \tilde{V}(x)$, where $x \in \{\hat{p}, \xi\}$. Since $\tilde{V}(p) \geq V_S(p)$ (lemma 4d), it suffices to show $V_{L0}(x; p_*) < V_S(x)$. Indeed, we have $V_{L0}(x; p_*) - V_S(x) = \{(x - p_*)/(p^* - p_*) - (x/p^*)\}v + C_S(x) - C_-(x; p_*) < C_S(x) - C_-(x; p_*)$, since $C_S(x)/c$ is independent of c and $C_-(x; p_*)/c \to \infty$ as $c \to 0$.[41]

By lemma 7a, we have $V_{L0}(p_*; p_*) = 0 > \tilde{V}(p_*)$. Since for $c < c_3^2$, $V_{L0}(\min\{\hat{p}, \xi\}; p_*) < \tilde{V}(\min\{\hat{p}, \xi\})$, there exists an intersection of $V_{L0}(p; p_*)$ and $\tilde{V}(p)$ at some $p \in (p_*, \min\{\hat{p}, \xi\})$. In the remainder of the proof, we show that $V_{L0}(\cdot; p_*)$ can cross $\tilde{V}(\cdot)$ only from above, which establishes uniqueness of the intersection on the whole interval (p_*, p^*).

[41] This is because $p_* \to 0$ as $c \to 0$ so that for the L0 strategy, the expected waiting time from any starting point x becomes infinite if the state is L.

We first consider $p^* < \eta$. In this case, $\tilde{V}(p) = V_R(p)$, and lemma 3 implies that $V_R(p) > V_S(p)$. Then, by lemma 2c, $V_{L0}(p; p_*)$ can cross $\tilde{V}(p)$ only from above.

Second, consider $p^* \geq \eta$. Since $\tilde{V}(p) = V_{LS}(p; \xi)$ for $p \in [\xi, \bar{\pi}_{LR}]$ and both V_{L0} and V_{LS} satisfy (C2), no intersection can occur in the interval $[\xi, \bar{\pi}_{LR}]$. Outside this interval, $\tilde{V}(p)$ satisfies (C1) and $\tilde{V}(p) > V_S(p)$ by lemma 4d. Therefore, again lemma 2c implies that $V_{L0}(p; p_*)$ can cross $\tilde{V}(p)$ only from above. QED

C1.2.1. Equilibrium Payoffs. The equilibrium value functions are given as follows:

$$V(p) := \begin{cases} 0 & \text{if } p \in [0, p_*), \\ V_{L0}(p; p_*) & \text{if } p \in [p_*, \underline{\pi}_{LR}), \\ \tilde{V}(p) & \text{if } p \in [\underline{\pi}_{LR}, p^*), \\ v & \text{if } p \geq p^*, \end{cases} \quad \text{and} \quad U(p) := \begin{cases} U_\ell(p) & \text{if } p \in [0, p_*), \\ U_{L0}(p; p_*) & \text{if } p \in [p_*, \underline{\pi}_{LR}), \\ \tilde{U}(p) & \text{if } p \in [\underline{\pi}_{LR}, p^*), \\ U_r(p) & \text{if } p \geq p^*. \end{cases}$$

LEMMA 8. When (Cond2) fails, $V(\cdot)$ is nonnegative and nondecreasing on $[0, p^*]$ and strictly convex on $[p_*, p^*]$.

Proof. Lemma 7b implies that $V(p) = \max\{V_{L0}(p; p_*), \tilde{V}(p)\}$ over $[p_*, p^*]$. This is strictly convex since it is the maximum of two strictly convex functions. Strict convexity of $V_{L0}(\cdot)$ on $[p_*, p^*]$ is routine to verify; we had already shown convexity of $\tilde{V}(p)$ in lemma 4a. Finally, by (C5), $V(p)$ is continuously differentiable at $p_* = \pi_{\ell L}$ and therefore convex on $[0, p^*]$. This also implies that $V(p)$ is nondecreasing. QED

C2. Verifying the Sender's Incentives

We show that for each p^*, the sender's strategy is a best response if the buyer waits if and only if $p \in W$.[42] To this end, we must show that in the waiting region, the sender's equilibrium value function solves the Hamilton-Jacobi-Bellmann (HJB) equation:[43]

$$\max_{\alpha(\cdot;p)} \sum_{q \neq p} \alpha(q; p) v(p; q) = c, \tag{HJB}$$

where $v(p; q)$ is as defined in section V.A. Outside the waiting region, the sender's value is independent of her strategy. Still, our refinement requires that her strategy maximize her instantaneous payoff normalized by dt; that is, her choice of experiment should solve

$$\max_{\alpha(\cdot;p)} \sum_{q \neq p} \alpha(q; p) v(p; q) - \mathbf{1}_{\{\sum \alpha(q;p_t) > 0\}} c. \tag{Ref}$$

Proposition 1b implies that if $V(p)$ meets certain conditions, then we can restrict attention to Poisson experiments with jump targets, 0, p_*, and p^*, which

[42] Recall that $W = [p_*, p^*)$ in proposition 2 and $W = (p_*, p^*)$ in proposition 3.

[43] More formally, since $V(p)$ has kinks, we show that it is a *viscosity solution* of (HJB). Together with $V(p) > 0$, this is necessary and sufficient for optimality of the sender's strategy. For necessity, see theorem 10.8 in Oksendal and Sulem (2009). While we are not aware of a statement of sufficiency that covers precisely our model, the arguments in Soner (1986) can be easily extended to show sufficiency.

greatly simplifies both (HJB) and (Ref). Here, we show that our equilibrium value function $V(\cdot)$ satisfies all properties required by proposition 1b, namely, that it is nonnegative, increasing, and strictly convex on $(p_*, p^*]$, and $V(p_*)/p_* \leq V'(p_*)$. If (Cond2) holds, the first two properties hold by lemma 6, while strict convexity of $V(p)$ follows from lemma 4a and $V(p) = \tilde{V}(p)$ for $p \in [p_*, p^*]$. The last property also holds because $\tilde{V}(p)$ is convex and $\lim_{p \to 0} \tilde{V}(p) = -c/\lambda < 0$. If (Cond2) fails, the first three properties follow from lemma 8 and $V(p_*)/p_* \leq V'(p_*)$ also holds, because $p_* = \pi_{\ell L} > 0$ and $V(\pi_{\ell L}) = V'(\pi_{\ell L}) = 0$.

Stopping region.—We first apply proposition 1b to the stopping region and verify (Ref). For $p \geq p^*$, the result is immediate from proposition 1b(iii). Now consider p below p_*. Proposition 1b(ii) implies that the sender has three choices: two L-drifting experiments with jump target p_* or p^* and simply passing. This reduces (Ref) to

$$\max_{\alpha_*, \alpha^* \geq 0} \lambda p(1-p)\left[\alpha_* \frac{V(p_*)}{p_* - p} + \alpha^* \frac{v}{p^* - p}\right] - c(\alpha_* + \alpha^*) \text{ subject to } \alpha_* + \alpha^* \leq 1.$$

i. Proposition 3: if (Cond2) fails, then $V(p_*) = 0$ so that $\alpha_* = 0$ is optimal. The coefficient of α^* is $\lambda v p(1-p)/(p^* - p) - c$. By (C5), this is negative for all $p < p_* = \pi_{\ell L}$, so $\alpha^* = 0$ is optimal. Therefore, for all $p \in [0, p_*]$, passing—the sender's strategy, as specified in proposition 3—satisfies (Ref).

ii. Proposition 2: if (Cond2) holds, then as discussed in section V.A and depicted in figure 7, there exists a cutoff $\pi_0 < p_*$ such that the coefficient of α_* is greater than the coefficient of α^* if and only if $p > \pi_0$.[44] The following lemma shows that $\pi_{\ell L} < \pi_0$.

LEMMA 9. If (Cond2) holds, then $\pi_{\ell L} < \pi_0$.

Proof. Let $\pi_{\ell R}$ be the value of p such that $\tilde{V}(p) = 0$. We show that $\pi_{\ell L} < \pi_{\ell R} < \pi_0$. The latter inequality is immediate from the strict convexity of $\hat{V}(\cdot)$ on $[0, p^*]$ (lemma 4a) and the definition of π_0. For the former inequality, it suffices to show that $\tilde{V}(\pi_{\ell L}) < 0$, which is shown as in the proof of lemma 7a. QED

As in the case of proposition 3, passing satisfies (Ref) for $p \leq \pi_{\ell L}$. Moreover, we have shown that $\alpha^* = 1$ satisfies (Ref) for $p \in (\pi_{\ell L}, \pi_0)$ and $\alpha_* = 1$ satisfies it for $p \in [\pi_0, p_*)$. Therefore, the sender's strategy in proposition 2 satisfies (Ref) for all $p < p_*$.

Waiting region.—If we apply proposition 1b(i) to $p \in W$, (HJB) simplifies to

$$c = \lambda p(1-p) \max_{\alpha \in [0,1]} \left[\alpha \frac{v - V(p)}{p^* - p} - (1-\alpha)\frac{V(p)}{p} - (2\alpha - 1)V'(p)\right]. \quad \text{(HJB-S)}$$

Our goal is to show that the value function $V(p)$ satisfies this equation at every $p \in W$. The key argument is the following unimprovability lemma:

LEMMA 10 (Unimprovability).

a. If $V_+(p)$ satisfies (C1) and $V_+(p) \geq V_S(p)$ at $p \in [0, p^*)$, then $V_+(p)$ satisfies (HJB-S) at p. If $V_+(p) > V_S(p)$, then $\alpha = 0$ is the unique maximizer in (HJB-S).

[44] Specifically, π_0 satisfies

$$\frac{V(p_*) - V(\pi_0)}{p_* - \pi_0} = \frac{V(p_*) - V(\pi_0)}{p_* - \pi_0} \Leftrightarrow \frac{V(p^*)}{p^* - \pi_0} = \frac{v}{p^* - \pi_0} \Leftrightarrow \pi_0 = \frac{p_* v - p^* V(p_*)}{v - V(p_*)}.$$

b. If $V_-(p)$ satisfies (C2) and $V_-(p) \geq V_S(p)$ at $p \in [0, p^*)$, then $V_-(p)$ satisfies (HJB-S) at p. If $V_-(p) > V_S(p)$, then $\alpha = 1$ is the unique maximizer in (HJB-S).

Proof. a. If we substitute $V'(p) = V'_+(p)$ from (C1), (HJB-S) simplifies to

$$\max_{\alpha \in [0,1]} \left[-\frac{p^*}{(p^* - p)p} (V(p) - V_S(p)) \right] \alpha = 0.$$

If $V(p) - V_S(p) \geq 0$, $\alpha = 0$ is a maximizer, so the above condition holds. Further, if $V(p) > V_S(p)$, then $\alpha = 0$ is the unique maximizer. The proof for part b is similar. QED

By lemmas 4d and 7b, $V(p) \geq V_S(p)$ holds for all $p \in (p_*, p^*)$. Therefore, the unimprovability lemma 10 implies that $V(p)$ satisfies (HJB) for all points where it is differentiable. At the remaining points $\underline{\pi}_{LR}$ and $\bar{\pi}_{LR}$, the value function satisfies (C1) and (C2), respectively, if we replace V'_+ by the right derivative and V'_- by the left derivative. As in the proof of the unimprovability lemma 10, this implies that (HJB) continues to hold if we insert directional derivatives. Using this observation together with the fact $V(p)$ is convex at the points $\underline{\pi}_{LR}$ and $\bar{\pi}_{LR}$ where it has kinks, we can show that $V(p)$ is a viscosity solution of (HJB), which is sufficient for optimality of the sender's strategy in the waiting region (see n. 43).

C3. Verifying the Receiver's Incentives

We now prove the optimality of the receiver's strategy for each belief p, taking as given the sender's strategy. If the sender passes, which occurs when $p \leq \pi_{\ell L}$ or $p \geq p^*$, then the receiver gains nothing from waiting. Since $\pi_{\ell L} \leq p_* < \hat{p}$ (assuming $c \leq \min\{c_1, c_2\}$) and $p^* > \hat{p}$, the receiver chooses ℓ if $p \leq \pi_{\ell L}$ and r if $p \geq p^*$.

Consider next the region $(\pi_{\ell L}, p^*)$ on which the sender does not pass. For this region, we prove that given the sender's strategy, the receiver's strategy solves her optimal stopping problem in the dynamic programming sense. By standard verification theorems, it is sufficient for optimality that the receiver's equilibrium payoff $U(p)$ satisfies the following HJB conditions for all p:[45]

$$c \geq \lambda p(1 - p) \left[\alpha(p) \frac{U(q(p)) - U(p)}{q(p) - p} + (1 - \alpha(p)) \frac{u_\ell^L - U(p)}{p} - (2\alpha(p) - 1)U'(p) \right], \quad \text{(R1)}$$

and

$$U(p) \geq \max\{U_\ell(p), U_r(p)\}, \quad \text{(R2)}$$

[45] The receiver's value function $U(p)$ is not continuously differentiable at p_* (in case [Cond2] holds), $\underline{\pi}_{LR}$, and $\bar{\pi}_{LR}$. At these nonsmooth points, we replace $U'(p)$ in (R1) by the right derivative $U'(p_+)$, which is the directional derivative in the direction of the belief dynamics given by the sender's strategy. With this modification, (R1) is well defined for all p.

By standard verification theorems, the conditions (R1) and (R2) are sufficient for optimality if $U(p)$ is continuously differentiable. To see that sufficiency also holds for the receiver's problem, note that we can verify the receiver's strategy separately for intervals that are closed under the belief dynamics given by the sender's strategy. For example, if (Cond2) holds and $p^* \geq \eta$, we can partition $(\pi_{\ell L}, p^*)$ into $P = \{(\pi_{\ell L}, p_*), [p_*, \bar{\pi}_{LR}), [\bar{\pi}_{LR}, p^*)\}$. If the

and at least one condition holds with equality. Here, $(\alpha(p), q(p))$ represents the sender's strategy as specified in propositions 2 and 3, respectively.[46]

Waiting region.—Suppose $p \in W$. For all points where the receiver's equilibrium payoff function $U(p)$ is differentiable, by construction, it satisfies (R1) with equality.[47] Hence, it suffices to prove (R2). We first show that at p^*, the slope of $U(p)$ is less than or equal to the slope of $U_r(p)$. To this end, observe

$$\begin{aligned} U'(p^*) &= U_R'(p^*) = \frac{U_r(p^*) - u_\ell^L}{p^*} + \frac{c}{\lambda p^*(1-p^*)} \\ &= U_r'(p^*) - \frac{u_\ell^L - u_r^L}{p^*} + \frac{c}{\lambda p^*(1-p^*)}. \end{aligned}$$

Since $u_\ell^L > u_r^L$, we have $U'(p^*) \le U_r'(p^*)$ whenever $c \le c_4 := (1-p^*)(u_\ell^L - u_r^L)$.

i. Proposition 2: when (Cond2) holds, $U(\cdot)$ is convex on $[p_*, p^*]$ since $U(p) = \tilde{U}(p)$ for $p \in [p_*, p^*)$ and $\tilde{U}(\cdot)$ is convex on $[0, p^*]$ (lemma 4a). Together with $U'(p^*) \le U_r'(p^*)$, this implies that $U(p) \ge U_r(p)$ for all $p \in [p_*, p^*]$, provided that $c \le c_4$. We have argued in footnote 40 that $U(p) \ge U_\ell(p)$ for all $p \in [p_*, p^*]$. Therefore, (R2) holds for all $p \in [p_*, p^*)$.

ii. Proposition 3: we begin by showing that $U_{L0}(p, p_*) > U_\ell(p)$ for all $p \in (p_*, p^*)$. Since $U_{L0}(p_*, p_*) = U_\ell(p_*)$, we have

$$\begin{aligned} U_{L0}'(p_*; p_*) &= \frac{U_r(p^*) - U_\ell(p_*)}{p^* - p_*} - \frac{c}{\lambda p_*(1-p_*)} \ge \frac{U_\ell(p^*) - U_\ell(p_*)}{p^* - p_*} + \frac{v}{p^* - p_*} - \frac{c}{\lambda p_*(1-p_*)} \\ &= \frac{U_\ell(p^*) - U_\ell(p_*)}{p^* - p_*} = u_\ell^R - u_\ell^L = U_\ell'(p_*), \end{aligned}$$

where the inequality holds since (Cond2) fails, and the second equality follows from (C5) and $p_* = \pi_{\ell L}$. Together with the fact that $U_{L0}(\cdot; p_*)$ is convex on $[p_*, p^*]$, this implies that $U_{L0}(p, p_*) > U_\ell(p)$ for all $p \in (p_*, p^*)$.

For $p \in (p_*, \underline{\pi}_{LR})$, $U(p) = U_{L0}(p, p_*)$. By lemma 7b, $\underline{\pi}_{LR} \le \hat{p}$, provided that $c \le c_3$. Hence $U_r(p) < U_\ell(p)$ for $p < \underline{\pi}_{LR}$, and (R2) holds since $U(p) = U_{L0}(p, p_*) > U_\ell(p)$ for $p \in (p_*, \underline{\pi}_{LR})$.

Next, suppose $p \in [\underline{\pi}_{LR}, p^*)$. Here $U(p) = \tilde{U}(p)$, and by the same arguments as in part i, we have $\tilde{U}(p) > U_r(p)$. To show that $\tilde{U}(p) > U_\ell(p)$, it suffices to show that $\tilde{U}(p) - U_{L0}(p; p_*) > 0$. Since the sender and the receiver incur the same cost for each strategy, we can rewrite this difference as

$$\tilde{U}(p) - U_{L0}(p; p_*) = \tilde{V}(p) - V_{L0}(p; p_*) + \frac{p_*(p^* - p)}{p^*(p^* - p_*)}(U_r(p^*) - U_\ell(p^*) - v) > 0.$$

---

prior belief is in one of these intervals, the posterior will never leave it unless a Poisson jump occurs, and the continuation value after a jump can be taken as fixed. This means that we can verify the optimality of the receiver's strategy separately for each interval; since $U(p)$ is continuously differentiable on each of the intervals, the standard verification theorems apply.

[46] Specifically, $\alpha(p) = 0$ if the sender plays the R-drifting experiment, $(\alpha(p), q(p)) = (1, q)$ if she plays the L-drifting experiment with jump target q, and $(\alpha(p), q(p)) = (1/2, p^*)$ if she plays the stationary strategy.

[47] At kinks, $U(p)$ satisfies (R1) if $U'(p)$ is replaced by $U'(p_+)$ (see n. 45).

The inequality holds since by lemma 7b, $\tilde{V}(p) - V_{L0}(p; p_*) \geq 0$ for $p \geq \underline{\pi}_{LR}$, and $U_r(p^*) - U_\ell(p^*) - v > 0$ if (Cond2) is violated.

The stopping region with $p \in (\pi_{\ell L}, p_*)$.—If (Cond2) fails, then $p_* = \pi_{\ell L}$, so this case does not arise. The proof of proposition 3 is thus complete.

Now suppose that (Cond2) holds and $p \in (\pi_{\ell L}, p_*)$. In this case, $U(p)$ satisfies (R2) with equality, so it suffices to show (R1). Consider first $p \in [\pi_0, p_*)$. For these beliefs, the sender adopts the L-drifting experiment with jump target p_*, that is, $(\alpha(p), q(p)) = (1, p_*)$. When we plug this into (R1) and use the fact that $U(p) = U_\ell(p)$ for all $p \leq p_*$, the right-hand side of (R1) is equal to zero so that (R1) is satisfied.

Finally, consider $p \in [\pi_{\ell L}, \pi_0)$, at which the sender plays the L-drifting experiment with jump target p^*, so $(\alpha(p), q(p)) = (1, p^*)$. Since $U(p) = U_\ell(p)$ for all $p \leq p_*$, (R1) reduces to

$$\lambda p(1-p)\left[\frac{U_r(p^*) - U_\ell(p)}{p^* - p} - U_\ell'(p)\right] = \frac{\lambda p(1-p)}{p^* - p}(U_r(p^*) - U_\ell(p^*)) \leq c,$$

which is equivalent to $p \leq \phi_{\ell L}$, where $\phi_{\ell L}$ is the unique value of p such that

$$\frac{\lambda p(1-p)}{p^* - p}(U_r(p^*) - U_\ell(p^*)) = c.$$

The following lemma shows that $\phi_{\ell L} \geq \pi_0$ if $c \leq c_5$ for some $c_5 > 0$. It then follows that if $c \leq \min\{c_1, \ldots, c_5\}$, the receiver has no incentive to deviate from his prescribed strategy in proposition 2, completing the proof.

LEMMA 11. Suppose that (Cond2) holds. There exists $c_5 > 0$ such that if $c \leq c_5$, then $\phi_{\ell L} \geq \pi_0$.

Proof. Let $\Delta U := U_r(p^*) - U_\ell(p^*)$. Since $\phi_{\ell L}$ is the lowest p such that $p(1-p)\lambda\Delta U/(p^* - p) \geq c$,

$$\pi_0 \leq \phi_{\ell L} \quad \Leftrightarrow \quad \frac{\pi_0(1-\pi_0)}{p^* - \pi_0}\frac{\lambda}{c}\Delta U < 1. \tag{C6}$$

It suffices to show that this inequality holds in the limit as $c \to 0$. Recall that

$$\frac{V(p^*)}{p^* - \pi_0} = \frac{v}{p^* - \pi_0} \quad \Leftrightarrow \quad \pi_0 = \frac{p^*}{v - V(p_*)}\left(\frac{p_*}{p^*}v - V(p_*)\right) = \frac{p^*}{v - V(p_*)}\tilde{C}(p_*),$$

where $\tilde{C}(p_*) = (p_*/p^*)v - \tilde{V}(p_*)$ denotes the total persuasion costs incurred when $p = p_* = \phi_{\ell R}$ and (Cond2) holds. By the definition of $\tilde{V}(p_*)$, $\tilde{C}(p_*)$ can be written as

$$\tilde{C}(p_*) = C_+(p_*; q_R) + \frac{p_*}{q_R}C_S(q_R) = \left(p_* \log\left(\frac{q_R}{1-q_R}\frac{1-p_*}{p_*}\right) + 1 - \frac{p_*}{q_R} + \frac{p_*}{q_R}\frac{2(p^* - q_R)}{p^*(1-q_R)}\right)\frac{c}{\lambda},$$

where $q_R := p^*$ if $p^* \leq \eta$ and $q_R := \xi$ if $p^* > \eta$. Importantly, as $c \to 0$, we have $p_* \to 0$, $\tilde{C}(p_*) \to 0$, and $\tilde{C}(p_*)\lambda/c \to 1$. It follows that $\pi_0 \to 0$ and $\pi_0\lambda/c \to p^*/v$, so

$$\frac{\pi_0(1-\pi_0)}{p^* - \pi_0}\frac{\lambda}{c}\Delta U \to \frac{\Delta U}{v} < 1,$$

where the inequality is due to (Cond2). This completes the proof. QED

C4. SMPE Uniqueness Given p^*

Fix any p^*. To show that for c sufficiently small, the strategy profiles in propositions 2 and 3 are the unique SMPEs, we prove that any other choice of p_* than specified in sections C1.1 and C1.2 (i.e., $p_* \neq \phi_{\ell R}$ if [Cond2] holds and $p_* \neq \pi_{\ell L}$ if [Cond2] fails) cannot yield an SMPE. This requires a full characterization of the sender's optimal dynamic strategy, given any lower bound p_* and upper bound p^*, and a thorough examination of the receiver's incentives in the stopping region as well as in the waiting region. The former closely follows our construction and analysis of the equilibrium value functions in sections C1 and C2, and the latter follows closely section C3. We relegate the full proof to the supplemental material.

References

Arrow, K. J., D. Blackwell, and M. A. Girshick. 1949. "Bayes and Minimax Solutions of Sequential Decision Problems." *Econometrica* 17 (3–4): 213–44.

Au, P. H. 2015. "Dynamic Information Disclosure." *RAND J. Econ.* 46:791–823.

Bizzotto, J., J. Rüdiger, and A. Vigier. 2021. "Dynamic Persuasion with Outside Information." *American Econ. J. Microeconomics* 13 (1): 179–94.

Brocas, I., and J. D. Carrillo. 2007. "Influence through Ignorance." *RAND J. Econ.* 38 (4): 931–47.

Che, Y.-K., and J. Hörner. 2018. "Recommender Systems as Mechanisms for Social Learning." *Q.J.E.* 133:871–925.

Che, Y.-K., K. Kim, and K. Mierendorff. 2021. "Keeping the Listener Engaged: A Dynamic Model of Bayesian Persuasion." Working paper.

Che, Y.-K., and K. Mierendorff. 2019. "Optimal Dynamic Allocation of Attention." *A.E.R.* 109:2993–3029.

Che, Y.-K., and J. Sákovics. 2004. "A Dynamic Theory of Holdup." *Econometrica* 72 (4): 1063–103.

Denti, T., M. Marinacci, and A. Rustichini. 2022. "Experimental Cost of Information." *A.E.R.* 112 (9): 3106–23.

Ely, J. C. 2017. "Beeps." *A.E.R.* 107:31–53.

Ely, J. C., and M. Szydlowski. 2020. "Moving the Goalposts." *J.P.E.* 128 (2): 468–506.

Escudé, M., and L. Sinander. 2023. "Slow Persuasion." *Theoretical Econ.* 18 (1): 129–62.

Fudenberg, D., P. Strack, and T. Strzalecki. 2018. "Speed, Accuracy, and the Optimal Timing of Choices." *A.E.R.* 108:3651–84.

Goodman, G. 1970. "Subfunctions and the Initial Value Problem for Differential Equations Satisfying Caratheodory's Hypotheses." *J. Differential Equations* 41:232–42.

Henry, E., and M. Ottaviani. 2019. "Research and the Approval Process: The Organization of Persuasion." *A.E.R.* 109 (3): 911–55.

Kamenica, E., and M. Gentzkow. 2011. "Bayesian Persuasion." *A.E.R.* 101 (6): 2590–615.

Ke, T. T., and J. M. Villas-Boas. 2019. "Optimal Learning before Choice." *J. Econ. Theory* 180:383–437.

Keller, G., and S. Rady. 2015. "Breakdowns." *Theoretical Econ.* 10 (1): 175–202.

Keller, G., S. Rady, and M. Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica* 73 (1): 39–68.

Klein, N., and S. Rady. 2011. "Negatively Correlated Bandits." *Rev. Econ. Studies* 78 (2): 693–732.
Kremer, I., Y. Mansour, and M. Perry. 2014. "Implementing the 'Wisdom of the Crowd.'" *J.P.E.* 122:988–1012.
Marinovic, I., and M. Szydlowski. 2020. "Monitor Reputation and Transparency." Working paper, Stanford Univ. Graduate School Bus.
Mayskaya, T. 2020. "Dynamic Choice of Information Sources." ICEF Working Paper no. WP9/2019/05, California Inst. Tech.
McClellan, A. 2022. "Experimentation and Approval Mechanisms." *Econometrica* 90 (5): 2215–47.
Moscarini, G., and L. Smith. 2001. "The Optimal Level of Experimentation." *Econometrica* 69 (6): 1629–44.
Nikandrova, A., and R. Pancs. 2018. "Dynamic Project Selection." *Theoretical Econ.* 13 (1): 115–44.
Oksendal, B., and A. Sulem. 2009. *Applied Stochastic Control of Jump Diffusions*, 3rd ed. Berlin: Springer.
Orlov, D., A. Skrzypacz, and P. Zryumov. 2020. "Persuading the Principal to Wait." *J.P.E.* 128 (7): 2542–78.
Renault, J., E. Solan, and N. Vieille. 2017. "Optimal Dynamic Information Provision." *Games and Econ. Behavior* 104 (4): 329–49.
Simon, L. K., and M. B. Stinchcombe. 1989. "Extensive Form Games in Continuous Time: Pure Strategies." *Econometrica* 57 (5): 1171–214.
Soner, H. M. 1986. "Optimal Control with State Space Constraint. II." *SIAM J. Control and Optimization* 24 (6): 1110–22.
Wald, A. 1947. "Foundations of a General Theory of Sequential Decision Functions." *Econometrica* 15 (4): 279–313.
Zhong, W. 2022. "Optimal Dynamic Information Acquisition." *Econometrica* 90 (4): 1537–82.