

Keeping the Listener Engaged: a Dynamic Model of Bayesian Persuasion*

Yeon-Koo Che Kyungmin Kim Konrad Mierendorff[†]

August 31, 2021

Abstract

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 (MPE) of this game. However, for persuasion costs sufficiently small, a version of a folk theorem holds: outcomes that approximate [Kamenica and Gentzkow \(2011\)](#)'s sender-optimal persuasion as well as full revelation and everything in between are obtained in MPE, as the cost vanishes.

Keywords: Bayesian persuasion, general Poisson experiments, Markov perfect equilibria, folk theorem.

JEL Classification Numbers: C72, C73, D83

1 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 receiver ultimately accepts that pitch—or is “persuaded”—depends on the underlying truth (the state of the world) but importantly,

*We thank Martin Cripps, Jeff Ely, Faruk Gul, Emir Kamenica, Stephan Laueremann, 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. Yeon-Koo Che is supported by National Science Foundation (SES-1851821); he and Kyungmin Kim are supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2020S1A5A2A03043516).

[†]Che: Department of Economics, Columbia University (email: yeonkooche@gmail.com); Kim: Department of Economics, Emory University (email: kyungmin.kim@emory.edu); Mierendorff: Department of Economics, University College London. Our dear friend and coauthor, Konrad Mierendorff, passed away in August 2021. All the ideas and results in this paper are collaborative work by three authors, but Che and Kim are responsible for any remaining errors.

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, with the friction that real information takes time to generate, 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 at any time 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 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), and 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 bad-news Poisson experiment of given accuracy (see, 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 Footnote 9). 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.¹ 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 (MPE)* 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 version of a folk theorem. 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 (2011)’s sender-optimal persuasion, to ones that approximate full revelation, and covers

¹A salesperson’s pitches, which include her manner, tones, and body languages, not just her messages, can reveal a lot about what she is “intending” to say, not just what she is saying. At the same time, whether the pitches succeed depends on the buyer’s specific needs, and is uncertain from the salesperson’s perspective. It is also reasonable that an experienced salesperson could tell the outcome of her pitches from the buyer’s reactions.

everything in between.

In the “persuasion failure” 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 outcomes that approximate KG’s optimal persuasion, full revelation, and anything in between, as the flow cost c tends to 0.²

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

In our MPEs, the sender optimally makes use of the following three strategies: (i) *confidence-building*, (ii) *confidence-spending*, and (iii) *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. This strategy triggers upward movement of the belief when the state is R but quite likely even when it is L ; in fact, it minimizes the probability of bad news, by insisting that the news be conclusive.³ 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,

²The 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 hold-up problem. In fact, the main tension in our dynamic persuasion problem can be interpreted as a hold-up 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 analogue in [Che and Sákovics \(2004\)](#).

³In the salesperson context, this may correspond to explaining a possible harm or a side-effect of a product that does not apply to most buyers, and to rule it out for a buyer, a likely event, can improve her odds of eventual persuasion.

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.⁴ 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 also accomplish persuasion fast and thus can be used for a range of beliefs, but the sender would be running 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 5 describes precisely how the sender uses them in different equilibria.

Related literature. This paper relates to several strands of literature. First, it 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](#); [Bizzotto, Rüdiger, and Vigier, 2020](#); [Orlov, Skrzypacz, and Zryumov, 2020](#)). Our focus is different from most of these papers since we consider gradual production of information and assume that there is no commitment.

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 adopts a drift-diffusion model in a continuous-time setting. 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 eco-

⁴In the salesperson context, this may correspond to touting a virtue of a product irrelevant for most buyers, and for it not to pan out, a likely event, deteriorates the odds of eventual persuasion.

conomic problems/contexts; for example, [Henry and Ottaviani \(2019\)](#) focus on regulatory approval, while we study persuasive *communication*. However, it leads to very different persuasion outcomes: in their models, complete persuasion failure never occurs, and there exists a unique equilibrium.⁵ 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 decides simply whether to continue or not, 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 Pans \(2018\)](#), [Che and Mierendorff \(2019\)](#) and [Mayskaya \(2019\)](#) study an agent’s stopping problem when she acquires information through Poisson experiments.⁶ [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 non-conclusive experiment).

Finally, the current paper is closely related to repeated/dynamic communication models. [Margaria and Smolin \(2018\)](#), [Best and Quigley \(2020\)](#), and [Mathevet, Pearce, and Stachetti \(2019\)](#) study repeated cheap-talk communication, and some of them establish versions of folk theorems. Their models consider repeated actions by receiver(s), serially independent states, and feedback on the veracity of the sender’s communication, based on which non-Markovian punishment can be inflicted to support a cooperative outcome. By contrast, the current model considers a fixed state, once-and-for-all action by the receiver (and hence no feedback), and Markov perfect equilibria.

The paper is organized as follows. Section 2 introduces the model. Section 3 illustrates the main ideas of our equilibria. Section 4 states our folk theorem. Section 5 explores the dynamics of our MPE strategies and their implications. Section 6 concludes.

⁵[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,” that is, 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 folk theorem as in our paper (see Footnote 24 below). Finally, their evaluator-authority case is obtained when the receiver can commit to an acceptance threshold.

⁶The Wald stopping problem has also been studied with drift-diffusion learning (e.g., [Moscarini and Smith, 2001](#); [Ke and Villas-Boas, 2019](#); [Fudenberg, Strack, and Strzalecki, 2018](#)), and in a model that allows for general endogenous experimentation (see [Zhong, 2019](#)).

2 Model

We consider a game in which a *Sender* (“she”) wishes to persuade a *Receiver* (“he”). There is an unknown state ω which 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 the receiver		
states/actions	ℓ	r
L	$(0, u_\ell^L)$	(v, u_r^L)
R	$(0, u_\ell^R)$	(v, u_r^R)

The receiver gets u_a^ω if he takes action $a \in \{\ell, r\}$ when the state is $\omega \in \{L, R\}$. The sender’s payoff depends only on the receiver’s action: she gets v if the receiver takes r and zero otherwise. We assume $u_\ell^L > \max\{u_r^L, 0\}$ and $u_r^R > \max\{u_\ell^R, 0\}$, 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 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 $a \in \{\ell, r\}$, let $U_a(p)$ denote the receiver’s expected 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, $U_\ell(\hat{p}) = U_r(\hat{p})$.⁷

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 blue line in the left panel of Figure 1. The flip side is that the receiver enjoys no rents from persuasion; his payoff is $\mathcal{U}(p) := \max\{U_\ell(p), U_r(p)\}$, 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 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 from a feasible set at a flow cost. The set of feasible experiments and cost structure are made precise below. Just as it is costly for the sender to produce information, it is also costly for the receiver to process it. Specifically,

⁷Specifically, for each $p \in [0, 1]$, $U_\ell(p) := pu_\ell^R + (1-p)u_\ell^L$ and $U_r(p) := pu_r^R + (1-p)u_r^L$. Therefore, $\hat{p} = (u_\ell^L - u_r^L) / (u_r^R - u_\ell^R + u_\ell^L - u_r^L)$, which is well-defined in $(0, 1)$ under our assumptions on the receiver’s payoffs.




Figure 1: Payoffs from static persuasion. Solid red curves: payoffs without persuasion (information). Dashed blue curve: the sender’s expected payoff in the KG solution. Blue dots: payoffs in the KG solution at prior p_0 . Dash-dotted brown curves: payoffs under a fully revealing experiment.

if the sender experiments, then the receiver also pays the same flow cost and observes the experiment and its outcome.⁸ 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 KG, wherein the receiver is allowed to take an action only after the sender finishes her information provision. Second, the players’ flow (information) costs are assumed to be the same. This is, however, just a normalization which allows us to directly compare the players’ payoffs, and all subsequent results can be reinterpreted as relative to each player’s individual information cost.⁹

Feasible experiments and flow costs. We endow the sender with a class of Poisson experiments. While the main thrust of our results holds more generally, we adopt this class since it allows us to capture many realistic and rich information processes. Specifically, at each point in time, the sender has a unit capacity to allocate across different experiments that generate Poisson signals. The experiments are indexed by $i \in \mathbb{N}$.¹⁰ Each Poisson experiment $i \in \mathbb{N}$ generates breakthrough news at rate λ^ω in each state $\omega = L, R$, where,

⁸In this sense, the flow cost is interpreted as a “listening cost” rather than a waiting cost. This distinction does not matter in the continuous-time game; our analysis below will not change even if the receiver incurs a flow waiting cost, regardless of the sender’s choice of experiment. However, it would be relevant in a discrete-time version of our model. See the third paragraph of Section 6.

⁹Suppose 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 0, 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 Footnote 33 for a relevant discussion.

¹⁰One can extend this to an uncountable set of experiments. However, in our model, the sender never mixes over an infinite number of experiments, and thus such extra generality is unnecessary.




Figure 2: Arrival rates of feasible Poisson experiments.

for some constant $\lambda > 0$,

$$(\lambda^L, \lambda^R) \in \Lambda := \{(\nu^L + \mu, \nu^R + \mu) : 0 \leq \nu^L + \nu^R \leq \lambda, \mu \geq 0\}.$$

In other words, the sender can split λ between the two states and also boost the arrival rates for both states equally by μ .¹¹ The component (ν^L, ν^R) produces real information about the state if $\nu^L \neq \nu^R$, while (μ, μ) represents “noise.” See the shaded area of Figure 2 for a visual representation of the set Λ .

For instance, suppose the sender chooses $(\mu, \lambda + \mu)$ for some $\mu \geq 0$. Breakthrough news would then favor state R , but by how much depends on μ . If $\mu = 0$, the news is conclusive R -evidence. As μ rises, the news becomes less conclusive about R but arrives with a higher rate, as depicted by different points on the dashed line in Figure 2. One can interpret μ as a “dilution” of R -evidence. With $\mu > 0$, the sender “overclaims” R -evidence even in state L . The more she does this, the more diluted the information content of the breakthrough news becomes, as depicted by the flattening of the associated ray from the origin in Figure 2 as μ increases from 0 to μ_1 to μ_2 . Indeed, the slope of the ray from the origin passing through (λ^L, λ^R) corresponds to the *likelihood ratio* λ^R/λ^L associated

¹¹One can interpret λ as the measure of maximal information that can be generated in a unit time. To see this, consider an experiment with binary signals that occur with probabilities:

Binary-signal experiment i		
states/signals	L -signal	R -signal
L	x	$1 - x$
R	$1 - y$	y

and let $E(x, y)$ denote such a binary-signal experiment, with \mathcal{E} denoting the set of all such experiments. Assume $1 \leq x + y \leq 1 + \lambda dt$ for a small interval of time $dt > 0$. Any feasible Poisson experiment we consider can be obtained in the limit of such binary experiments as $dt \rightarrow 0$. Observe that the quantity $x + y - 1$ “measures” the informativeness of the experiment. Formally, this *measure* of information induces an order on \mathcal{E} that “completes” the Blackwell order: *if $E(x', y')$ Blackwell dominates $E(x, y)$, then $x' + y' \geq x + y$* . In this sense, our model caps the “amount” of information the sender may generate per unit of time (from an experiment in \mathcal{E}) by λ .

with the news. Meanwhile, the *prevalence* of the news increases with the distance from the origin. Since for any likelihood ratio chosen, the sender will wish to maximize the prevalence of that signal, it is routine to see that the sender will choose either $\nu^L = \lambda$ or $\nu^R = \lambda$ so that (λ^L, λ^R) is on the boundary of the shaded area in Figure 2.

The set Λ is quite rich. Given any prior p , the sender can generate any posterior $q_i \in [0, 1]$ by choosing arrival rates in Λ appropriately. Suppose $q_i > p$. (The opposite case is analogous.) Since the ratio of arrival rates equals the likelihood ratio, Bayes rule implies that $\frac{q_i}{1-q_i} = \frac{\lambda + \mu_i}{\mu_i} \frac{p}{1-p}$. Solving for μ_i , one can obtain the arrival rates of posterior q_i in states L and R , respectively:¹²

$$\lambda^L(q_i) := \lambda \frac{p(1-q_i)}{|q_i-p|} \text{ and } \lambda^R(q_i) := \lambda \frac{q_i(1-p)}{|q_i-p|}.$$

Indeed, for our purpose it is more convenient to view the sender as choosing posterior q_i directly, and represent an experiment i by the posterior q_i triggered by the corresponding news. If the sender chooses $\alpha_i \in [0, 1]$ units of experiment i with posterior q_i , then the news arrives at rate $\alpha_i \lambda^\omega(q_i)$ in state $\omega = L, R$, or unconditionally at rate:

$$(1-p)\alpha_i \lambda^L(q_i) + p\alpha_i \lambda^R(q_i) = \alpha_i \lambda \frac{p(1-p)}{|q_i-p|}. \quad (1)$$

We assume $\sum_{i=1}^{\infty} \alpha_i \leq 1$; namely, total quantity of experiments cannot exceed a unit. We call a collection of experiments $(\alpha_i, q_i)_{i \in \mathbb{N}}$ an *information structure*.¹³

We assume that each Poisson experiment (α_i, q_i) requires a flow cost $c > 0$ per unit. Specifically, if the sender chooses an information structure $(\alpha_i, q_i)_{i \in \mathbb{N}}$, then both players incur total flow cost $(\sum_{i: q_i \neq p} \alpha_i) c$. In equilibrium, either the capacity constraint binds with $\sum_i \alpha_i = 1$, in which case the associated flow cost is c , or no informative experiment is chosen with $\alpha_i = 0$ for all $i \in \mathbb{N}$, in which case the associated flow cost is 0. We refer to the latter choice as the sender's "passing."

Our model generalizes Poisson models considered in the existing literature. To see this, suppose that the sender allocates the entire unit capacity to one Poisson experiment with jump target q . The jump to q then occurs at the rate of $\lambda p(1-p)/|q-p|$. Conclusive R -evidence ($q = 1$) is obtained at the rate of λp , as is assumed in "good"

¹²If $q_i > p$, these arrival rates are obtained by substituting μ_i required for q_i into the arrival rates $(\mu_i, \lambda + \mu_i)$ of the news. The case of $q_i < p$ is analogous. If $q_i = p$, then the news is uninformative; since, as will be seen later, this is never an optimal choice, we leave this case unspecified.

¹³Even though we denote information in terms of the current belief p and jump-target beliefs q_i , recall that the feasible set of experiments is independent of beliefs. This feature distinguishes our approach from the rational inattention model and its generalizations (Sims, 2003; Matejka and McKay, 2015; Caplin, Dean, and Leahy, 2013), in which costs or constraints depend on the current belief, not just on the experiment. More generally, they consider cost functions that are *posterior separable* and based on a measure of information that quantifies the uncertainty in the posterior beliefs induced by an information structure (see also Frankel and Kamenica, 2019). One can show that the information measure discussed in Footnote 11 cannot be obtained from a posterior separable cost function.

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 non-conclusive news with $q \in (0, 1)$, as well as any arbitrary mixture among such experiments. Further, our arrival rate assumption captures the intuitive idea that more accurate information takes longer to generate. For example, assuming $q > p$, the arrival rate increases as the news becomes less precise (q falls), and it approaches infinity as the news becomes totally uninformative (i.e., in the limit as q tends to p). Lastly, limited arrival rates, together with the capacity constraint $\sum_i \alpha_i \leq 1$, capture an important feature of our model that any meaningful persuasion takes time and requires delay.

If no Poisson jump arrives when the sender uses the information structure $(\alpha_i, q_i)_{i \in \mathbb{N}}$, the belief drifts according to the following law of motion:¹⁴

$$\dot{p} = - \left(\sum_{i:q_i > p} \alpha_i - \sum_{i:q_i < p} \alpha_i \right) \lambda p(1 - p). \quad (2)$$

Note that the drift rate depends only on the difference between the fractions of the capacity allocated to “right” versus “left” Poisson signals. In particular, the rate does not depend on the precisions q_i 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.¹⁵ 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 Figure 3, will prove particularly relevant for our purposes. They formalize the three modes of persuasion discussed in the introduction:

- **R -drifting experiment** (confidence building): $\alpha_1 = 1$ with $q_1 = 0$. The sender devotes all her capacity to a Poisson experiment with the (posterior) jump target $q_1 = 0$. In the absence of a jump, the posterior drifts to the right, at rate $\dot{p} = \lambda p(1 - p)$.
- **L -drifting experiment** (confidence spending): $\alpha_1 = 1$ with $q_1 = q$ for some $q > p$. The sender devotes all her capacity to a Poisson experiment with jumps targeting

¹⁴Since the belief is a martingale, we have

$$\begin{aligned} \sum_{i=1}^{\infty} q_i \alpha_i \lambda \frac{p_t(1 - p_t)}{|q_i - p_t|} dt + \left(1 - \sum_{i=1}^{\infty} \alpha_i \lambda \frac{p_t(1 - p_t)}{|q_i - p_t|} dt \right) (p_t + \dot{p}_t dt) &= p_t \\ \iff \left(1 - \sum_{i=1}^{\infty} \alpha_i \lambda \frac{p_t(1 - p_t)}{|q_i - p_t|} dt \right) \dot{p}_t &= \left(\sum_{i=1}^{\infty} \frac{p_t - q_i}{|q_i - p_t|} \alpha_i \right) \lambda p_t(1 - p_t). \end{aligned}$$

Letting $dt \rightarrow 0$, we obtain the updating formula.

¹⁵Suppose $q_i > p$. This means that the sender has chosen $\nu^R = \lambda$ 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.




Figure 3: Three prominent feasible experiments.

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): $\alpha_1 = \alpha_2 = 1/2$ with $q_1 = 0$ and $q_2 = q$ for some $q > p$. The sender assigns an equal share of her capacity to an experiment targeting $q_1 = 0$ and one targeting $q_2 = q$. Absent jumps, the posterior remains unchanged.

Solution concept. We study (pure-strategy) Markov Perfect equilibria (MPE, hereafter) of this dynamic game in which both players’ strategies depend only on the current belief p .¹⁶ Formally, a profile of Markov strategies specifies for each belief $p \in [0, 1]$, an information structure $(\alpha_i, q_i)_{i \in \mathbb{N}}$ chosen by the sender, and an action $a \in \{\ell, r, \text{wait}\}$ chosen by the receiver.¹⁷ An MPE is a strategy profile that, starting from any belief $p \in [0, 1]$, forms a subgame perfect equilibrium.¹⁸ Naturally, this solution concept limits the use of (punishment) strategies depending solely on the payoff-irrelevant part of the histories, and serves to discipline strategies off the equilibrium path.

We impose a restriction that captures the spirit of “perfection” in our continuous time framework. Suppose that for an interval around some low belief p , the receiver would choose action ℓ immediately. In continuous time, this implies that the sender’s strategy at p is inconsequential for the players’ payoffs—with probability one, the game ends with the receiver taking action ℓ . Nevertheless, we require the sender to choose a strategy that maximizes her *flow payoff* in such a situation. This can be seen as selecting an MPE that is robust to a discrete-time approximation. In discrete time, when the belief is in

¹⁶For non-Markov equilibria, see our discussion in Section 6.

¹⁷Note that defining the strategies as functions of p fully determines behavior on and off the equilibrium path; our Markovian restriction means that the same strategy is followed for each belief p regardless of the history.

¹⁸There are well known technical issues in defining a game in continuous time (see Simon and Stinchcombe, 1989; Bergin and MacLeod, 1993). In Online Appendix B, we formally define admissible strategy profiles that guarantee a well defined outcome of the game and also define Markov perfect equilibria.

the region where the receiver stops, the sender has one period to generate information that may change the receiver’s action or induce him to wait. Thus the sender’s strategy has non-trivial payoff consequences. In the same spirit, our refinement requires that in continuous time the sender choose the strategy that maximizes her flow payoff if she has an infinitesimal amount of time to generate information. See Online Appendix B for a formal definition. In what follows, we simply use *MPE* or *equilibrium* to refer to an MPE with this refinement.

3 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 $p_0 < \hat{p}$, such 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 0 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.




Figure 4: Replicating the KG outcome through R -drifting experiments.

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 even 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 payoff from the static KG experiment is $U_\ell(p_0)$. Due to 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.¹⁹ Indeed, this problem leads to the existence of a

¹⁹The KG outcome can also be replicated by other dynamic strategies. For instance, the sender could repeatedly choose a stationary strategy with jumps targeting $q_1 = \hat{p}$ and $q_2 = 0$ until either jump occurs.

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 (if the receiver chooses ℓ).²⁰*

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, discussed at the end of Section 2, requires that the sender choose an information structure that maximizes her flow payoff, which is given by²¹

$$\max_{(\alpha_i, q_i)_{i \in \mathbb{N}}} \sum_{q_i \neq p} \alpha_i \left(\lambda \frac{p(1-p)}{|q_i - p|} \mathbf{1}_{\{q_i \geq \hat{p}\}} v - c \right) \text{ subject to } \sum_{\alpha_i} \alpha_i \leq 1.$$

If the sender chooses any nontrivial experiment, its jump target must be $q_i = \hat{p}$. Hence the optimal information structure is either $(\alpha_1 = 1, q_1 = \hat{p})$ or $\alpha_i = 0$ for all i . The former is optimal if and only if $\frac{\lambda p(1-p)}{\hat{p}-p} v \geq c$, or equivalently $p \geq \hat{\pi}_{\ell L}$. \square

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

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.

²⁰Specifically, $\hat{\pi}_{\ell L}$ equates the sender's flow cost c to the flow 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 $\hat{\pi}_{\ell L}$ (as given by (1)). Solving the equation yields

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

²¹The 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_i , a Poisson jump occurs at rate $\lambda p(1-p)/|q_i - p|$.

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 belief reaches $p^* > \hat{p}$. If the receiver waits until the belief either jumps to 0 or reaches p^* , then her expected payoff is equal to²²

$$U_R(p) = \frac{p^* - p}{p^*} u_\ell^L + \frac{p}{p^*} U_r(p^*) - \left(p \log \left(\frac{p^*}{1 - p^*} \frac{1 - p}{p} \right) + 1 - \frac{p}{p^*} \right) \frac{c}{\lambda}.$$

Importantly, if p^* is sufficiently large (and c is sufficiently small), then $U_R(p)$ (the dashed curve in Figure 5) stays above $\max\{U_\ell(p), U_r(p)\}$ (the black kinked curve) 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 $c \rightarrow 0$. In this fashion, a dynamic persuasion strategy can be constructed to virtually implement the KG solution when c is sufficiently small.

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} , due to

²²To understand this explicit solution, first notice that under the prescribed strategy profile, the receiver takes action ℓ when p jumps to 0, which occurs with probability $(p^* - p)/p^*$, and action r when p reaches to 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

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




Figure 5: Persuasive R -drifting experiments

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:²³ 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.²⁴

4 Folk Theorem

The equilibrium logic outlined in the previous section 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 version of a folk theorem: *any payoff for the sender between her payoff under the KG solution and her payoff from full revelation can be virtually supported as an MPE payoff.*

²³We will show in Section 5.2 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.

²⁴It is clear that this logic extends beyond the Poisson model we employ here. Consider Henry and Ottaviani (2019)’s 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 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 Footnote 5, this model is similar to Henry and Ottaviani (2019)’s no-commitment regime, except that our sender has the option to “pass” without ending the game and the receiver incurs a flow cost. A simple 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}, ∞) as $c \rightarrow 0$.




Figure 6: Implementable payoff set for each player at each p_0 .

Theorem 2 (Folk theorem). *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 (with our refinement) 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 5.2. 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 0. Theorem 2 states that any payoffs in the green and red 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 $\mathcal{U}(p_0)$. In both panels, the thick blue lines correspond to the players' payoffs in the no-persuasion equilibrium of Theorem 1.

Note that Theorem 2 is silent about payoffs in the gray shaded 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).

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 $p_0 v - 2c/\lambda$.²⁵ If the receiver's strategy is to stop at p_0 , then

²⁵In order to understand this payoff, notice that the strategy fully reveals the state, and thus the


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 $p_0v - 2c/\lambda$.

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 exceed at least $p_0v - 2c/\lambda$. 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 $p_0v - 2c/\lambda$. \square

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

Definition 1. A Markov perfect equilibrium is a *simple MPE* (henceforth, 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 $p \geq p^*$.²⁶

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 non-simple MPEs exist or not is irrelevant for our results. While we construct SMPEs to establish our folk theorem, Theorem 3 is valid for *all* MPEs. Finally, we continue to require our refinement with SMPEs.

To prove the folk theorem, 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 $c \rightarrow 0$, p_* approaches 0 as well (see Propositions 2 and 3 in Section 5.2). 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 on the line segment that connects $(0, u_\ell^L)$ and $(p^*, U_r(p^*))$ 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 0. We state this in the following corollary.

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

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

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.

²⁶We do not restrict the receiver's decision at the lower bound p_* , so that the waiting region can be either $W = (p_*, p^*)$ or $W = [p_*, p^*)$. Requiring $W = (p_*, p^*)$ can lead to non-existence of an SMPE (see Proposition 2 below). Requiring $W = [p_*, p^*)$ can lead to non-admissibility of the sender's best response in Proposition 3 (see the discussion of admissibility in Online Appendix B).

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

5 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 the Introduction and Section 2. Then, using them as building blocks, we construct full SMPE strategy profiles and also discuss several behavioral implications of the equilibrium persuasion dynamics.

5.1 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 $q_i \neq p$ when the current belief is p . Then, the belief jumps to q_i at rate $\lambda p(1-p)/|q_i - p|$ (see (1)). Absent jumps, it moves continuously according to (2). Therefore, her flow benefit is given by

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

where $\text{sgn}(x) = x/|x|$, and $V(\cdot)$ is the sender's value of playing the candidate equilibrium strategy.²⁷ Specifically, for $q_i > p$, the flow benefit consists of the value increase from a breakthrough which arises at rate $\lambda \frac{p(1-p)}{|q_i - p|}$ (the first term) and the decay of value in its absence (the second term). For $q_i < p$, the first term captures the value decrease from a breakdown, while the second term represents the gradual appreciation in its absence.

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

$$v(p) := \max_{(\alpha_i, q_i)_i} \sum_{q_i \neq p} \alpha_i v(p; q_i) \text{ subject to } \sum_{i \in \mathbb{N}} \alpha_i \leq 1. \quad (3)$$

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

²⁷Note 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 Appendix A.

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 $v(p) \geq c$.²⁸ 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 (3), to searching among a small subset of all 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 uses at most two experiments, (α_1, q_1) and (α_2, q_2) .*
- (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 $\alpha_1 = \alpha_2 = 0$ (i.e., the sender passes), or $\alpha_1 = 1$ and $q_1 = p_*$ or $q_1 = p^*$;*
 - (iii) *for $p > p^*$, $\alpha_1 = \alpha_2 = 0$ (i.e., the sender passes).*

For part (a) of Proposition 1, notice that the right-hand side in equation (3) is linear in each α_i and the constraint $\sum_{i \in \mathbb{N}} \alpha_i \leq 1$ is also linear. Therefore, by the standard linear programming logic, there exists a solution that makes use of at most two experiments, one below p and the other above p .²⁹ This result implies that

$$v(p) = \max_{(\alpha_1, q_1), (\alpha_2, q_2)} \lambda p(1-p) \left[\alpha_1 \frac{V(q_1) - V(p)}{q_1 - p} - \alpha_2 \frac{V(p) - V(q_2)}{p - q_2} - (\alpha_1 - \alpha_2)V'(p) \right], \quad (4)$$

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 RHS of (4) 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 $q_1 = p^*$

²⁸Suppose 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 $p \in W$.

²⁹One 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 Online Appendix B for details.




Figure 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^* .

whereas the latter strategy leads to $q_2 = 0$ (see p_3 and the dashed lines in Figure 7).³⁰ 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 $p < p^*$.

***R*-drifting experiment (confidence building):** This corresponds to choosing $\alpha_2 = 1$ and $q_2 = 0$. The sender uses an experiment that generates Poisson jumps to 0 at (unconditional) rate $\lambda(1-p)$, upon which the receiver immediately takes action ℓ . In the absence of Poisson jumps, p continuously drifts rightward at rate $\dot{p} = \lambda p(1-p)$. Targeting $q_2 = 0$ is explained by the same logic as selecting zero as the belief that induces action ℓ in the static model: the jump to zero is less likely than a jump to any other $q < p$, as the arrival rate $\lambda p(1-p)/(p-q)$ is increasing in q . Intuitively, this experiment can be interpreted as the strategy of building up the receiver's confidence slowly but steadily while minimizing the chance of a breakdown.

***L*-drifting experiment (confidence spending):** This corresponds to choosing $\alpha_1 = 1$ and $q_1 = p^*$ or possibly $q_1 = p_*$ if $p < p_*$. The sender generates rightward Poisson jumps that lead to either p_* or p^* . In the absence of Poisson jumps, the belief continuously drifts leftward at rate $\dot{p} = -\lambda p(1-p)$. This strategy is the polar opposite of the *R*-drifting experiment. It can yield fast success, but success is unlikely to happen. In addition, when there is no success, the receiver's confidence diminishes—the sender “spends the receiver's confidence.”

³⁰Note 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 $V(p_*)/p_* \leq V'(p_*)$.

Stationary experiment (confidence preserving): This arises when the sender targets two beliefs, $q_1 = p^*$ and $q_2 = 0$, with equal weights ($\alpha_1 = \alpha_2$).³¹ In this case, unless the belief jumps to 0 or p^* , it stays constant (thus, “stationary”). This can be interpreted as a mixture between R -drifting and L -drifting experiments. It can lead to immediate persuasion via a jump to p^* , while not eroding the receiver’s confidence.

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., right-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.³² 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 jumps to p^* until the belief reaches p_*) offers a similar tradeoff 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 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

³¹We will show that any other mixture (in which $\alpha_1 \neq \alpha_2$) never arises in equilibrium.

³²As p tends to p^* , the expected waiting time shrinks to zero under both strategies: under R -drifting, the drift time goes to zero whereas under stationary strategy, the arrival rate of belief p^* goes to infinity. Yet, given our arrival rate specification, the expected waiting time converges to zero faster under the former strategy.

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.

5.2 Equilibrium Characterization

We now explain how the sender’s equilibrium strategy deploys the three modes of persuasion introduced in Section 5.1, 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:

$$p^* \leq \eta \approx 0.943. \tag{C1}$$

This condition determines whether the sender always prefers the R -drifting strategy to the stationary strategy or not. 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 Appendix A.1 for the 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:

$$v > U_r(p^*) - U_\ell(p^*). \tag{C2}$$

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.³³ If (C2) holds, then the sender has a stronger incentive to experiment than the receiver has to listen, so the belief p_* below which some player wishes to stop is determined by the receiver’s incentives. Conversely, if (C2) 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 (C2) 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*

³³As explained in Section 2 (see Footnote 9), the payoffs of the two players are directly comparable, because their information cost c is normalized to be the same. With different information costs, (C2) has to be stated using each player’s payoff relative to their information cost. In the extreme case when the sender’s cost is zero but the receiver’s is not, (C2) 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.

strategy is as follows:³⁴

(a) Suppose 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 0 for all $p \in [p_*, p^*)$.

(ii) If $p^* \in (\eta, 1)$,³⁵ 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 0; for $p = \xi$, she uses the stationary strategy with jumps to 0 and p^* ; and for $p \in (\xi, \bar{\pi}_{LR}]$, she adopts the L -drifting strategy with right-jumps to p^* .

(b) Suppose the belief is outside the waiting region with $p < p_*$. There exist cutoffs $0 < \pi_{\ell L} < \pi_0 < p_*$ such that for $p \leq \pi_{\ell L}$, 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 $c \rightarrow 0$.

Figure 8 below summarizes the sender's SMPE strategy in Proposition 2, depending on whether $p^* < \eta$ or not. 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 the figure. 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 (confidence-building) experiments but switches to the stationary experiment when the belief reaches ξ . For beliefs above ξ , but below $\bar{\pi}_{LR}$, she employs L -drifting (confidence-spending) experiments and also switches to the stationary experiment when the belief reaches ξ .




Figure 8: The sender's SMPE strategies in Proposition 2, that is, when $v > U_r(p^*) - U_\ell(p^*)$.

To understand these different patterns, recall from Section 5.1 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 experiment when p is rather close to p^* , in fact for all p in the waiting region if $p^* < \eta$, but only for p in $[\bar{\pi}_{LR}, p^*)$ if $p^* \geq \eta$. If p^* is above η , then for p below $\bar{\pi}_{LR}$, other experiments become optimal. For $p < \xi$, the sender starts by building confidence, but instead of continuing with

³⁴We 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 $p_* \in W$.

³⁵Notice that in the knife-edge case when $p^* = \eta$, there are two SMPEs, one as in (a.i) and another as in (a.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 $p_0 \in [p_*, \xi]$.

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 $p \in (\xi, \bar{\pi}_{LR})$, 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 0). 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 $p_0 \in [\bar{\pi}_{LR}, p^*)$, 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.

Condition (C2) 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, (C2) 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 Figure 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 only aims to make the buyer stay. If p is rather far away from p_* (below π_0 such as p_1 in Figure 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 (C2) 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:³⁶*

- (a) *Suppose the belief is in the waiting region with $p \in (p_*, p^*)$.*
 - (i) *If $p^* \in (\hat{p}, \eta)$, then there exists a cutoff $\underline{\pi}_{LR} \in W$ such that for $p \in (\underline{\pi}_{LR}, p^*)$, the sender uses the R-drifting strategy with left-jumps to 0; and for $p \in (p_*, \underline{\pi}_{LR})$, she uses the L-drifting strategy with right-jumps to p^* .*
 - (ii) *If $p^* \in (\eta, 1)$, then there exist cutoffs $p_* < \underline{\pi}_{LR} < \xi < \bar{\pi}_{LR} < p^*$ such that for $p \in [\underline{\pi}_{LR}, \xi) \cup [\bar{\pi}_{LR}, p^*)$, the sender plays the R-drifting strategy with left-jumps to 0; for $p = \xi$, she adopts the stationary strategy with jumps to 0 or p^* ; and for $p \in (p_*, \underline{\pi}_{LR}) \cup (\xi, \bar{\pi}_{LR})$, 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 0.

Figure 9 describes the persuasion dynamics in Proposition 3.




Figure 9: The sender's SMPE strategy in Proposition 3, that is, when $v \leq U_r(p^*) - U_\ell(p^*)$.

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

³⁶We 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. See Online Appendix B.

$p \in (\pi_{\ell L}, p_*)$. 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 (C2): if $v \leq U_r(p^*) - U_\ell(p^*)$ 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 (C2) 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 (C2) holds (as in Proposition 2), as the belief nears p_* , the receiver loses interest in listening, but the sender still sees a significant value in staying "in the game." Hence, the sender tries to build, instead of running down, the receiver's confidence in that case.

5.3 Behavioral Implications

Propositions 2 and 3 fully describe equilibrium persuasion dynamics that arise in SMPEs. Here, we explore their implications for several outcome variables. For this purpose, we consider SMPEs with arbitrary persuasion targets $p^* > \hat{p}$, and assume that for each p^* , c is low enough for a unique SMPE to exist.

Persuasion, accuracy, and waiting time. Of particular interest are the sender's persuasion probability, the accuracy of the action taken by the receiver, and the time it takes for the latter to take an action. The results are summarized in the next proposition.

Proposition 4. *Consider the SMPE for any persuasion target $p^* \in (\hat{p}, 1)$.*

- (a) *If (C2) holds or if $p_0 \geq \underline{\pi}_{LR}$, then the posteriors take values in $\{0, p^*\}$, resulting in persuasion with probability p_0/p^* , and the expected waiting time is strictly concave in p_0 for $p_0 \in [\underline{\pi}_{LR}, p^*]$, attaining an interior maximum.*
- (b) *If (C2) fails and $p_0 < \underline{\pi}_{LR}$, then the posteriors take values in $\{p_*, p^*\}$, with persuasion succeeding with probability $\frac{p_0 - p_*}{p^* - p_*}$. If (C2) fails, the expected waiting time is quasi-concave in p_0 for $p_0 \in [p_*, p^*]$ and involves a discontinuous jump at $p_0 = \underline{\pi}_{LR}$.*

Suppose (C2) holds so that part (a) applies. For any prior $p_0 \in (p_*, p^*)$, the sender generates posteriors 0 and p^* , and enjoys persuasion with probability p_0/p^* , as depicted in Panel (a) of Figure 10 below. From the receiver's perspective, this means that his choice of ℓ is perfectly accurate but his choice of r is less accurate. As mentioned earlier, this pattern conforms to the familiar logic of Bayesian persuasion, according to which the sender limits unfavorable information to be sent only when it is perfectly accurate so as to minimize its frequency. In fact, the receiver would have preferred the opposite type




Figure 10: Persuasion and waiting time. In this figure, $p^* = 0.95 > \eta$.

Note: The parameter values used for all panels are $u_\ell^L = u_r^R = 1$, $u_\ell^R = u_r^L = 0$, $\lambda = 1$, and $c = 0.035$. In the left panels, $v = 1$, while $v = 0.3$ in the right panels.

of information when R is unlikely (i.e., when p_0 were low): if the receiver had chosen information himself, he would have settled for less accurate information for L in exchange for more accurate information for R (see Theorem 1 of [Che and Mierendorff, 2019](#)).³⁷

Another interesting observation is that even though the persuasion dynamics the sender employs *do* vary with p_0 to minimize persuasion cost, this does not alter the support of the posterior distribution chosen by the sender. Consequently, a decision takes a longer time, the more uncertain the prior is. This can be seen by the inverse “U”-shaped expected waiting time in Panel (c) of Figure 10.³⁸

Part (b) of the proposition, where (C2) fails, is similar to Part (a), except for one crucial difference. When (C2) fails, the persuasion dynamics “interfere” with the posterior distribution chosen by the sender: for a low range of priors p_0 , the sender generates less efficient beliefs $\{p_*, p^*\}$ than $\{0, p^*\}$. See Panel (b) of Figure 10. If (C2) fails, the

³⁷Specifically, when p_0 is low, a decision maker requires beliefs $p \in (0, \hat{p})$ and 1 for choosing ℓ and r , respectively; likewise, when p_0 is large, he requires beliefs 0 and $\bar{p} \in (\hat{p}, 1)$ for choosing ℓ and r , respectively. Intuitively, he “demands” stronger evidence to justify an action that is unlikely to be optimal. While their model assumes conclusive Poisson signals, their result remains valid within the current class of Poisson signals, as noted in Proposition 6 of [Che and Mierendorff \(2019\)](#).

³⁸This feature is similar to the quasi-concavity of “delay” exhibited by the decision maker in the Wald problem (see Proposition 1 of [Che and Mierendorff, 2019](#)). It is also useful to note that the waiting time is proportional to the persuasion costs incurred by the sender. More precisely, $V(p) = v \cdot [\text{Persuasion prob}] - c \cdot [\text{Expected waiting time}]$. The concavity of the waiting time then follows from the convexity of the value function.




Figure 11: The conditional distribution of the persuasion time. The parameter values used for this figure are identical to those of Figure 10 (with $v = 0.3$).

persuasion cost “looms” large in the mind of the sender, so she sacrifices the persuasion probability to reduce her persuasion costs; this is seen by the discontinuous drop as p_0 falls below $\underline{\pi}_{LR}$ in Panel (d) of Figure 10. This means that as p_0 falls below $\underline{\pi}_{LR}$, the accuracy of action ℓ suffers from the receiver’s perspective. Indeed, the receiver would have preferred that switch to occur at a lower belief, if at all. That is, a *speed-accuracy substitution* occurs at $\underline{\pi}_{LR}$, not for the reason familiar from a single-person decision context,³⁹ but for the *speed-persuasion tradeoff* that motivates the sender’s information choice.

Persuasion time. A novel feature of our dynamic model is a rich prediction it yields for the persuasion time (i.e, how long it takes for the sender to persuade the receiver). Figure 11 shows the distribution of the persuasion time conditional on the sender successfully persuading the receiver, for four different priors.⁴⁰

As evident in the figure, the distribution of the persuasion time varies with prior p_0 . If p_0 is close to p^* (the red curve), the persuasion time is deterministic; it is also back-loaded with the receiver never persuaded before the deterministic time. This pattern is a consequence of the sender adopting the confidence-building strategy that repeats R -drifting experiments. If $p_0 \in (\xi, \bar{\pi}_{LR})$ (the green case), persuasion time is dispersed over time. In fact, the rate of persuasion is high initially but drops at a certain point. This is because the sender initially employs the confidence-spending strategy which repeats L -drifting experiments, but switches to the stationary experiment at ξ . If $p_0 \in (\underline{\pi}_{LR}, \xi)$ (the blue case), then no persuasion occurs for a period of time as the sender repeats R -drifting experiments until the belief reaches ξ ; thereafter persuasion occurs at a certain rate as stationary experiments commence. Finally, if p_0 is close to p_* (the brown case),

³⁹See Ratcliff and McKoon (2008) for a standard explanation and Che and Mierendorff (2019) and Zhong (2019) for the Poisson signal models.

⁴⁰Figure 11 depicts the most sophisticated and comprehensive case where $p^* > \eta$ and (C2) fails. For the other cases, it suffices to consider the relevant subset of four different priors. For example, if $p^* > \eta$ and (C2) holds, then the case when $p_0 = (p_* + \underline{\pi}_{LR})/2$ (the brown curve) is not relevant, while the other three cases effectively stay unchanged.

then persuasion is unlikely to succeed, but when it does, it is relatively front-loaded: the sender plays L -drifting experiments until the belief reaches p_* .

6 Concluding Discussions

We have restricted attention to Markov perfect equilibria, which by definition do not rely on incentives provided by off-path equilibria. Certainly, other (non-Markov) equilibria could be used so as to enlarge the set of sustainable payoffs.⁴¹ Then, it seems plausible that as the players' persuasion costs vanish, one could implement all individually rational payoffs, including the gray 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 simpler 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).⁴² Still, our main economic lessons will continue to apply: all three theorems in Section 4 would continue to hold.⁴³ 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 5.

Our continuous-time game has a straightforward discrete-time analogue and can be interpreted as its limit. One important caveat is that the receiver's per-period (flow) listening cost should be proportional to the amount of information the sender generates,

⁴¹As 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.

⁴²Specifically, 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 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^* .

⁴³The 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 Section 3) holds unchanged with discounting.

as we model in Section 2. If the receiver’s listening cost is independent of the amount of generated information, then our “power-of-beliefs” logic no longer holds in discrete time: 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 0, in which case she would prefer to stop at $p < p^*$. Thus, any persuasion equilibrium with target $p^* > \hat{p}$ would unravel, leaving the no-persuasion equilibrium in Theorem 1 as the unique MPE. If the receiver’s listening cost is proportional to the amount of new information, however, 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 analogue.⁴⁴

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 the folk theorem—namely, the incentivizing power of equilibrium beliefs—applies equally well to such models (see Footnote 24).

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 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 3, it is straightforward that the sender can approximately implement the KG outcome. However, it is non-trivial 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 acquiring information but more importantly *what type* of information to search for. This modeling innovation may fruitfully apply to other dynamic settings.

⁴⁴The same logic applies when there is discounting in terms of the period length Δ . If Δ is independent of the sender’s information structure, then all persuasive SMPs 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.

Appendix A 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 A.1 constructs the value functions that correspond to the equilibrium strategies in Propositions 2 and 3. Sections A.2 and A.3 respectively verify the sender's and the receiver's incentives. Uniqueness of SMPE is proven in Online Appendix C. A brief sketch is provided in Section A.4.

A.1 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 subsection.

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 p 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:⁴⁵

$$c = \lambda p(1-p) \left(\frac{-V_+(p)}{p} + V'_+(p) \right) \quad \text{and} \quad c = \lambda p(1-p) \left(\frac{u_\ell^L - U_+(p)}{p} + U'_+(p) \right). \quad (5)$$

Similarly, suppose for any belief p 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

$$c = \lambda p(1-p) \left(\frac{v - V_-(p)}{p^* - p} - V'_-(p) \right) \quad \text{and} \quad c = \lambda p(1-p) \left(\frac{U_r(p^*) - U_-(p)}{p^* - p} - U'_-(p) \right). \quad (6)$$

R -drifting strategy: Suppose the sender plays R -drifting experiments until the belief reaches p^* . In this case, the players' payoffs are obtained as the solutions to (5) with boundary conditions $V_+(p^*) = v$ and $U_+(p^*) = U_r(p^*)$, respectively. We obtain

$$V_R(p) = \frac{p}{p^*} v - C_+(p; p^*) \quad \text{and} \quad U_R(p) = \frac{p^* - p}{p^*} u_\ell^L + \frac{p}{p^*} U_r(p^*) - C_+(p; p^*),$$

where $C_+(p; q) := \left(p \log \left(\frac{q}{1-q} \frac{1-p}{p} \right) + 1 - \frac{p}{q} \right) \frac{c}{\lambda}$ represents the expected cost of using R -drifting experiments until the belief moves from p to either 0 or q .

⁴⁵The 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.

Stationary strategy: Suppose the sender uses the stationary experiment with jump targets 0 and p^* at p . Then, the players' value functions, $V_S(p)$ and $U_S(p)$, are respectively given by

$$V_S(p) = \frac{p}{p^*}v - C_S(p) \quad \text{and} \quad U_S(p) = \frac{p^* - p}{p^*}u_\ell^L + \frac{p}{p^*}U_r(p^*) - C_S(p), \quad (7)$$

where $C_S(p) := \frac{2c(p^* - p)}{\lambda p^*(1 - p)}$ represents the expected cost of playing the stationary strategy.⁴⁶

RS strategy (R-drifting followed by stationary): Suppose the sender plays the R-drifting strategy until $q (> p)$ and then switches to the stationary strategy. Then, the players' value functions solve (5) with boundary conditions $V_+(q) = V_S(q)$ and $U_+(q) = U_S(q)$, yielding

$$V_{RS}(p; q) = \frac{p}{p^*}v - C_+(p; q) - \frac{p}{q}C_S(q) \quad \text{and} \quad U_{RS}(p; q) = \frac{p^* - p}{p^*}u_\ell^L + \frac{p}{p^*}U_r(p^*) - C_+(p; q) - \frac{p}{q}C_S(q).$$

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 the sender plays the L-drifting strategy until $q (< p)$ and then switches to the stationary strategy. Then, the players' value functions solve (6) with boundary conditions $V_-(q) = V_S(q)$ and $U_-(q) = U_S(q)$, resulting in

$$\begin{aligned} V_{LS}(p; q) &= \frac{p}{p^*}v - C_-(p; q) - \frac{p^* - p}{p^* - q}C_S(q) \quad \text{and} \\ U_{LS}(p; q) &= \frac{p^* - p}{p^*}u_\ell^L + \frac{p}{p^*}U_r(p^*) - C_-(p; q) - \frac{p^* - p}{p^* - q}C_S(q), \end{aligned}$$

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 1 (Crossing Lemma). *Let $V_+(p)$ and $V_-(p)$ be solutions to (5) and (6), respectively.*

(a) *Let $p^* < 8/9$. For all $p < p^*$, if $V_+(p) = V_S(p)$, then $V'_+(p) < V'_S(p)$. Similarly, if*

$V_-(p) = V_S(p)$ then $V'_-(p) < V'_S(p)$.

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

⁴⁶Under the stationary strategy, the total arrival rate of Poisson jumps is equal to $\lambda_S(p) = \frac{\lambda}{2}(1 - p) + \frac{\lambda}{2} \frac{p(1 - p)}{p^* - p} = \frac{\lambda}{2} \frac{p^*(1 - p)}{p^* - p}$. $C_S(p)$ is equal to c times the expected arrival time $1/\lambda_S(p)$.

- (i) For all $p < p^*$, if $V_+(p) = V_S(p)$, then $V'_+(p) = V'_S(p)$ if and only if $p \in \{\xi_1, \xi_2\}$, and $V'_+(p) > V'_S(p)$ if and only if $p \in (\xi_1, \xi_2)$;
- (ii) For all $p < p^*$, if $V_-(p) = V_S(p)$, then $V'_-(p) = V'_S(p)$ if and only if $p \in \{\xi_1, \xi_2\}$, and $V'_-(p) > V'_S(p)$ if and only if $p \in (\xi_1, \xi_2)$.

(c) For all $p < p^*$, if $V_+(p) = V_-(p)$, then $\text{sign}(V'_+(p) - V'_-(p)) = \text{sign}(V_-(p) - V_S(p))$. All parts also hold for the receiver's value functions $U_+(\cdot)$, $U_-(\cdot)$, and $U_S(\cdot)$.

Proof. We focus on the sender's value functions, as the same proofs apply to the receiver. From (5), (6), and (7), we can obtain expressions for $V'_+(p)$, $V'_-(p)$, and $V'_S(p)$. Combining these with $V_+(p) = V_S(p)$ and $V_-(p) = V_S(p)$, we obtain

$$V'_+(p) - V'_S(p) = V'_-(p) - V'_S(p) = -\frac{c(2p^2 - 3p^*p + p^*)}{\lambda p^*p(1-p)^2} \begin{matrix} \geq \\ \leq \end{matrix} 0 \Leftrightarrow -2p^2 + 3p^*p - p^* \begin{matrix} \geq \\ \leq \end{matrix} 0.$$

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 (b).

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

$$V'_+(p) - V'_-(p) = \frac{p^*}{p(p^* - p)} (V_-(p) - V_S(p)),$$

which leads to (c). □

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 $V_{RS}(p; \xi)$ meets $V_S(p)$ from above at $p = \xi$ (as p rises toward ξ). In particular, together with the Crossing Lemma 1.(b), 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 $V_R(\xi(p^*)) = V_S(\xi(p^*))$.⁴⁷ (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 2.

- (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 $V_R(\xi) < V_S(\xi)$.

⁴⁷ To show that η is well-defined, we can define a function $g : (8/9, 1) \rightarrow \mathbb{R}$ 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 $p^* \in (8/9, 1)$.

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 1.(b), 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 Footnote 47). \square

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

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

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

Lemma 3.

- (a) Both $\widetilde{V}(p)$ and $\widetilde{U}(p)$ are strictly convex in p over $[0, p^*]$.
- (b) If $p^* \leq \eta$ then $\widetilde{V}(p) = V_R(p)$ and $\widetilde{U}(p) = U_R(p)$ for all $p \in [0, p^*]$.
- (c) If $p^* > \eta$ then there exists $\bar{\pi}_{LR} \in (\xi, p^*)$ such that $\widetilde{V}(p) = \widehat{V}(p)$ and $\widetilde{U}(p) = \widehat{U}(p)$ for $p \leq \bar{\pi}_{LR}$ and $\widetilde{V}(p) = V_R(p)$ and $\widetilde{U}(p) = U_R(p)$ for $p \in [\bar{\pi}_{LR}, p^*]$.
- (d) $\widetilde{V}(p) \geq V_S(p)$ 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 $\widehat{V}(p) = V_{RS}(p; p^*) = V_R(p)$, which implies (b). Since $V_R(p)$ is strictly convex, (a) holds as well. In what follows, we consider $p^* \geq 8/9$, in which case $\widehat{V}(p) \neq V_R(p)$.

(a) Since $\widetilde{V}(p)$ is the upper envelope of two functions and $V_R(p)$ is strictly convex over $[0, p^*]$, it suffices to prove that $\widehat{V}(p)$ 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 $\widehat{V}(p)$ 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 1.(c).

(b) If $p^* < \eta$, $V_{RS}(\xi; \xi) = V_S(\xi) < V_R(\xi)$ by Lemma 2.(a). Together with the fact that both $V_{RS}(p; \xi)$ and $V_R(p)$ satisfy the ODE (5), this implies that $\widehat{V}(p) = V_{RS}(p; \xi) < V_R(p)$ for all $p \leq \xi$.⁴⁸ For $p \in (\xi, p^*]$, observe that $V_{LS}(\xi; \xi) = V_S(\xi) < V_R(\xi)$ (Lemma 2.(a)),

⁴⁸It is easy to see that (5) satisfies the Lipschitz condition for uniqueness on $(0, p^*)$.

$V_{LS}(p^*; \xi) = V_R(p^*)$, and $V'_{LS}(p^*; \xi) > V'_R(p^*)$. Therefore, either $\widehat{V}(p) = V_{LS}(p; \xi) < V_R(p)$ 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 1.(c), $V_R(p) = V_{LS}(p; \xi) = V_-(p) < V_S(p)$, contradicting Lemma 2.(a).

The result for $p^* = \eta$ follows from a continuity argument: both $\widehat{V}(p)$ and $V_R(\cdot)$ change continuously in p^* . Since $\widehat{V}(p) < V_R(p)$ for all $p < p^*$ whenever $p^* < \eta$, it must be that $\widehat{V}(p) \leq V_R(p)$ 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) $\widehat{V}(p) \geq V_S(p)$ 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 $\widehat{V}(p) = V_{RS}(p; \xi) > V_S(p)$. To see this, pick $q < \xi$. Then by Lemma 1.(b).(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 1.(b).(i) then implies that $V_R(p)$ and $V_S(p)$ cannot intersect at $p \geq \xi_2$. \square

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 2.(c)), which immediately implies that $\widehat{V}(p) = V_{RS}(p; \xi) > V_R(p)$ 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 $p \in (\xi, p^*)$, $V_{LS}(p; \xi) = \widehat{V}(p) = V_R(p) > V_S(p)$, and hence by Lemma 1.(c), $V_{LS}(p; \xi)$ 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 2.(a) and 3.(b). If $p^* > \eta$ the result is immediate from Lemma 3.(c) and Claim 1. \square

A.1.1 Equilibrium payoffs and construction of p_* in Proposition 2

When (C2) 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 by⁴⁹

$$U_\ell(\phi_{\ell R}) = \widetilde{U}(\phi_{\ell R}). \quad (8)$$

We focus on the case in which c is sufficiently small. In the limit as $c \rightarrow 0$, $\widetilde{U}(p) = \frac{p^* - p}{p^*} u_\ell^L + \frac{p}{p^*} U_r(p^*) > U_r(p)$ for all p . Therefore, there exists $c_1 > 0$ such that $p_* = \phi_{\ell R} < \hat{p}$

⁴⁹To see that $\phi_{\ell R}$ is well defined, observe that, whether $p^* \leq \eta$ or $p^* > \eta$, $\lim_{p \rightarrow 0} \widetilde{U}(p) = u_\ell^L - \frac{c}{\lambda} < u_\ell^L = U_\ell(0)$, while $\widetilde{U}(p^*) = U_r(p^*) > U_\ell(p^*)$ (because $p^* > \hat{p}$). In addition, $\widetilde{U}(p)$ is strictly convex over $[0, p^*]$ (Lemma 3.(a)), while $U_\ell(p)$ is linear. Therefore, $\widetilde{U}(p)$ crosses $U_\ell(p)$ from below only once.

for all $c \leq c_1$. We assume that $c \leq c_1$ in the sequel. The following Lemma shows that the sender's payoff is positive at p_* if Condition (C2) holds.

Lemma 4. $\tilde{V}(\phi_{\ell R}) > 0$ if and only if Condition (C2) holds.

Proof. By (8), we have

$$\tilde{V}(\phi_{\ell R}) = \frac{\phi_{\ell R}}{p^*}v + \tilde{U}(\phi_{\ell R}) - \left(\frac{p^* - \phi_{\ell R}}{p^*}u_\ell^L + \frac{\phi_{\ell R}}{p^*}U_r(p^*) \right) \stackrel{(8)}{=} \frac{\phi_{\ell R}}{p^*} (v - (U_r(p^*) - U_\ell(p^*))),$$

where the first equality holds because both players incur the same costs, so that $\tilde{V}(p) - \frac{p}{p^*}v = \tilde{U}(p) - \left(\frac{p^* - p}{p^*}u_\ell^L + \frac{p}{p^*}U_r(p^*) \right)$ whenever $p \in (0, p^*)$. The last expression is positive if and only if (C2) holds. \square

We set the players' value functions as follows:

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

Lemma 5. When (C2) 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 $\tilde{V}(p_*) \geq 0$ by Lemma 4, $\tilde{V}(\cdot)$ is increasing on $[p_*, p^*]$. Hence $V(\cdot)$ is nondecreasing on $[0, 1]$, and nonnegative since $V(0) = 0$. \square

A.1.2 Equilibrium payoffs and construction of p_* in Proposition 3

When (C2) fails, the same construction as above does not work; for example, $\tilde{V}(p_*) < 0$ by Lemma 4. The right construction requires us to consider another L -drifting strategy.

$L0$ strategy (L -drifting followed by passing): Suppose 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 (6) with boundary conditions $V_-(q) = 0$ and $U_-(q) = U_\ell(q)$, which yields

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

Note that this strategy leads to q with probability $\frac{p^* - p}{p^* - q}$ and p^* with probability $\frac{p - q}{p^* - q}$.

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

$$V'_{L0}(q; q) \geq 0 \Leftrightarrow \frac{\lambda q(1 - q)}{p^* - q}v \geq c \Leftrightarrow q \geq \pi_{\ell L} := \frac{1}{2} + \frac{c}{2\lambda v} - \sqrt{\left(\frac{1}{2} + \frac{c}{2\lambda v} \right)^2 - \frac{cp^*}{\lambda v}}. \quad (9)$$

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 (C2) fails, we set $p_* := \pi_{\ell L}$. Clearly, $\lim_{c \rightarrow 0} p_* = 0$. 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 6. *Suppose (C2) 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 (b).

(a) Suppose $p^* \leq \eta$ so that $\tilde{V}(p) = V_R(p)$ for all $p \leq p^*$. Since (9) 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 $f_1(p) := \log \left(\frac{p^*}{1-p^*} \frac{1-p}{p} \right) - \frac{p^* - p}{p^*(1-p)}$. 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,

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

Define $f_2(p) := \frac{2p(p^* - \xi)}{p^2 \xi (1 - \xi)} + p \log \left(\frac{\xi}{1 - \xi} \frac{1 - p}{p} \right) + 1 - \frac{p}{\xi} - \frac{p^* - p}{p^*(1 - p)}$. 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 3.(d)), it suffices to show $V_{L0}(x; p_*) < V_S(x)$. Indeed, we have $V_{L0}(x; p_*) - V_S(x) = \left(\frac{x - p_*}{p^* - p_*} - \frac{x}{p^*} \right) v + C_S(x) - C_-(x; p_*) < C_S(x) - C_-(x; p_*)$, since $C_S(x)/c$ is independent of c and $C_-(x; p_*)/c \rightarrow \infty$ as $c \rightarrow 0$.⁵⁰

By Lemma 6.(a) 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^*) .

We first consider $p^* < \eta$. In this case $\tilde{V}(p) = V_R(p)$ and Lemma 2 implies that $V_R(p) > V_S(p)$. Then, by Lemma 1.(c), $V_{L0}(p; p_*)$ 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 (6), no intersection can occur in the interval $[\xi, \bar{\pi}_{LR}]$. Outside this interval $\tilde{V}(p)$ satisfies (5) and $\tilde{V}(p) > V_S(p)$ by Lemma 3.(d). Therefore, again Lemma 1.(c) implies that $V_{L0}(p; p_*)$ can cross $\tilde{V}(p)$ only from above. \square

⁵⁰This is because $p_* \rightarrow 0$ as $c \rightarrow 0$ so that for the $L0$ strategy the expected waiting time from any starting point x becomes infinite if the state is L .

A.1.2.1 Equilibrium payoffs. The equilibrium value functions are given as follows:

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

Lemma 7. *When (C2) fails, $V(\cdot)$ is nonnegative and nondecreasing on $[0, p^*]$, and strictly convex on $[p_*, p^*]$*

Proof. Lemma 6.(b) 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 3.(a). Finally, by (9), $V(p)$ is continuously differentiable at $p_* = \pi_{\ell L}$ and therefore convex on $[0, p^*]$. This also implies that $V(p)$ is nondecreasing. \square

A.2 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 $p \in W$.⁵¹ To this end, we must show that in the waiting region the sender's equilibrium value function solves the Hamilton-Jacobi-Bellmann (HJB) equation:⁵²

$$\max_{(\alpha_i, q_i)_{i \in \mathcal{I}} \in \mathcal{I}} \sum_{q_i \neq p} \alpha_i v(p; q_i) = c, \quad (\text{HJB})$$

where $\mathcal{I} := \{(\alpha_i, q_i)_{i \in \mathbb{N}} \mid \alpha_i \geq 0; \sum_{i=1}^{\infty} \alpha_i \leq 1; q_i \in [0, 1]\}$ denotes the set of feasible information structures and $v(p; q_i)$ is as defined in Section 5.1. Outside the waiting region, the sender's value is independent of her strategy. Still, our refinement requires that her strategy maximize her flow payoff; that is, her choice of experiment should solve

$$\max_{(\alpha_i, q_i)_{i \in \mathbb{N}} \in \mathcal{I}} \sum_{q_i \neq p} \alpha_i (v(p; q_i) - c). \quad (\text{Ref})$$

Proposition 1.(b) implies that if $V(p)$ meets certain conditions, then we can restrict attention to Poisson experiments with jump targets, 0 , p_* , and p^* , which greatly simplifies both (HJB) and (Ref). Here, we show that our equilibrium value function $V(\cdot)$ satisfies all properties required by Proposition 1.(b), namely that it is nonnegative, increasing, and strictly convex on $(p_*, p^*]$, and $V(p_*)/p_* \leq V'(p_*)$. If (C2) holds, the first two

⁵¹Recall that $W = [p_*, p^*)$ in Proposition 2, and $W = (p_*, p^*)$ in Proposition 3.

⁵²More 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.

properties hold by Lemma 5, strict convexity of $V(p)$ follows from Lemma 3.(a) and $V(p) = \tilde{V}(p)$ for $p \in [p_*, p^*]$. The last property also holds because $\tilde{V}(p)$ is convex and $\lim_{p \rightarrow 0} \tilde{V}(p) = -c/\lambda < 0$. If (C2) fails, the first three properties follow from Lemma 7 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 1.(b) to the stopping region and verify (Ref). For $p \geq p^*$, the result is immediate from Proposition 1.(b).(iii). Now consider p below p_* . Proposition 1.(b).(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 (C2) 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 (9), 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 (C2) holds, then as discussed in Section 5.1 and depicted in Figure 7 there exists a cutoff $\pi_0 < p_*$ such the coefficient of α_* is greater than the coefficient of α^* if and only if $p > \pi_0$.⁵³ The following Lemma shows that $\pi_{\ell L} < \pi_0$.

Lemma 8. *If (C2) 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 3.(a)) 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 6.(a). \square

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 satisfied (Ref) for all $p < p_*$.

Waiting region. Applying Proposition 1.(b).(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:

⁵³Specifically π_0 satisfies

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

Lemma 9 (Unimprovability).

- (a) If $V_+(p)$ satisfies (5) 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).
- (b) If $V_-(p)$ satisfies (6) 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) Substituting $V'(p) = V'_+(p)$ from (5), (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 (b) is similar. \square

By Lemmas 3.(d) and 6.(b), $V(p) \geq V_S(p)$ holds for all $p \in (p_*, p^*)$. Therefore, the Unimprovability Lemma 9 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 (5) and (6), respectively, if we replace V'_+ by the right derivative and V'_- by the left derivative. As in the proof of the Unimprovability Lemma 9 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, it can be shown that $V(p)$ is a viscosity solution of (HJB), which is sufficient for optimality of the sender's strategy in the waiting region (see Footnote 52 above).

A.3 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 :⁵⁴

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

and

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

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.⁵⁵

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.⁵⁶ 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

$$U'(p^*) = U'_R(p^*) = \frac{U_r(p^*) - u_\ell^L}{p^*} + \frac{c}{\lambda p^*(1-p^*)} = U'_r(p^*) - \frac{u_\ell^L - u_r^L}{p^*} + \frac{c}{\lambda p^*(1-p^*)}.$$

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

(i) Proposition 2: When (C2) 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 3.(a)). Together with $U'(p^*) \leq U'_r(p^*)$, this implies that $U(p) \geq U_r(p)$ for all $p \in [p_*, p^*]$, provided that $c \leq c_4$. We have argued in Footnote 49 that $U(p) \geq U_\ell(p)$ 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

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

⁵⁴The receiver's value function $U(p)$ is not continuously differentiable at p_* (in case (C2) holds), $\underline{\pi}_{LR}$, and $\bar{\pi}_{LR}$. At these non-smooth 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 which are closed under the belief dynamics given by the sender's strategy. For example if (C2) 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 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.

⁵⁵Specifically, $\alpha(p) = 0$ 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.

⁵⁶At kinks, $U(p)$ satisfies (R1) if $U'(p)$ is replaced by $U'(p_+)$ (see footnote 54).

where the inequality holds since (C2) fails, and the second equality follows from (9) 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 6.(b), $\underline{\pi}_{LR} \leq \hat{p}$, provided that $c \leq c_3$. 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 (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

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

The inequality holds since by Lemma 6.(b), $\tilde{V}(p) - V_{L0}(p; p_*) \geq 0$ for $p \geq \underline{\pi}_{LR}$; and $U_r(p^*) - U_\ell(p^*) - v > 0$ if (C2) is violated.

The stopping region with $p \in (\pi_{\ell L}, p_*)$. If (C2) fails, then $p_* = \pi_{\ell L}$, so this case does not arise. The proof of Proposition 3 is thus complete.

Now suppose that (C2) 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, $(\alpha(p), q(p)) = (1, p_*)$. Plugging this into (R1) and using 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 that $\phi_{\ell L} \geq \pi_0$ if $c \leq c_5$ for some $c_5 > 0$. It then follows that if $c \leq \min\{c_1, \dots, c_5\}$, the receiver has no incentive to deviate from his prescribed strategy in Proposition 2, completing the proof.

Lemma 10. *Suppose (C2) 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 $\frac{p(1-p)\lambda\Delta U}{p^*-p} \geq c$,

$$\pi_0 \leq \phi_{\ell L} \iff \frac{\pi_0(1-\pi_0)\lambda}{p^*-\pi_0} \Delta U < 1. \quad (10)$$

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

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

where $\tilde{C}(p_*) = \frac{p^*}{p^*} v - \tilde{V}(p_*)$ denotes the total persuasion costs incurred when $p = p_* = \phi_{\ell R}$ and (C2) 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 $c \rightarrow 0$, we have $p_* \rightarrow 0$, $\tilde{C}(p_*) \rightarrow 0$ and $\tilde{C}(p_*) \frac{\lambda}{c} \rightarrow 1$. It follows that $\pi_0 \rightarrow 0$ and $\pi_0 \frac{\lambda}{c} \rightarrow \frac{p^*}{v}$, so

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

where the inequality is due to (C2). This completes the proof. \square

A.4 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 A.1.1 and A.1.2 (i.e., $p_* \neq \phi_{\ell R}$ if (C2) holds and $p_* \neq \pi_{\ell L}$ if (C2) 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 A.1 and A.2, and the latter follows closely A.3. We relegate the full proof to the Online Appendix C.

References

- ARROW, K. J., D. BLACKWELL, AND M. A. GIRSHICK (1949): "Bayes and Minimax Solutions of Sequential Decision Problems," *Econometrica*, pp. 213–244. 6
- AU, P. H. (2015): "Dynamic Information Disclosure," *RAND Journal of Economics*, 46, 791–823. 5
- BERGIN, J., AND W. B. MACLEOD (1993): "Continuous Time Repeated Games," *International Economic Review*, 34(1), 21–27. 12
- BEST, J., AND D. QUIGLEY (2020): "Persuasion for the Long Run," unpublished manuscript, University of Oxford. 6

- BIZZOTTO, J., J. RÜDIGER, AND A. VIGIER (2020): “Dynamic persuasion with outside information,” *Available at SSRN 3154313*. 5
- BROCAS, I., AND J. D. CARRILLO (2007): “Influence through ignorance,” *The RAND Journal of Economics*, 38(4), 931–947. 5
- CAPLIN, A., M. DEAN, AND J. LEAHY (2013): “Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy,” . 10
- CHE, Y.-K., AND J. HÖRNER (2018): “Recommender Systems as Mechanisms for Social Learning,” *Quarterly Journal of Economics*, 133, 871–925. 5
- CHE, Y.-K., AND K. MIERENDORFF (2019): “Optimal Dynamic Allocation of Attention,” *American Economic Review*, 109, 2993–3029. 3, 6, 28, 29
- CHE, Y.-K., AND J. SÁKOVICS (2004): “A dynamic theory of holdup,” *Econometrica*, 72(4), 1063–1103. 4
- ELY, J. C. (2017): “Beeps,” *American Economic Review*, 107, 31–53. 5
- ELY, J. C., AND M. SZYDŁOWSKI (2020): “Moving the goalposts,” *Journal of Political Economy*, 128(2), 468–506. 5
- FRANKEL, A., AND E. KAMENICA (2019): “Quantifying information and uncertainty,” *American Economic Review*, 109(10), 3650–80. 10
- FUDENBERG, D., P. STRACK, AND T. STRZALECKI (2018): “Speed, Accuracy, and the Optimal Timing of Choices,” *American Economic Review*, 108, 3651–84. 6
- HENRY, E., AND M. OTTAVIANI (2019): “Research and the approval process: the organization of persuasion,” *American Economic Review*, 109(3), 911–55. 5, 6, 16, 31
- KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101(6), 2590–2615. 1, 2, 3, 5
- KE, T. T., AND J. M. VILLAS-BOAS (2019): “Optimal learning before choice,” *Journal of Economic Theory*, 180, 383–437. 6
- KELLER, G., AND S. RADY (2015): “Breakdowns,” *Theoretical Economics*, 10(1), 175–202. 3, 11
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 73(1), 39–68. 3, 11
- KLEIN, N., AND S. RADY (2011): “Negatively Correlated Bandits,” *The Review of Economic Studies*, 78(2), 693–732. 1
- KREMER, I., Y. MANSOUR, AND M. PERRY (2014): “Implementing the “Wisdom of the Crowd”,” *Journal of Political Economy*, 122, 988–1012. 5
- MARGARIA, C., AND A. SMOLIN (2018): “Dynamic Communication with Biased Senders,” *Games and Economic Behavior*, 110, 330–339. 6
- MATEJKA, F., AND A. MCKAY (2015): “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model,” *The American Economic Review*, 105(1), 272–298. 10
- MATHEVET, L., D. PEARCE, AND E. STACHETTI (2019): “Reputation and Information Design,” unpublished manuscript, New York University. 6
- MAYSKAYA, T. (2019): “Dynamic choice of information sources,” . 6
- MOSCARINI, G., AND L. SMITH (2001): “The Optimal Level of Experimentation,” *Econometrica*, 69(6), 1629–1644. 6
- NIKANDROVA, A., AND R. PANCS (2018): “Dynamic Project Selection,” *Theoretical Economics*, 13(1), 115–144. 6
- OKSENDAL, B., AND A. SULEM (2009): *Applied Stochastic Control of Jump Diffusions*. Springer, 3rd edition edn. 39

- ORLOV, D., A. SKRZYPACZ, AND P. ZRYUMOV (2020): “Persuading the principal to wait,” *Journal of Political Economy*, 128(7), 000–000. 5
- RATCLIFF, R., AND G. MCKOON (2008): “The Diffusion Decision Model: Theory and Data for Two-choice Decision Tasks,” *Neural computation*, 20(4), 873–922. 29
- RENAULT, J., E. SOLAN, AND N. VIEILLE (2017): “Optimal Dynamic Information Provision,” *Games and Economic Behavior*, 104(4), 329–349. 5
- SIMON, L. K., AND M. B. STINCHCOMBE (1989): “Extensive Form Games in Continuous Time: Pure Strategies,” *Econometrica*, 57(5), 1171–1214. 12, 30
- SIMS, C. A. (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50(3), 665–690. 10
- SONER, H. M. (1986): “Optimal Control with State Space Constraint II,” *SIAM Journal of Control and Optimization*, 24(6), 1110–1122. 39
- WALD, A. (1947): “Foundations of a General Theory of Sequential Decision Functions,” *Econometrica*, pp. 279–313. 6
- ZHONG, W. (2019): “Optimal Dynamic Information Acquisition,” unpublished manuscript, Stanford University. 6, 29

Supplemental Material

(for online publication)

Appendix B Formal Description of the Game in Continuous Time

An information structure was defined as a collection $(\alpha_i, q_i)_{i \in \mathbb{N}}$ of experiments, where each (α_i, q_i) specifies a Poisson experiment with jump-targets q_i and associated weight α_i . The set of feasible information structures is thus given by

$$\mathcal{I} = \left\{ (\alpha_i, q_i)_{i \in \mathbb{N}} \mid \alpha_i \geq 0; \sum_{i=1}^{\infty} \alpha_i \leq 1; q_i \in [0, 1] \right\}.$$

We define a game in Markov strategies. The sender's strategy is a measurable function $\sigma^S : [0, 1] \rightarrow \mathcal{I}$ that maps the belief p to an information structure $\sigma^S(p)$.⁵⁷ The receiver's strategy is a measurable function $\sigma^R : [0, 1] \rightarrow A$, that maps the belief p to an action $\sigma^R(p) \in A := \{\ell, r, w\}$. We impose the following *admissibility* restrictions in order to ensure that a strategy profile $\sigma = (\sigma^S, \sigma^R)$ yields a well defined outcome.

Admissible Strategies for the Sender. Our first restriction ensures that σ^S gives rise to a well-defined evolution of the (common) belief about the state.⁵⁸ For a Markov strategy $\sigma^S(p)$, with experiments $(\alpha_i(p; \sigma^S), q_i(p; \sigma^S))_{i \in \mathbb{N}}$, Bayesian updating leads to the following integral equation for the belief p_t (conditional on non-arrival):⁵⁹

$$p_t = \frac{p_0 e^{-\lambda \int_0^t (\alpha_s^+ - \alpha_s^-) ds}}{p_0 e^{-\lambda \int_0^t (\alpha_s^+ - \alpha_s^-) ds} + (1 - p_0)}, \quad (11)$$

where $\alpha_t^+ = \sum_{i: q_i(p_t; \sigma^S) > p_t} \alpha_i(p_t; \sigma^S)$ is the total weight on upward jumps at time t , and $\alpha_t^- = \sum_{i: q_i(p_t; \sigma^S) < p_t} \alpha_i(p_t; \sigma^S)$ is the total weight on downward jumps at time t . To define admissibility formally, we also introduce the following discrete time approximation. For

⁵⁷We can take \mathcal{I} to be a subset of $\mathbb{R}^{2\mathbb{N}}$, the set of sequences $(\alpha_1, q_1), (\alpha_2, q_2), \dots$ in \mathbb{R}^2 , with the product σ -algebra $\mathcal{B}(\mathbb{R}^{2\mathbb{N}}) = \mathcal{B}(\mathbb{R}^2) \otimes \mathcal{B}(\mathbb{R}^2) \otimes \dots$, where $\mathcal{B}(\mathbb{R}^2)$ is the Borel σ -algebra on \mathbb{R}^2 .

⁵⁸In this part, we follow [Klein and Rady \(2011\)](#), with the difference that in their model, the evolution of beliefs is jointly controlled by two players. Given that in our model, only the sender controls the information structures, we can dispense with their assumption that Markov strategies are constant on the elements of a finite interval partition of the state space.

⁵⁹The corresponding differential equation is given by $\dot{p}_t = -(\alpha_t^+ - \alpha_t^-) \lambda p_t (1 - p_t)$.

period length $\Delta > 0$, let

$$\tilde{p}_{(k+1)\Delta} = \frac{\tilde{p}_k e^{-\lambda\Delta(\alpha_{k\Delta}^+ - \alpha_{k\Delta}^-)}}{\tilde{p}_k e^{-\lambda\Delta(\alpha_{k\Delta}^+ - \alpha_{k\Delta}^-)} + (1 - \tilde{p}_k)}.$$

This can be used to define $\tilde{p}_{k\Delta}$ recursively for each $\tilde{p}_0 = p_0$, and yields a step-function $p_t^\Delta := \tilde{p}_{\lfloor t/\Delta \rfloor \Delta}$.

Definition 2. A measurable function $\sigma^S : [0, 1] \rightarrow \mathcal{I}$ is an *admissible strategy for the sender* if for all $p_0 \in [0, 1]$,

- (a) there exists a solution to (11), and
- (b) if there are multiple solutions to (11), then the pointwise limit $\lim_{\Delta \rightarrow 0} p_t^\Delta$ exists and solves (11).

This definition imposes two restrictions on Markov strategies. First, there must be a solution to (11). Indeed, there are Markov strategies for which no solution exists. Consider, for example, a strategy of the following form:

$$\sigma^S(p) = \begin{cases} (\alpha = 1, q = 1), & \text{if } p \geq p', \\ (\alpha = 1, q = 0), & \text{if } p < p'. \end{cases}$$

This strategy does not lead to a well-defined evolution of the belief if p_0 is given by the “absorbing belief” $p_0 = p'$. To satisfy admissibility, we can set $\sigma^S(p') = ((\alpha_1 = 1/2, q_1 = 1), (\alpha_2 = 1/2, q_1 = 0))$, while keeping the strategy otherwise unchanged.

The second restriction guarantees that if there are multiple solutions, we can select one of them by taking the pointwise limit of the discrete time approximation. Consider, for example, the following strategy:

$$\sigma^S(p) = \begin{cases} (\alpha = 1, q = 0), & \text{if } p \geq p', \\ (\alpha = 1, q = 1), & \text{if } p < p'. \end{cases} \quad (12)$$

If $p_0 = p'$, then there is an “obvious” solution $p_t^1 = \frac{p' e^{\lambda t}}{p' e^{\lambda t} + (1 - p')} > p'$ for $t > 0$. However, there exists another solution $p_t^2 = \frac{p' e^{-\lambda t}}{p' e^{-\lambda t} + (1 - p')}$ consistent with $p_0 = p'$. But, in discrete time, $\tilde{p}_\Delta > p'$ for any $\Delta > 0$, and thus $\lim_{\Delta \rightarrow 0} p_t^\Delta = p_t^1$. This means that the strategy in (12) is admissible, while the latter strategy with p_t^2 is not. In general, when there are multiple solutions, admissibility enables us to select the “obvious” one that would be obtained from the discrete time approximation. With this selection, admissibility of the sender’s strategy guarantees a well defined belief for all $t > 0$ and all prior beliefs p_0 .

Admissible Strategy Profiles. In addition to a well defined evolution of beliefs, we need to ensure that a strategy profile $\sigma = (\sigma^S, \sigma^R)$ leads to a well defined stopping time

for any initial belief p_0 . Consider for example the function

$$\sigma^R(p) = \begin{cases} w & \text{if } p \leq p', \\ r & \text{if } p > p'. \end{cases}$$

If the sender uses the (admissible) Markov strategy given by $\sigma^S(p) = (\alpha = 1, q = 0)$ for all p , and the prior belief is $p_0 < p'$, then the function $\sigma^R(p)$ does not lead to a well-defined stopping time. To be concrete, suppose that the true state is $\omega = R$. In this case, no Poisson jumps occur, and the belief drifts upwards. Let t' denote the time at which the belief reaches p' . The receiver's strategy implies that for any $t \leq t'$, the receiver plays w and for any $t > t'$, the receiver has stopped before t . Hence, the stopping time is not well defined. Clearly, the following modified strategy fixes the problem:

$$\hat{\sigma}^R(p) = \begin{cases} w & \text{if } p < p', \\ r & \text{if } p \geq p'. \end{cases}$$

This example demonstrates that we need a joint restriction on the sender's and the receiver's strategies to ensure a well defined outcome.

To formally define admissibility, we need the following notation: for a given strategy of the receiver σ^R , let $W = \{p \in [0, 1] \mid \sigma^R(p) = w\}$ and $S = [0, 1] \setminus W$ be the receiver's *waiting region* and *stopping region*, respectively, and denote the closures of these sets by \overline{W} and \overline{S} .

Definition 3. A strategy profile $\sigma = (\sigma^S, \sigma^R)$ is *admissible* if (i) σ^S is an admissible strategy for the sender, and (ii) for each $p \in \overline{W} \cap \overline{S}$, either $p \in S$, or if $p \notin S$, then there exists $\varepsilon > 0$ such that $p_t(p) \in W$ for all $t < \varepsilon$, where $p_t(p) = \lim_{\Delta \rightarrow 0} p_t^\Delta$ is the selected solution to (11) with $p_0 = p$.

Requirement (i) guarantees that the sender's strategy gives rise to a well defined belief at all $t > 0$ for all prior beliefs regardless of the receiver's strategy. Requirement (ii) ensures that for any belief $p \in W$, the belief evolution is such that absent jumps the belief remains in the waiting region.

One may wonder why we do not simply require that the stopping region is a closed set. This is stronger than requirement (ii) and it turns out that in some cases it can lead to non-existence of an equilibrium.⁶⁰

Payoffs and Equilibrium. Let $\sigma = (\sigma^S, \sigma^R)$ be a profile of strategies. If σ is not admissible, then both players receive $-\infty$ from playing the strategy profile. If σ is admissible,

⁶⁰For example, in the equilibrium characterized in Proposition 2, we have $S = [0, p_*) \cup [p^*, 1]$. If we require S to be closed and set $S = [0, p_*] \cup [p^*, 1]$ instead, the sender does not have a best response for $p \in (\pi_0, p_*)$ since $v(p; q_i)$ fails to be upper semi-continuous in q_i at $q_i = p_*$.

then for each prior belief p_0 , both players' expected payoffs are well defined:

$$V^\sigma(p_0) = v \mathbb{P} [\sigma^R(p_\tau) = r | p_0] - c \mathbb{E} \left[\int_0^\tau \mathbf{1}_{\{\sum \alpha_i(p_t) \neq 0\}} dt \middle| p_0 \right]$$

for the sender, and

$$U^\sigma(p_0) = \mathbb{E} \left[U_{\sigma^R(p_\tau)}(p_\tau) - c \int_0^\tau \mathbf{1}_{\{\sum \alpha_i(p_t) \neq 0\}} dt \middle| p_0 \right]$$

for the receiver, where τ is the stopping time defined by the strategy profile and p_τ is the belief when the receiver stops.

Definition 4 (Markov Perfect Equilibrium). An admissible strategy profile $\sigma = (\sigma^S, \sigma^R)$ is a *Markov perfect equilibrium* (MPE), if

- (i) for any $p_0 \in [0, 1]$ and any admissible strategy profile $\hat{\sigma} = (\hat{\sigma}^S, \sigma^R)$, $V^{\hat{\sigma}}(p_0) \leq V^\sigma(p_0)$,
- (ii) for any $p_0 \in [0, 1]$ and any admissible strategy profile $\hat{\sigma} = (\sigma^S, \hat{\sigma}^R)$, $U^{\hat{\sigma}}(p_0) \leq U^\sigma(p_0)$,
- and
- (iii) for any $p \in S$: (refinement)

$$\sigma^S(p) \in \arg \max_{(\alpha_i, q_i) \in \mathcal{I}} \sum_{i: q_i \neq p} \alpha_i \frac{\lambda p(1-p)}{|q_i - p|} (V^\sigma(q_i) - \mathbf{1}_{\{\sigma^R(p)=r\}} v) - \mathbf{1}_{\{\sum \alpha_i \neq 0\}} c.$$

Parts (i) and (ii) in this definition require that no player have a profitable deviation to a Markov strategy that, together with the opponent's strategy, forms an admissible strategy profile. Part (iii) formalizes our refinement. We do not explicitly require that deviations to non-Markov strategies should not be profitable. This requirement is in fact hard to formulate since we do not define a game that allows for non-Markov strategies. However, given the opponent's strategy, each player faces a Markov decision problem. Therefore, if there is a policy in this decision problem that yields a higher payoff than the candidate equilibrium strategy, then there is also a profitable deviation that is Markov.

Appendix C Uniqueness of SMPE for given p^*

We prove that for each fixed $p^* \in (\hat{p}, 1)$, the equilibrium in Propositions 2 and 3 is the unique SMPE in each case, provided that c is sufficiently small (i.e., $c \leq \min\{c_1, \dots, c_5\}$). We first characterize the sender's best response given any lower bound p_* of the waiting region and then show that it can be part of an equilibrium if and only if it is as specified in Propositions 2 and 3, respectively.

In the following we use p_* to denote an exogenously given lower bound of W , which may be different from the lower bound ϕ_{LR} in Proposition 2 or π_{LL} in Proposition 3. We will also use $V(p)$ and $U(p)$ to denote generic value functions for the sender and receiver

that are obtained from the sender's best response to a given waiting region with bounds p_* and p^* . Hence in the following $V(p)$ and $U(p)$ may be different from the functions defined in Sections A.1.1 and A.1.2. Throughout we assume that $c \leq \min\{c_1, \dots, c_5\}$.

A necessary condition. One crucial observation is that at p_* , either $V(p_*) = 0$ or $U(p_*) = U_\ell(p_*)$ —that is, at least one player should not expect a strictly positive net expected payoff from continuing. Toward a contradiction, suppose that $V(p_*) > 0$ and $U(p_*) > U_\ell(p_*)$.⁶¹ In this case, if p is just below p_* then, as when $p \in (\pi_0, \phi_{\ell R})$ in Proposition 2, the sender's flow payoff is maximized by her playing the L -drifting experiment with jump target p_* . But then, since $U(p_*) > U_\ell(p_*)$, the sender has no incentive to stop at p , contradicting that $p < p_*$ is not in the waiting region.

Now we proceed by characterizing the sender's value $V(p)$ and the receiver's value $U(p)$, if the sender plays a best response. In particular we characterize $V(p_*)$ and $U(p_*)$ which will enable us to use the necessary condition to narrow down possible equilibrium values of p_* .

The sender's best response in the waiting given (any) p_* . Let $\bar{V}(\cdot)$ and $\bar{U}(\cdot)$ denote the value functions given in Section A.1.2; that is, $\bar{V}(\cdot)$ and $\bar{U}(\cdot)$ represent the equilibrium value functions for Proposition 3. They play an important role in the subsequent analysis, because they coincide with the players' payoffs in a hypothetical situation where given p^* , the sender chooses both her dynamic strategy and the boundary of the waiting region p_* , ignoring the receiver's incentives.⁶² This implies that in any SMPE with fixed upper bound p^* , the sender's payoff can never exceed $\bar{V}(p)$. The following result is then immediate.

Lemma 11. *Fix p^* . Then in any SMPE $p_* \geq \pi_{\ell L}$.*

⁶¹Note that in equilibrium, $V(p)$ must be right-continuous at p_* , i.e., $V(p_*) = V(p_{*+}) = \lim_{p \downarrow p_*} V(p)$. If $V(p_{*+}) = 0$ this is obvious. Next suppose $V(p_{*+}) > 0$. For $p \in W$, the value of any strategy of the sender is continuous in p and hence $V(p)$ must be continuous in the waiting region. Therefore, a lack of right continuity at p_* can only arise if $p_* \notin W$, so that $V(p_*) = 0$. But if $V(p_{*+}) > 0$ and $V(p_*) = 0$, then for $p < p_*$ close to p_* , there is no strategy for the sender that satisfies (Ref), hence such a discontinuity cannot arise in equilibrium. We thus conclude that either $V(p_*) = V(p_{*+}) = 0$, or $V(p_*) = V(p_{*+}) > 0$ and in the latter case $p_* \in W$. Noting right-continuity at p_* makes the necessary condition much stronger and is key to the arguments below.

⁶²If the sender controls both the information structure, and the stopping decision, her value function is given by the viscosity solution of:

$$\max \left\{ \max_{(\alpha_i, q_i)_{i \in \mathbb{N}} \in \mathcal{I}} \sum_{q_i \neq p} \alpha_i (v(p; q_i) - c), -V(p) \right\} = 0. \quad (13)$$

It is easily verified that $\bar{V}(\cdot)$ satisfies this condition. For $p > \pi_{\ell L}$, $\bar{V}(\cdot)$ is a viscosity solution of (HJB) and $\bar{V}(p) > 0$ so that (13) holds. For $p \leq \pi_{\ell L}$, $\bar{V}(p) = 0$ and (Ref) holds with $\alpha_i = 0$ for all $i \in \mathbb{N}$, so that the first term in the outer max $\{\dots\}$ in (13) is equal to zero.

Proof. Suppose instead that $p_* < \pi_{\ell L}$. Then the sender's value is $V(p) = \bar{V}(p) = 0$ for $p \in (p_*, \pi_{\ell L})$ which is achieved by passing. Any other strategy leads to a negative value, so passing is the unique best response. Given this, the receiver's best response is to take action ℓ for $p \in (p_*, \pi_{\ell L})$, contradicting the hypothesis that the infimum of the waiting region is $p_* < \pi_{\ell L}$. \square

Next we characterize the sender's best response if $p_* > \pi_{\ell L}$. We begin with cases where it coincides with the strategy prescribed in Proposition 3

Lemma 12. *Fix p^* and suppose that p_* lies in the region where the sender either plays the stationary strategy (ξ for $p^* > \eta$), or plays the R -drifting experiment in Proposition 3. Then, the sender's value over W of her best response to p_* is given by $\bar{V}(p)$ and the receiver's value is given by $\bar{U}(p)$.*

Proof. If the prior is $p_0 \in (p_*, p^*]$ and the sender mimics her equilibrium strategy from Proposition 3, then the receiver will stop at the same time as if $p_* = \pi_{\ell L}$. Therefore the sender's payoff is equal to $\bar{V}(p_0)$ which is an upper bound for her optimal payoff. Hence the strategy remains a best response. To show that the receiver's value is given by $\bar{U}(p_*)$ we must also characterize the sender's best response (and not just her value). For $p \notin \{\xi, \underline{\pi}_{LR}, \bar{\pi}_L\}$ the Unimprovability Lemma 9 implies that the sender has a unique best response in Proposition 3, since by Lemmas 2 and 3.(c), $\bar{V}(p) > V_S(p)$ for $p \neq \xi$, hence we get uniqueness also if $p_* \notin \{\xi, \underline{\pi}_{LR}, \bar{\pi}_L\}$. For $p_* = \xi$ uniqueness of the sender's best response follows since choosing $\alpha > 1/2$ at p_* yields a value of zero for the sender, and $\alpha < 1/2$ violates admissibility since the sender uses the L -drifting experiment for $p > \xi$. If $p_* \in \{\underline{\pi}_{LR}, \bar{\pi}_{LR}\}$, non-uniqueness in Proposition 3 arises because the sender is indifferent between the $L0$ and RS strategies (or LS and R at $\bar{\pi}_{LR}$). This is no longer the case if $p_* \in \{\underline{\pi}_{LR}, \bar{\pi}_{LR}\}$ since the $L0$ or LS strategies are no longer feasible, so that uniqueness obtains. Since the sender has a unique best response given by the strategy from Proposition 3, the receiver's value from the sender's best response is given by $\bar{U}(p_*)$. \square

This Lemma immediately allows us to apply the necessary condition. Recall that $\bar{V}(p) > 0$ and $\bar{U}(p) > \mathcal{U}(p)$ for all $p \in (\pi_{\ell L}, p^*]$. Hence if p_* is in the region where the sender either plays the stationary strategy (ξ for $p^* > \eta$), or plays the R -drifting experiment in Proposition 3, $\bar{V}(p_*) > 0$ and $\bar{U}(p_*) > U_\ell(p_*)$.⁶³ Therefore p_* in this region leads to a violation of the necessary condition.

Next we consider the case where p_* lies in the region where the sender plays an L -drifting experiment in Proposition 3. This is the case if $p_* \in (\pi_{\ell L}, \underline{\pi}_{LR})$; and when $p^* \geq \eta$ also if $p_* \in (\xi, \bar{\pi}_{LR})$. In this case, the sender cannot simply replicate her strategy in Proposition 3. For example, suppose $p_* \in (\pi_{\ell L}, \underline{\pi}_{LR})$ and $p_0 \in (p_*, \underline{\pi}_{LR})$. If the sender used

⁶³As argued in Footnote 61, we must have $p_* \in W$ since $\bar{V}(p_{*+}) > 0$, and hence $\bar{V}(p_*) > 0$ and $\bar{U}(p_*) > U_\ell(p_*)$.

the $L0$ strategy, the receiver would stop when the belief drifts to p_* while in Proposition 3 he would wait until the belief reaches $\pi_{\ell L}$. Therefore, the sender's best response may be different from the strategy in Proposition 3.

The following lemma reports a set of observations about the best response that will allow us to use the necessary condition for an SMPE. To state this precisely, let $\pi_{\ell R}$ be the unique value such that $\tilde{V}(\pi_{\ell R}) = 0$.⁶⁴

Lemma 13. *Let $V(p)$ and $U(p)$ denote the players' expected payoffs from the sender's best response to a fixed lower bound p_* of the waiting region. Then there exists $c_6 > 0$ such that for all $c \leq c_6$:*

- (a) *If $p^* \geq \eta$ and $p_* \in (\xi, \hat{p})$, then $V(p_*) \geq V_S(p_*) > 0$ and $U(p_*) \geq U_S(p_*) > U_\ell(p_*)$.*
- (b) *If $p_* \in (\pi_{\ell L}, \pi_{\ell R})$ then $V(p) = V_{L0}(p; p_*)$ and $U(p) = U_{L0}(p; p_*)$ for all $p \in [p_*, \pi_{\ell R})$.*
- (c) *If $p_* \in [\pi_{\ell R}, \underline{\pi}_{LR})$ then $V(p_*) = \tilde{V}(p_*)$ and $U(p_*) = \tilde{U}(p_*)$.*

Proof. If $p_* > \xi$, let $c_6 > 0$ be chosen such that for all $c \leq c_6$, $V_S(p) > 0$ for all $p \in [p_*, p^*)$, and $U_S(p_*) > \max\{U_\ell(p_*), U_r(p_*)\}$.

The sender's best response can be constructed in a similar way as in Section A.1, taking the lower bound p_* as a constraint. If $p_* \in (\pi_{\ell L}, \xi)$, the $L0$ strategy now uses $p_* \neq \pi_{\ell L}$ as a stopping bound and the value is given by $V_{L0}(p; p_*)$ as before. No further modifications are needed in this case. If $p_* > \xi$, we replace the LS strategy by a modified version which we denote LS_* . According to this strategy, the sender uses the L -drifting experiment for $p > p_*$ and switches to the stationary strategy when the belief drifts to p_* . The value of this strategy is given by $V_{LS}(p; p_*)$. We also set $\hat{V}_*(p) = V_{LS}(p; p_*)$ if $p_* > \xi$. With this we can then characterize the value of the sender's best response as $\tilde{V}_*(p) = \max\{V_R(p), \hat{V}_*(p)\} = \max\{V_R(p), V_{LS}(p; p_*)\}$. Verification of the sender's best response proceeds using similar steps to those in Section A.2.

(a) The value of the sender's best response in this case is given by $V(p) = \max\{V_{LS}(p; p_*), V_R(p)\}$. If $V_R(p_*) > V_S(p_*)$ this implies $V(p) = V_R(p)$ since $V_R(p) > V_S(p)$ by Lemma 1.(b) and therefore $V_{LS}(p; p_*)$ cannot cross $V_S(p)$ from below by Lemma 1.(c). With $V(p) = V_R(p) > V_S(p)$, the Unimprovability Lemma 9 implies that there is a unique best response, the R -drifting strategy, and therefore $U(p_*) = U_R(p_*) > U_S(p_*) > U_\ell(p_*)$

If $V_R(p_*) \leq V_S(p_*)$, we must have $p_* < \xi_2$ since $V_R(p) > V_S(p)$ for $p \in [\xi_2, p^*)$ by Lemma 1.(b). Also by Lemma 1.(b), $V_{LS}(p; p_*) > V_S(p)$ for $p \in (p_*, \xi_2]$. Hence $V(p) > V_S(p)$ for all $p \in (p_*, p^*)$. By the Unimprovability Lemma 9 this implies that there is a unique best response for all $p \in [p_*, p^*]$ except at the belief where $V_{LS}(p; p_*)$ and $V_R(p)$ intersect (which could occur at p_*), and in the latter case $U_{LS}(p; p_*) = U_R(p)$ since the sender and the receiver incur the same cost. Hence $U(p_*) = U_{LS}(p_*; p_*) = U_S(p_*) > U_\ell(p_*)$,

(b) In this case, $V(p) = V_{L0}(p; p_*) > 0 > \tilde{V}(p) > V_S(p)$ for $p \in (\pi_{\ell L}, \pi_{\ell R})$. Therefore, the Unimprovability Lemma 9 implies that there is a unique best response for

⁶⁴From our previous results, we know (i) $\pi_{\ell R} \in (\pi_{\ell L}, \underline{\pi}_{LR})$ (Lemma 6.(b)); (ii) if (C2) holds, then $\pi_{\ell R} < \phi_{\ell R}$ (Lemma 4) and $\pi_{\ell R} < \pi_0$ (Lemma 3.(a)); and (iii) if (C2) fails, then $\pi_{\ell R} \geq \phi_{\ell R}$ (Lemma 4).

$p \in (\pi_{\ell L}, \pi_{\ell R})$ and the sender uses the $L0$ strategy with stopping bound p_* . Therefore we have $U(p) = V_{L0}(p; p_*)$ for all $p \in [p_*, \pi_{\ell R})$.

(c) In this case $V(p) = \tilde{V}(p) > V_S(p)$ for $p \in [p_*, \xi)$,⁶⁵ and the Unimprovability Lemma 9 implies that there is a unique best response: the sender uses the RS -strategy if $p^* > \eta$ and the R -drifting strategy if $p^* < \eta$.⁶⁶ Therefore we have $U(p_*) = \tilde{U}(p_*)$ \square

By Lemmas 11, 12, and 13.(a), we must have $\pi_{\ell L} \leq p_* < \underline{\pi}_{LR}$ if $c < \min\{c_1, \dots, c_6\}$. Now we further narrow down possible equilibrium values of p_* and show that in Proposition 2 we must have $p_* = \phi_{\ell R}$, and in Proposition 3 we must have $p_* = \pi_{\ell L}$. With that it only remains to show uniqueness of the sender's equilibrium strategy which follows from the Unimprovability Lemma.

Proposition 2. Proposition 2 concerns the case where (C2) holds. First we rule out $p_* \in (\pi_{\ell R}, \underline{\pi}_{LR}) \setminus \{\phi_{\ell R}\}$. If $p_* \in (\pi_{\ell R}, \underline{\pi}_{LR})$, $V(p_*) = \tilde{V}(p_*)$ and $U(p_*) = \tilde{U}(p_*)$ by Lemma 13.(c). Since $p_* > \pi_{\ell R}$, $V(p_*) = \tilde{V}(p_*) > 0$ and hence the necessary condition for an SMPE implies that $U(p_*) = \tilde{U}(p_*) = U_{\ell}(p_*)$. But this condition only holds when $p_* = \phi_{\ell R}$, in which case the equilibrium is as specified in the Proposition 2.

Next, we rule out $p_* \in (\pi_{\ell L}, \pi_{\ell R})$. We begin by showing that $p_* < \pi_{\ell R}$ implies $p_* < \phi_{\ell L}$ if $c \leq c_5$. To see this, recall that by Lemma 8, $\pi_0 < \phi_{\ell L}$ if $c \leq c_5$. Since $\tilde{V}(p)$ is convex, the construction of π_0 therefore implies $\tilde{V}(\phi_{\ell L}) > 0$. On the other hand $\tilde{V}(p_*) < 0$ if $p_* < \pi_{\ell R}$. Therefore $p_* < \pi_{\ell R}$ implies $p_* < \phi_{\ell L}$.

Now we proceed given that $p_* < \phi_{\ell L}$: By Lemma 13.(b), $p_* \in (\pi_{\ell L}, \pi_{\ell R})$ implies that $U(p) = U_{L0}(p; p_*)$ for all $p \in [p_*, \pi_{\ell R})$. Therefore, $U(p) = U_{L0}(p; p_*) < \max\{U_{\ell}(p), U_T(p)\}$ for $p \in (p_*, \phi_{\ell L})$ which means that it is not optimal for the receiver to wait for beliefs $p \in (p_*, \phi_{\ell L})$ in the waiting region. Therefore $p_* \in (\pi_{\ell L}, \phi_{\ell L})$ cannot arise in equilibrium.

Proposition 3. Suppose (C2) fails. We first rule out $p_* \in (\pi_{\ell R}, \underline{\pi}_{LR})$. If $p_* \in (\pi_{\ell R}, \underline{\pi}_{LR})$, then Lemma 13.(c) implies $V(p_*) = \tilde{V}(p_*) > 0$ and with (C2) failing, this implies $U(p_*) > U_{\ell}(p_*)$. Hence, $p_* \in (\pi_{\ell R}, \underline{\pi}_{LR})$ leads to a violation of the above necessary condition.

Next we rule out $p_* \in (\pi_{\ell L}, \pi_{\ell R}]$. If $p_* \in (\pi_{\ell L}, \pi_{\ell R}]$, then by Lemma 13.(b) and (c), $V(p_*) = 0$. Hence, for any $p \in (\pi_{\ell L}, p_*)$, by the refinement, the sender uses the L -drifting experiment with jumps to p^* . Since (C2) fails, by the same argument as in the analysis of the waiting region in Section A.3, the receiver prefers to wait. This contradicts $p < p_*$.

We have ruled out all values for p_* except $p_* = \pi_{\ell L}$. Hence, the equilibrium specified in Proposition 3 is unique (up to tie breaking at $\underline{\pi}_{LR}$ and $\bar{\pi}_{LR}$, and if $p^* = \eta$). Uniqueness follows from the Unimprovability Lemma 9 since $V(p) > V_S(p)$ for all $p \in [p_*, p^*)$.

⁶⁵To see this note that $\tilde{V}(p_*) \geq V_{L0}(p_*; p_*) = 0$ and $\tilde{V}(p) > V_S(p)$ for $p \in [p_*, \xi)$. Hence by the Crossing Lemma 1.(d), $V_{L0}(p; p_*) < \tilde{V}(p)$ for all $p \in (p_*, \xi)$.

⁶⁶In the case $p^* = \eta$, both the sender and the receiver are indifferent between the RS -strategy and the R -drifting strategy.