Cheap Talk in Complex Environments[*]

Yunus C. Aybas[†]   Steven Callander[‡]

February 29, 2024

Abstract

Decision making in practice is often difficult, with many actions to choose from and much that is unknown. Experts play a particularly important role in such complex environments. We study the strategic provision of expert advice in a variation of the classic sender-receiver game in which the environment is complex. We identify an equilibrium that is efficient and sender-optimal. The outcome is exactly what the sender would choose if she held full decision making authority. This contrasts with the simple environment of Crawford and Sobel (1982) in which equilibrium outcomes are inefficient and favor the receiver. The equilibrium we identify satisfies the neologism-proof and announcement-proof refinements, and all equilibria satisfying the latter requirement are outcome equivalent to our equilibrium.

---

[*] For helpful comments, we thank Anton Kolotilin, Johannes Schneider, Juan Ortner, Avi Acharya, Dave Baron, Dana Foarta, Spencer Pantoja, Takuo Sugaya, Ilya Segal, Ravi Jagadeesan, Matt Jackson, Bob Wilson, Ian Ball, Yuliy Sannikov, Arjada Bardhi, Can Urgun, Weijie Zhong, Mitchell Watt, Emir Kamenica, Marina Halac, Ben Brooks, Daniel Rappoport, Doron Ravid, Eran Shmaya, Archishman Chakraborty, Emiel Awad, Gilat Levy, Ricardo Alonso, Leeat Yariv, Nicolas Lambert, João Ramos, Wouter Dessein, Joel Sobel, seminar audiences, and especially Mike Harrison. We also thank Christoph Carnehl and Paul Milgrom for their close readings of earlier drafts.

[†] Stanford University, Department of Economics; aybas@stanford.edu.

[‡] Stanford University, Stanford Graduate School of Business; sjc@stanford.edu.

1 Introduction

Expert advice is vital to decision makers in many aspects of economic, political, and social life. Expertise not only improves the quality of decision making, it also delivers power to those who hold it. In domains ranging from medical care to real estate, from car repair to business investment, experts are able to steer decisions toward their own interests even when those conflict with the interests of decision makers. Weber (1958, p. 232) went so far as to say that the "power position" of an expert is always "overtowering" and that the decision maker "finds himself in the position of the 'dilettante' opposite the 'expert'."[1]

Models of expert-guided decisions specify both the extra information that the expert knows and the differences between the parties' respective interests. In past models, these differences are simple in the sense that the expert knows only a single thing that the decision maker does not, and this advantage is such that if the decision maker knew the expert's optimal choice, she could infer her own optimal choice (Crawford and Sobel, 1982; Milgrom, 1981). This parsimonious approach has yielded useful insights into the structure of expert advice and provided the foundation for innumerable studies of expertise in markets and institutions.

In simplifying the decision problem to such a degree, however, these models generate several stark properties that do not resonate with how expertise operates in the real world. In practice, a doctor not only knows much more than her patient, but a single diagnosis, no matter how precise, does not reveal to a patient what his most preferred treatment would be. Moreover, in the classic model of Crawford and Sobel (1982), the expert's minimal informational advantage leads her to communicate imprecisely, and the information that she does convey promotes a decision that is optimal for the decision maker and not the expert herself. Rather than "overtowering" the decision maker, the expert would be better off if she could simply transfer her information to the decision maker for free.

In this paper we develop a model of expertise in complex environments. In a

---

[1] See Milgrom and Roberts (1988) for division managers manipulating headquarters into funding too many projects, Levitt and Syverson (2008) for realtors manipulating homeowners into selling too quickly and cheaply, and Gruber and Owings (1996) for OBGYNs manipulating patients into having too many C-sections.

complex environment the expert knows many things that the decision maker does not. Moreover, the expert's knowledge is imperfectly invertible. Thus, should the decision maker learn that an action produces the expert's ideal (or any) outcome, he cannot infer perfectly the outcome of other actions.

We show that in complex environments an expert can hold power over a decision maker. We revisit the classic sender-receiver game of Crawford and Sobel (1982) for complex environments and identify an equilibrium that is expert-optimal and efficient. In fact, the outcome is exactly what the expert would obtain were she to hold full decision making authority herself.

A complex environment empowers the expert (sender) as it allows her to communicate precisely yet imperfectly. She can recommend her most-preferred action without revealing all of her information to the decision maker (receiver). In a complex environment the *informational spillover* from her recommendation is incomplete. In contrast, in a simple environment the informational spillover is complete. Precise communication about one action is precise communication about all actions. In a complex environment, the sender can use her information while keeping some of it private. We show that this ability is the source of expert power.

Expert power in equilibrium is, however, not maximized when the informational spillover is minimized. A key part of our result is to identify a sender strategy and environments that allow the sender to not just reduce the informational spillover, but to shape it in such a way that it systematically favors the recommendation itself. By shaping the spillover, the sender is able to dissuade the receiver from other actions, and by so doing, more effectively persuade him to accept the recommendation that the receiver knows is the sender's optimal. We show that when the informational spillover is minimized the efficient sender-optimal equilibrium exists only if the sender's bias is small, but when the spillover is larger and shaped in the sender's favor, the equilibrium exists for small and large bias.

An expert's *power* is distinct from her *influence*. An expert is influential when her information shapes the decision maker's decision (Sobel, 2010). The more finely her advice communicates the state, the more influence she has, and the more efficient is communication. Expert power is the degree to which outcomes favor the expert over the decision maker. In the equilibrium we identify, the expert is influential *and* powerful. In the simple environment of Crawford and Sobel (1982), the expert is influential but not powerful, and less influential than she is in complex environments.

These pairings are not coincidental. When the decision maker holds power, the expert has an incentive to deceive him. Crawford and Sobel (1982) show how imprecise communication is necessary to remove that temptation. In a simple environment the expert must sacrifice influence to gain credibility. In a complex environment there need not be such a trade-off. When the expert holds power she has no need to deceive and, therefore, she need not sacrifice her influence. We show that in equilibrium the expert can simultaneously hold maximum power and maximum influence. Inefficient communication is, therefore, not an inherent property of cheap talk with a biased expert. Rather, it is a function of which player holds power.

A complex environment enables an expert to hold power and influence, but does not guarantee it. We prove that, even in complex environments, if the decision maker holds any power in equilibrium, the expert's influence must be reduced. In simple environments the decision maker holds power in all equilibria and this bounds the expert's influence. In complex environments the expert can hold all the power in equilibrium and, when she does, she is able to communicate efficiently with maximum influence.

Expert power and influence are intimately related to equilibrium refinement notions in cheap talk games. The "neologism proofness" concept of Farrell (1993) requires that natural meanings of words are available when players share a common language. For example, a sender can say "I am a high type" and the receiver will understand the meaning of that statement. As intuitive as neologism-proofness is, it is a demanding requirement and, notably, it eliminates *all* equilibria in the environment of Crawford and Sobel (1982). The equilibrium we identify satisfies neologism-proofness, as well as the related concept of "announcement proofness" due to Matthews et al. (1991). In fact, we show that the *only* equilibria that satisfy announcement-proofness are outcome-equivalent to the equilibrium we identify.[2]

Our equilibrium satisfies these refinements precisely because the expert holds

---

[2] The richness of the state space—a continuum of random variables—is such that it is difficult to prove that *no* other neologism-proof equilibrium exists, although it is unclear what form it would take if one did exist.

power *and* influence. The existence of a natural language means that the sender's ability to communicate cannot be limited by the equilibrium itself. When the expert lacks power and influence, she has an incentive to access additional meaning and the natural language undermines equilibrium. It is exactly when the expert has power and influence that she has no need to access additional meaning and an equilibrium is neologism- and announcement-proof. This fact motivates our focus on an expert-optimal equilibrium.

Complex environments come in many different forms, varying in the size and nature of the informational gap between the players and even the set of actions available. We provide examples to illustrate the possibilities. To develop the ideas more deeply, we focus for much of the paper on a particular representation of complex environments in which the possible mappings from actions to outcomes are paths of Brownian motion. The expert knows the path while the decision maker does not.

The methodological advantage of the Brownian motion is that it allows us to tune the degree of informational spillover from the expert's recommendation. The scale of the Brownian motion parameterizes the correlation of payoffs for different pairs of actions and, thus, how much the decision maker can infer about other actions when the expert recommends his most-preferred action. By varying the scale relative to the drift of the Brownian motion, we can dial up and dial down the degree of informational spillover within a common framework. In this way, we can characterize when the expert-optimal equilibrium exists and, as importantly, when it does not. In the canonical model the decision maker learns too much from a recommendation for the expert to retain power. Using the Brownian motion, we show how much learning is too much for the expert to hold power in equilibrium.

The behavior in our equilibrium resonates with the many situations in practice in which a decision maker acquiesces to an expert's recommendation. It explains why a board of directors accepts unchanged a CEO's recommended strategy or a patient adopts a doctor's recommended treatment. Although such behavior may appear as deference, or even the rubber-stamping of a recommendation with little thought, our equilibrium shows why a rational decision maker acquiesces to an expert even when he is keenly aware that the expert is biased and knows that the recommendation serves the expert's interests and not his own. In his famous study of health care, (Arrow, 1963, p. 965) observed that as a “consequence of information inequality between physician and patient, ... the patient must delegate to the physician much of his freedom of choice.” Our model provides a formal basis to this observation. We show that it is the complexity of the environment that allows the doctor to “overtower” the patient in the sense of Weber (1958) and that the doctor can attain this position purely through the power of her advice.

Relationship to the Literature. We build on the seminal contribution of Crawford and Sobel (1982), hereafter CS, expanding their model to complex environments. It is notable that our results do not contradict their conclusion that “perfect communication is not to be expected in general unless agents’ interests completely coincide, ...” (p. 1450)[3] Our contribution is to observe that in complex environments a gap emerges between efficient and perfect communication and that the former does not require the latter. We show how in complex environments this gap can be leveraged by the expert, leading to favorable outcomes even when her interests do not coincide with those of the decision maker.

Our approach is most closely related to models in which the decision maker is unsure of the expert’s bias. Morgan and Stocken (2003) assume that with some probability the interests of the expert and decision maker are aligned. Later work extends this to uncertainty over the direction as well as the magnitude of the expert’s bias and allows bias to depend on the state of the world (Li and Madarász, 2008; Gordon, 2010). As the expert now knows two things the decision maker does not—her bias and the single-dimensional state—these models are equivalent to a minimal complexification of the canonical simple environment where the knowledge gap between the players takes a particular structure. This empowers the expert but only to a limited degree, and equilibria remain of the partitional form and are inefficient. In Section 5 we construct examples of complex environments in which the expert’s advantage is two pieces of information and show that efficient communication is possible but can be fragile. A central message of our paper is that a much larger informational advantage for the expert arises naturally when the environment is complex and that this advantage can robustly support efficient cheap talk.

An expert can also be empowered by institutional structure and commitment

---

[3] Sobel (2010, 2012) also emphasize the existence of fully revealing equilibria.

power. Gilligan and Krehbiel (1987) demonstrate how institutions, by constraining the receiver's action space, empower the sender and incentivize the acquisition of expertise.[4] Kamenica and Gentzkow (2011) show that if the sender can commit to an information structure then she can persuade a decision maker even when she is not an expert.

We follow CS in modeling players with state-dependent preferences. A separate line of work supposes the sender's preferences are state-independent. The notion of power is less clear in this setting as the players desire different things. (Kamenica and Gentzkow (2011) have popularized the use of *persuasion* to describe a sender convincing a receiver to take an action he otherwise would not.) Chakraborty and Harbaugh (2010) and Lipnowski and Ravid (2020) show how an expert can be influential if her preferences are quasiconvex in the receiver's action and, in some situations, that the expert can be better off and the receiver worse off than they would be with full information revelation.[5]

Embedded with Aghion and Tirole's (1997) famous model is a cheap talk game with two states and discrete actions. By assuming the quality of projects are drawn independently, they rule out informational spillover, in what can be thought of as an infinitely complex environment. Using the Brownian motion, we capture the full range of complexity in an environment. We show how expert power can emerge in the presence of informational spillovers, even when the action space is continuous and the receiver can adjust a recommendation as finely as he wishes. We quantify how much spillover is too much for an expert to hold power. Moreover, we show how the expert's strategy can shape the spillover in a way that favors the recommendation, delivering power to the expert even when her bias is large and in situations where she wouldn't have power in the absence of spillovers.

The Brownian motion has been used to represent the action-outcome mapping

---

[4] More generally, a discrete action space can empower the sender. See Chakraborty and Harbaugh (2007) for an example of binary-action cheap talk in which the sender obtains her ideal action.

[5] Equilibria are nevertheless inefficient for an open set of receiver priors unless the sender and receiver preferences are aligned such that both prefer full information revelation.

in a variety of applications.[6] Callander et al. (2021) analyze a model of verifiable information and show how the expert can obtain leverage by providing information in addition to a recommendation. The verifiability of information is essential to their result. Callander (2008) studies a model of cheap talk and identifies the efficient equilibrium when the informational spillover is minimized and the equilibrium exists only for small bias. Moreover, in that case, the equilibrium relies on the risk aversion of the decision maker. We show how the efficient equilibrium can exist for large bias, how it relies on the expert shaping the informational spillover in a particular and beneficial way, and that it need not rely on risk aversion of the decision maker.

2 The Model

We consider the classic sender-receiver game of Crawford and Sobel (1982) extended to complex environments. For clarity, we present the results for the workhorse domain of constant bias and quadratic utility.

Timing: An expert (sender) sends a message, $r \in \mathcal{M}$, to the decision maker (receiver), who chooses an action $a \in \mathcal{A}$ that affects the utility of both players.

The Environment: The set of available actions is an interval, $\mathcal{A} = [0, q]$, for $q \in \mathbb{R}_+ \cup \infty$. Each action produces an outcome given by the mapping, $\psi : \mathcal{A} \to \mathbb{R}$. The status quo is action 0 with outcome $\psi(0) > 0$. The mapping is given by the realized path of a Brownian motion with drift $\mu < 0$ and scale σ, and that passes through the status quo point. One possible path is depicted below in Figure 1. The state is the realized path and the state space is the set of all such paths, which we denote by Ψ.[7] The message space, $\mathcal{M}$, is arbitrary and large.

Information: The sender knows the realized path $\psi(\cdot)$. The receiver knows only the drift and scale parameters, and the status quo point (and that $\psi(\cdot)$ is generated as a Brownian motion over $\mathcal{A}$).

Preferences: Utility functions for the sender and receiver are denoted, respectively,

---

[6] Applications include search and experimentation (Callander, 2011; Garfagnini and Strulovici, 2016; Urgun and Yariv, 2021a,b; Cetemen et al., 2023) and "attributes"problems (Bardhi, 2022; Bardhi and Bobkova, 2023; Carnehl and Schneider, 2021).

[7] Formally, the state space is the set of all continuous functions with domain $\mathcal{A}$ and range $\mathbb{R}$.

by: $u^S, u^R : \mathcal{A} \times \Psi \to \mathbb{R}$. Throughout the paper, we focus on the particular form: $u^R(a,\psi) = -\psi(a)^2$ and $u^S(a,\psi) = -(\psi(a) - b)^2$, where $b > 0$ is the sender's bias. Our main results extend to receiver utility functions that exhibit weak concavity in outcomes, with a unique maximum at outcome 0, and sender utility functions that are maximized at outcome b.[8] We assume that the sender's preferred outcome is better than the status quo for the receiver, $b < \psi(0)$, and address the case of larger bias separately after our main result.

Strategies and Equilibrium: Strategies for the sender and receiver are maps, $m : \Psi \to \mathcal{M}$ and $a : \mathcal{M} \to \mathcal{A}$, respectively. The receiver updates his beliefs via Bayes' rule on the equilibrium path conditional on the realization of the message $m(\psi)$. We provide the formal description of these beliefs in the appendix. We say, informally, that the expert *recommends* action a if by sending message r the expert intends that the receiver choose action a. For simplicity, we refer to a recommendation r and the action it recommends interchangeably. Hereafter, *equilibrium* refers to a Perfect Bayesian Equilibrium.

Remark 1. In CS actions map directly to utility. In most applications actions map to an outcome, from which agents draw utility. Formalizing this intermediate step allows a clearer view of the decision making environment. Viewed through this lens, the state space in CS in the fixed bias case is equivalent to a mapping $\hat{\psi}(a) = \theta - a$, where $\theta \in [0,1]$ is the expert's single piece of private information. This reflects two differences with our setting. CS is equivalent to setting $\sigma = 0$ but allowing for the status quo outcome to be uncertain. In CS a known status quo outcome would fully reveal the mapping. In more complex environments, knowledge of a single point amongst a continuum of unknowns is less important and is immaterial when the outcome is further from the receiver's ideal than is the sender's bias.[9]

Remark 2. The Brownian motion has found application in a variety of settings as it provides a tractable and appealing representation of information rich environments.

---

[8] We point out the results that are special to the quadratic form where relevant.

[9] Our results extend to the case in which the support of uncertainty over the status quo outcome does not include the sender's ideal outcome. For convenience we assume $b < \psi(0)$ and that the status quo is known. Later, we show that our equilibrium exists only on a degenerate action space

One attractive property is that the mapping is partially invertible.[10] Learning the outcome of one action reveals some information about the outcomes of other actions but not everything. Moreover, the amount of information revealed is higher for actions that are nearby and lower for actions that are more distant. The degree of invertibility depends on the variance of the Brownian motion, given by σ^2, with higher variance meaning that less information spills over from a recommendation to other actions. As the cost of uncertainty due to σ^2 is scaled against the drift, we parameterize the *complexity* of the decision making environment by the ratio $\frac{\sigma^2}{|\mu|}$.

With the Brownian motion, the sender's advantage is a continuum of information and complexity is the correlation across that information. An alternative representation of complexity is by the number of discrete pieces of information a sender knows that a receiver does not.[11] In Section 5.2 we present several environments that extend CS in this way and which support efficient cheap talk.

Remark 3. We develop the analysis by varying the size of the action space rather than the parameters of the Brownian motion (although see the comparative statics in Section 4.4). The sender's power derives from the complexity of the environment but, as will become evident, her power is not in direct proportion to complexity. Varying the size of the action space, q, allows us to cleanly separate the effects the sender's strategy has on the receiver's beliefs in a way that varying the scale parameter, σ, doesn't.

Remark 4. For clarity of presentation, we focus on positive bias ($b > 0$), anchor the action space at 0, and impose quadratic utility. These assumptions are not essential to the underlying logic of our results, and we relax each later in the paper.

3 Decision Making Without an Expert

Suppose the expert is not present and the receiver is on his own. The receiver faces the choice of the certain outcome of the status quo or an uncertain outcome from any other action. His beliefs over outcomes follow from the properties of the Brownian

---

if $b \geq \psi(0)$. We conjecture that for an uncertain status quo outcome with b in its support, the equilibrium we identify exists as long as enough mass in the support is greater than b.

[10] See Callander (2011) for a more complete description of the properties of the Brownian motion.

[11] Such a representation implies that a decision maker becomes completely informed after observing

motion and are normally distributed for each action a with expected outcome and variance as follows:

$$\mathbb{E}\left[\psi\left(a\right)\right] = \psi\left(0\right) + \mu a,$$
$$\text{Var}\left(\psi\left(a\right)\right) = \sigma^2 a.$$

The expected outcome is determined by the drift line, which by assumption is negative. Figure 1 depicts the environment. Variance is increasing in the distance an action is from the status quo, capturing the idea that uncertainty is increasing the more distant an action is from what has been tried before. We say that beliefs of this form are *neutral*.

In evaluating actions, the receiver faces a trade-off between risk and return. The larger the action he chooses, the better the expected outcome, at least up to the point at which it crosses his ideal outcome at zero, but the greater is the variance. His optimal action depends on the ratio of variance to drift of the Brownian motion, thus, on the complexity of the decision making environment. The critical threshold is exactly half of this ratio, which we define by α such that $\alpha = \frac{\sigma^2}{2|\mu|}$.

Lemma 1 *In the absence of expertise, the receiver chooses* a^{no} *such that:*

(i) For $\psi\left(0\right) > \alpha$, $a^{no} = \frac{\psi(0)-\alpha}{|\mu|} > 0$ *and* $\mathbb{E}\psi\left(a^{no}\right) = \alpha$.

(ii) For $\psi\left(0\right) \in \left[0, \alpha\right]$, $a^{no} = 0$.

Lemma 1 reflects the reality that the alternative to advice is experimentation. If the status quo point is sufficiently unattractive, the receiver will forge out on his own and try something new in the hope that it delivers a better outcome. Quadratic utility delivers a particularly simple form to this choice.[12]

The threshold α represents the point at which the marginal benefit in expected outcome equals the marginal cost of greater risk. For a status quo less extreme than α, the risk of experimentation is not worth the return and the receiver accepts the certainty of the known outcome. For a status quo outcome beyond α, the risk is worth the return, and the receiver experiments to the point that the expected

---

a finite number of points in the mapping. An appealing property of the Brownian representation is that knowledge of the world remains incomplete after any (finite) number of observations.

[12] Our results extend to arbitrary weakly concave utility with a unique maximum, although the threshold in Lemma 1 is only constant for the quadratic case. We refer to α as a constant throughout the paper, though all statements hold for a generalized threshold. Quadratic utility

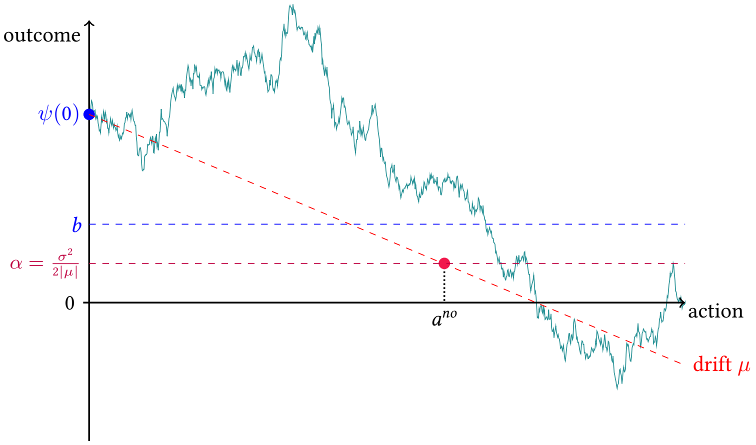

Figure 1: The Mapping From Actions to Outcomes.

outcome is exactly equal to α. Notably, the receiver could obtain his ideal outcome in expectation, though he chooses not to.

The receiver's optimal action in the absence of an expert is marked in Figure 1. His expected utility is strictly decreasing in $\psi(0)$. This is immediate for $\psi(0) \leq \alpha$ as his utility is simply $-\psi(0)^2$. For $\psi(0) > \alpha$ his expected utility takes the simple mean-variance form: $\mathbb{E}[u_R(a^{no})] = -\alpha^2 - \sigma^2 a^{no}$. As $\psi(0)$ increases so does a^{no}, and while the expected outcome remains constant at α, the variance increases in a^{no} and, thus, in $\psi(0)$.

4 Efficient Cheap Talk

In any sender-optimal equilibrium with full support, the sender recommends an action that is one of her most preferred. This action may not be unique. We study the *first-point* strategy in which she recommends the smallest of her most-preferred actions.

Definition 1 *In the first-point strategy the recommendation for each $\psi \in \Psi$ is:*

$$m^*(\psi) = \min\left\{a : |\psi(a) - b| \leq |\psi(a') - b| \ \textit{for all}\ a' \in [0, q]\right\}.$$

The first-point strategy requires that the sender recommend the smallest action that obtains outcome b whenever possible. If such an action does not exist, she recommends

---

matters at one other point; see footnote 18 on the comparative static of Proposition 1.

the action whose outcome gets as close to b as possible. We say an equilibrium is the *first-point equilibrium* if the sender uses the first-point strategy and the receiver follows the recommendation. We denote a generic realization of $m^*(\psi)$ by r^*.

The first-point equilibrium is Pareto efficient (ex post and ex ante) as it always delivers the sender's (weakly) most preferred action and no other action can make both players better off. Thus, it is clearly incentive compatible for the sender to follow the strategy (conditional on the receiver following her recommendation). The receiver's incentive to follow the recommendation is more subtle. The logic of his decision can be seen most clearly by varying the size of the action space. We begin with the case of an unbounded action space.

4.1 Unbounded Action Space

If $q = \infty$ and the action space is the entire real half-line, the negative drift of the Brownian motion implies that the path crosses b almost surely for at least one action. The receiver believes, therefore, that the recommendation from the sender using the first-point strategy delivers outcome b with probability one.

The information revealed is not limited to the recommendation and spills over to other actions as well. It is helpful to distinguish between what we refer to as direct and indirect informational spillover.[13] *Direct* informational spillover comes from knowledge that the mapping passes through the point (r^*, b). This shapes the receiver's beliefs about all other actions.

Indirect informational spillover is what the receiver infers from the fact that r^* was the recommendation and not some other action.[14] For an unbounded action space, the indirect informational spillover is contained in one region of the action space. Specifically, the receiver infers indirectly that actions to the left of r^* must produce outcomes above b—if they didn't, the recommendation would have been an action to the left of r^* instead. The receiver is able to infer indirectly, therefore, that actions to the left of the recommendation are strictly worse for him with certainty than the recommendation itself. Thus, if he is to override the recommendation, it must be

---

[13] In the simple environments of CS this distinction disappears.

[14] To see this distinction between the recommendation and the strategy, imagine the sender instead used a *last-point* strategy, revealing the largest action that produces outcome b. The direct spillover would be identical but the indirect spillover very different. See the discussion at Section 4.5.

with an action to the right.

To the right of the recommendation, however, there is no indirect informational spillover. Because the sender recommends the first point that crosses b, the receiver learns nothing about the mapping to the right beyond the direct spillover from the recommendation. To the right of r^* the receiver's beliefs remain neutral, albeit now anchored by the recommendation rather than the status quo.

This is important as neutral beliefs to the right of the recommendation mean that the logic of Lemma 1 applies. It follows that the receiver is willing to accept the recommendation, but only if the expected outcome is close enough to her ideal at zero. Only if, therefore, the expert's bias is not too large relative to the complexity of the environment.

Lemma 2 *If $q = \infty$, the first-point equilibrium exists if and only if $b \leq \alpha$.*

In equilibrium, the receiver knows that the expert is recommending her ideal outcome—and that it is different from his own—yet he is willing to accept because the risk of overriding the recommendation and experimenting on his own is not worth the return. The receiver knows that actions to the right deliver in expectation a better outcome, and with probability one that a better action exists, but he doesn't know with certainty which actions deliver a better outcome. He faces what we refer to as *response uncertainty.* For small enough bias, his response uncertainty is enough that he prefers the certainty of the sender's ideal action. The logic of this equilibrium does not depend on the quadratic utility form, though it does depend on risk aversion.

Lemma 2 can be stated equivalently in terms of the scale of the Brownian motion. The requirement that $b \leq \alpha = \frac{\sigma^2}{2|\mu|}$ is equivalent to a requirement that $\sigma^2 \geq 2b|\mu|$. This equivalence is not general. It holds here for two reasons. First, for any σ, the receiver draws the same inference from the recommendation as it depends on b, and second, there is no indirect informational spillover to the right of the recommendation. The receiver's beliefs are neutral, therefore, and an increase in σ maps directly to an increase in the receiver's residual uncertainty. Thus, if the receiver is willing to accept a recommendation for some σ, he is willing to accept the recommendation for larger σ.

That the receiver's beliefs are neutral means the recommendation does not dissuade the receiver from taking an action to the right. In fact, these actions are more attractive to the receiver than they were initially due to the direct informa-

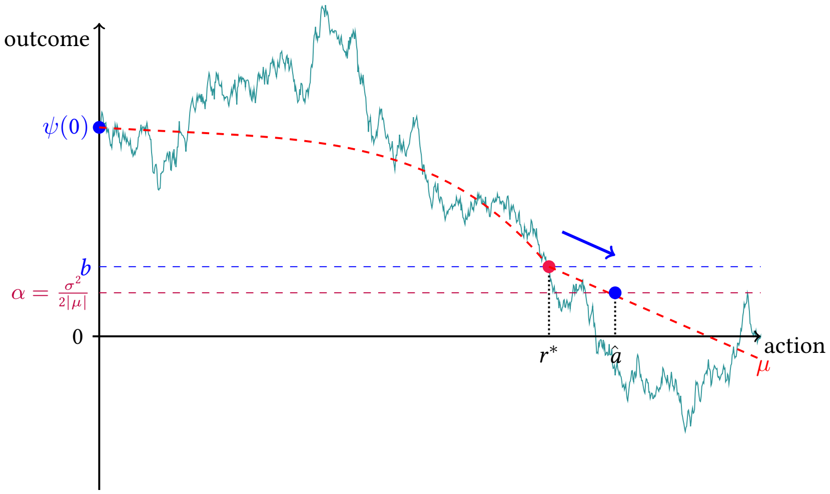

Figure 2: First-Point Strategy: Recommendation r^* & Optimal Receiver Response $\hat{a}$.

tional spillover from the recommendation. The equilibrium in Lemma 2 works purely through persuasion. This decreases the sender's power and only works because the receiver is risk averse. The sender convinces the receiver that the certainty of the recommendation is better than the risky alternatives and, even then, this is only possible when her bias is small.

For larger bias the risk of overriding the recommendation is worth the return and the equilibrium fails. Figure 2 depicts the situation in which the receiver overrides recommendation r^* with action $\hat{a}$. The equilibrium fails even if the recommendation delivers higher utility than the receiver would obtain without the expert. This is because the receiver would use the information contained in the recommendation to obtain an expected outcome of α with lower variance than had the expert not made her recommendation.

This example gets to the heart of the sender's challenge. In giving advice, the sender must use her information, but by using her information, she makes it possible for the receiver to repurpose that information to his own ends. In simple environments informational spillover is complete and this undermines efficient cheap talk. In complex environments, the sender's challenge does not go away although it is ameliorated. It is intuitive that the sender should want to minimize the spillover as much as she can. On an unbounded space she achieves this with the first-point strategy, confining the indirect spillover to one region of the action space. As we will see, however, on a bounded space the sender gains even greater leverage when there is more informational spillover from the first-point strategy.

4.2 Bounded Action Space

A bounded action space brings several changes to the decision problem. The state contains fewer variables and the receiver is constrained in his choice should he override a recommendation. The most important difference, however, is that with positive probability there may no longer be an action that produces outcome b. Formally, on the interval $[0, q]$, the Brownian path crosses b with probability less than one.

With positive probability, therefore, the best outcome the sender can obtain is above b. The sender is worse off when this happens but so too is the receiver, and, critically, the action that is optimal for the sender is also optimal for the receiver. By modeling the mapping from actions to outcomes, we can see how misaligned preferences over outcomes can translate endogenously into aligned preferences over actions.

This possibility implies that on a bounded space the sender reveals even less information about the recommendation from the first-point strategy. The sender reveals precisely her most preferred action, but only imprecisely the outcome it produces. The receiver does not know whether the outcome is above or at b and, thus, he does not know whether his action preference is aligned with the sender or not.

That the players share a common action preference in some states is, by itself, not important for efficient cheap talk. (Indeed, here and in Section 5.2 we show it is neither necessary nor sufficient for efficient cheap talk.) What is important is what the possibility implies about other actions. Although the sender reveals *less* about the recommendation itself, she reveals *more* information about other actions. The indirect informational spillover now extends to the right as well as to the left of the recommendation. In the event that the recommendation produces an outcome above b, the receiver infers that all actions to the right are worse than the recommendation itself. That this occurs with positive probability implies the receiver's beliefs are skewed upwards and away from his ideal outcome relative to the neutral beliefs he held on an unbounded action space. In this way the sender is able to dissuade the receiver from taking actions to the right as well as to the left. How much the receiver's beliefs are skewed is critical to supporting equilibrium.

For bias less than α, the indirect spillover to the right only reinforces the receiver's incentive to accept the recommendation. Either the recommendation produces outcome b, in which case the return from overriding is not worth the risk, or the outcome is above b and all other actions produce outcomes worse than the recommendation. Therefore, if the first-point equilibrium exists on an unbounded space, it also exists on a bounded space.

For larger bias, the receiver's calculus depends on the nature of the uncertainty. With some probability the outcome of the recommendation is at b, and the receiver's best response is $\hat{a} = r^* + \frac{b-\alpha}{|\mu|}$, following Lemma 1. With complementary probability the outcome is above b and the receiver's interests are aligned with the sender on the recommendation r^*.

The essential requirement for efficient cheap talk is that the receiver resolves his response uncertainty by choosing the recommendation itself. This means that for equilibrium to hold the receiver must choose one of his two possible best responses and not the other, and, more delicately, that he must not choose an intermediate action even though all are available. Were the receiver to deviate from the recommendation to any degree, the sender, anticipating this response, would shade her recommendation to the left, and efficient cheap talk would unravel as it does in simple environments.

For the equilibrium to hold, it must be that the indirect informational spillover is strong enough that even an incremental deviation is unprofitable. Theorem 1 shows that this is possible in the Brownian environment for larger bias so long as the action space is not too large.

Theorem 1 *The first-point equilibrium exists if and only if* $q \leq q_b^{\max}$*, where:*

(i) $q_b^{\max} = \infty$ *for* $b \in [0, \alpha]$*.*

(ii) $0 < q_b^{\max} < \infty$ *for* $b \in (\alpha, \psi(0))$*.*

That more informational spillover can improve communication is surprising given the intuition from CS. In the simple environment of CS, informational spillover undermines efficient cheap talk as the sender cannot use her information and also keep it private. This intuition carries over to an unbounded action space as the sender obtains her ideal action by limiting the informational spillover and containing it in one part of the action space. The reason more informational spillover improves communication here is that the spillover shapes the receiver's beliefs in a way that systematically favors the recommendation. The sender is able to dissuade the receiver from taking actions to the right and this increases her ability to persuade the receiver to accept the recommendation. The deeper insight, therefore, is that expert power comes not just from how much information the expert can keep private, but what information she can keep private and what she can reveal.

For bias beyond $\psi(0)$, the interests of the players relative to the status quo are directly opposed and the first-point equilibrium exists only on the degenerate space of $q = 0$.[15] Interestingly, the upper bound on bias is independent of the complexity of the underlying process.[16] Thus, whenever the interests of the sender and receiver are aligned relative to the status quo, efficient cheap talk is possible if the action space is not too large.

4.3 The Mechanics of Efficient Cheap Talk

In this section we develop the key steps in the proof of Theorem 1. We decompose the theorem into two lemmas. In Lemma 3 we establish that, given q_b^{max}, the first-point equilibrium exists for all narrower action spaces. In Lemma 4 we establish that the equilibrium exists for some $q > 0$. Combined with Lemma 2, Lemmas 3 and 4 prove the theorem. We begin by characterizing the receiver's inference problem.

The Receiver's Inference Problem. We refer to Event =b as the situation in which the sender's recommendation produces outcome b, and Event >b as situations in which the outcome is strictly above b.

Event =b occurs at a recommendation r^* if the mapping first reaches outcome b at r^*. To coin a phrase, r^* represents a "first minimum" of the mapping at b. As Event =b demands nothing from the mapping beyond that, the probability that Event =b occurs at r^* can be formalized as the probability that the Brownian motion first hits b at action r^*.

---

[15] For $\psi(0) \leq \alpha$ this implies a discontinuity in q_b^{max} as b crosses $\psi(0)$.

[16] To see this note that $\psi(a) = \psi_0 + \mu a + \sigma W(a)$ on $\mathcal{A} = [0, q]$ is equivalent to $\hat{\psi} = \psi_0 + \mu q a + \sigma\sqrt{q} W(a)$ on $\mathcal{A} = [0, 1]$, and both environments have the same complexity $\alpha = \sigma^2/2|\mu|$.

Defining the first hitting action for outcome y as:

$$\tau(y) = \inf\{a \in [0,q] \mid \psi(a) = y\}.$$

We have the probability density:

$$\mathbb{P}(\text{Event =b at } m^*(\psi) = r^*) = \mathbb{P}\{\tau(b) \in dr^*\}. \tag{1}$$

In the appendix, we provide a closed form expression for this density from the hitting time formula of the Brownian motion (see Harrison (2013) for details).[17]

Event >b at r^* also represents a first-minimum of the mapping, although it differs in two respects. Working in favor of Event >b is that the first-minimum can occur at any outcome between b and $\psi(0)$. Thus, loosely speaking, there are many more paths that satisfy the first-minimum for Event >b than for Event =b. Working against Event >b is that the recommendation also represents a "last-minimum" of the path. All actions to the right produce outcomes further from b than the recommendation itself.

The probability of Event >b at r^* is the probability that a first-minimum and a last-minimum occur at the recommendation r^* for an outcome in the interval $(b, \psi(0))$. The Markov property of the Brownian motion implies that these requirements are separable. The last-minimum requirement is the probability that the Brownian path does not drop below the outcome of the recommendation in the remaining part of the action space, $(r^*, q]$.

Defining the minimum of a path over an interval $[w, x]$ as:

$$\iota(w,x) = \inf\{\psi(a) \mid a \in [w,x]\},$$

we have the probability density:

$$\mathbb{P}(\text{Event >b at } m^*(\psi) = r^*) = \int_b^{\psi(0)} \underbrace{\mathbb{P}\{\tau(y) \in dr^*\}}_{\text{first-minimum}} \cdot \underbrace{\mathbb{P}\{\iota(r^*, q) \in dy\}}_{\text{last-minimum}} dy. \tag{2}$$

---

[17] Throughout this paper, we sometimes use the term "probability" as a shorthand when discussing probability densities. This is especially relevant when referring to the probability density of the first hitting time or the location of the minimum. To highlight that these are probability densities in the formulas, we employ the notation "$\in dx$" rather than "= x."

The probability density in (2) represents a new identity: the joint distribution of the hitting time of a Brownian motion and that the hitting time is a minimum of the path. Extending a result of Shepp (1979), we derive a closed form expression for (2) in the appendix.

Upon observing a recommendation r^*, the receiver uses Equations (1) and (2) to calculate his conditional beliefs over the relative likelihood of Events =b and >b. Bayes' rule implies the receiver's belief in Event =b conditional on recommendation r^* is:

$$\mathbb{P}(\text{Event }{=}\text{b} \mid m^*(\psi) = r^*) = \frac{\mathbb{P}(\text{Event }{=}\text{b at } m^*(\psi) = r^*)}{\mathbb{P}(\text{Event }{=}\text{b at } m^*(\psi) = r^*) + \mathbb{P}(\text{Event } > \text{b at } m^*(\psi) = r^*)}. \tag{3}$$

The Size of the Action Space: The decomposition in (2) relative to (1) leads directly to the result that the first-point equilibrium exists for all action spaces narrower than q_b^{max}.

Lemma 3 *If the first-point equilibrium exists for the set of actions* $[0, q]$*, then it exists for the set of actions* $[0, q']$ *for all* $q' < q$*.*

The first-minimum requirement depends only on the mapping to the left of the recommendation, whereas the last-minimum requirement depends on the mapping to the right. Therefore, conditional on a particular recommendation r^*, a narrower action space affects only the probability of Event >b and not Event =b. In particular, as a narrower action space makes the last-minimum requirement easier to satisfy, it increases the probability of Event >b.

This can be seen through the outcome paths that satisfy the two events, as depicted in Figure 3. As the action space narrows, the set of paths that satisfy the first-minimum requirement is unchanged for a given r^*. That is to say, no paths are lost or added as q is reduced.

This is not the case for the last-minimum requirement. There are paths that fail the last-minimum requirement on a wider action space but satisfy it on a narrower space. The red path in Figure 3 is one such path. The path obtains a first minimum at r^* but fails the last-minimum requirement at $\hat{r}$ (and would, therefore, generate recommendation $\hat{r}$ rather than r^*). However, for the action space bounded by q', the red path does satisfy the last-minimum requirement, generating recommendation r^* and Event >b.

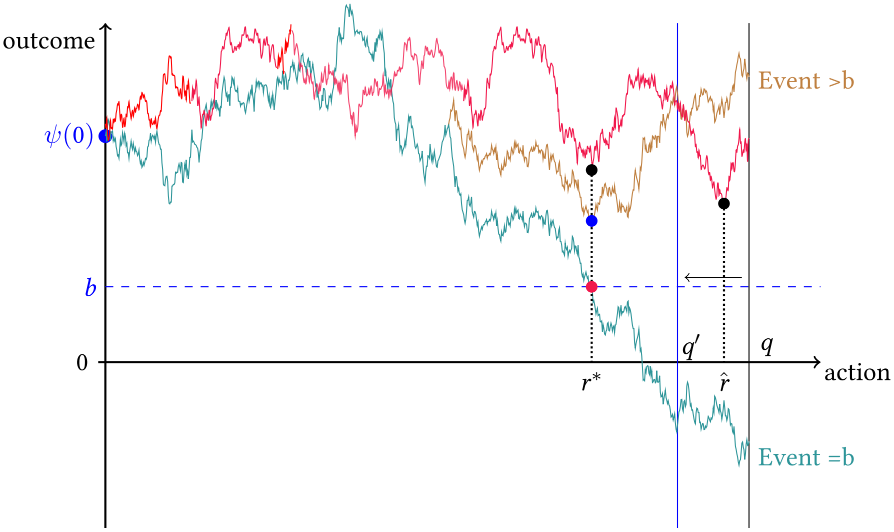

Figure 3: Paths that Induce r^* as $q \to q'$.

This implies that, following a recommendation r^*, if it is unprofitable for the receiver to deviate to $a > r^*$ in action space $[0, q]$, it is unprofitable to do so in action space $[0, q']$ when $q' < q$. By the law of total probability, the receiver either gets the same payoff as for q or an outcome from the additional paths that is strictly worse than the recommendation. The other possibility is that action a is itself no longer available in the narrower action space, in which case the deviation is moot. As actions to the left of the recommendation are dominated in both events, the dominance result in Lemma 3 follows.

Equilibrium Existence: To establish equilibrium existence, we must show, for some q, that for any possible recommendation, all deviations from the recommendation are unprofitable. It is immediate that overriding a recommendation to the left is dominated in both events. For actions to the right, overriding in Event >b is unprofitable, whereas it is profitable in Event =b when bias is larger than α.

As noted earlier, the receiver faces different best responses for the two events, either r^* or $r^* + \frac{b-\alpha}{|\mu|}$, and the challenge of efficient cheap talk is that the receiver must resolve his uncertainty by choosing r^* exactly and not $r^* + \frac{b-\alpha}{|\mu|}$ or some action between it and $r^* + \frac{b-\alpha}{|\mu|}$. Thus, the cost of deviating in Event >b must dominate the benefit in Event =b, even for small deviations.

To see why this is possible, observe that in Event =b the receiver's beliefs are neutral. Thus, deviations impose a variance cost that is linear and an expected outcome benefit that is quadratic in b, where the benefit outweighs the cost for $b > \alpha$.

In contrast, the cost of overriding the recommendation in Event >b increases much faster. The last-minimum requirement implies that the receiver's beliefs are non-neutral as outcomes are bounded below by the outcome of the recommendation. Formally, this defines a type of stochastic process known as a Brownian meander. We obtain expressions for the expected value of a Brownian meander with a known terminal value c at $a = q$. This value is continuous in a and c and we show that the derivative at the recommendation r^* is infinite for any $c > b$. Thus, by the law of iterated expectations, the marginal cost of deviating from the recommendation in Event >b is infinite.

For sufficiently small action spaces, the only way for the equilibrium to not exist is for Event =b to become infinitely more likely than Event >b. This is not true, and, in fact, the opposite holds. As the action space contracts, the probability of Event >b becomes infinitely more likely than Event =b for all available actions. Thus, for some $q > 0$, overriding the recommendation with any action is unprofitable and the first-point equilibrium exists.

Lemma 4 *For $b < \psi(0)$, the first-point equilibrium exists for some $q > 0$.*

The likelihood ratio of Event =b to Event >b is non-monotonic in the recommendation. For larger recommendations, it is more likely that the path crosses b, and so the relative probability increases that the first-minimum is at b rather than above. At the same time, the last-minimum requirement of Event >b becomes easier to satisfy as there are fewer actions to the right. At either end of the action space, Event >b dominates. It is for a recommendation internal to the action space that Event =b is most likely and the receiver's incentive to deviate is highest.

Lemmas 3 and 4, along with the earlier Lemma 2, deliver Theorem 1. We conclude this section with two notes.

Dissuasion, Persuasion, and Residual Uncertainty: Theorem 1 cannot be restated simply in terms of σ, in contrast to Lemma 2. The difference on a bounded space is that a change in σ changes what the receiver infers about the recommendation itself. The probability of Event >b relative to Event =b changes, whereas on an unbounded space the receiver's inference about the recommendation is constant in σ

(that the outcome is b with probability one). On a bounded space, this change in inference can favor Event =b for some recommendations and Event >b for others. Thus, as an increase in σ increases the receiver's residual uncertainty for other actions, it is only in an unbounded space that this translates directly into a greater willingness to accept the recommendation.

The breakdown of a tight link between σ and the equilibrium reinforces that on a bounded space the sender is doing more than simply maximizing the receiver's residual uncertainty. On a bounded space, the sender dissuades as well as persuades. He convinces the receiver that other actions are *per se* unattractive and not just that they are risky.

As a result, the equilibrium on a bounded space does not depend on the receiver being risk averse, although risk aversion does change the exact domain of existence. In Event >b the receiver is worse off with certainty, not just in expectation, should he deviate to the right. Therefore, with sufficiently high likelihood of Event >b, even a risk neutral receiver would accept the recommendation. The sender's increased power comes not from the degree of the receiver's residual uncertainty, but from the shape of that uncertainty.

Less Knowledge vs. a Smaller Action Space: In a complex environment defined by a Brownian path, bounding the action space is important because it creates the possibility for common preferences over actions despite different outcome preferences. The bounded action space also implies that the expert, in a sense, knows less, as now she knows an interval of measure q rather than the entire real line. It is important that despite knowing less, the expert still knows everything. This is important as it creates the aligned action-preference that supports the equilibrium.

To see why, suppose the action space is the real half-line but the sender knows only the interval $[0, q_b^{max}]$. Now suppose that the recommended action is q_b^{max}. The last minimum requirement is satisfied trivially in this case, making Event >b much more likely. With the last-minimum requirement redundant, Event >b reveals no information to the right of the recommendation and the receiver's beliefs are neutral. Given he has neutral beliefs in Event =b as well, it follows from Lemma 1 that the receiver will override the recommendation and choose an action to the right whenever $b > \alpha$.

Thus, the expert knowing less than the real line is by itself not sufficient to support efficient cheap talk. It is important that, in Event >b, the receiver believes that other actions deliver a worse outcome and that overriding the recommendation will be costly. In the Brownian environment, therefore, there must be enough indirect as well as direct informational spillover. The expert must dissuade as well as persuade to obtain her full power.

4.4 Welfare and Comparative Statics

The maximum size of the action space: The maximum size of the action space, q_b^{max}, decreases in the sender's bias when bias is larger than α. As the interests of the players diverge, the action space on which the first-point equilibrium exists contracts, shrinking to the status quo action itself as the bias approaches $\psi(0)$. For a fixed action space the path is more likely to cross b the larger is b, giving the receiver a greater incentive to override the recommendation. To maintain equilibrium the action space must contract in b.[18]

Proposition 1 *$q_b^{\max}$ is strictly decreasing in b for $b > \alpha$. Moreover, $q_b^{\max}$ approaches 0 as $b \to \psi(0)$ from below and approaches ∞ as $b \to \alpha$ from above.*

Equilibrium requires only that it is not too likely that the path crosses b. Thus, q_b^{max} decreases in b at a slow enough rate such that the path crosses b—and the sender and receiver have opposing interests—with substantial probability.[19]

The following comparative statics address welfare *within* the first-point equilibrium. The statements are valid for parameters for which first-point equilibria continue to exist, thus within the range of $q < q_b^{max}$.

The size of the action space: Theorem 1 establishes the upper bound on the size of the action space for the first-point equilibrium to exist. Within that bound, the utility of both players strictly increases in the size of the action space.

Corollary 1 *In the first-point equilibrium, both sender and receiver utility strictly increase in q.*

---

[18] The limiting behavior of q_b^{max} in Proposition 1 holds for arbitrary weakly concave utility with a unique maximum. The monotonicity of q_b^{max} requires an additional condition that encompasses quadratic utility.

[19] For example, setting $\mid \mu \mid = -1, \sigma^2 = 1, \psi(0) = 2$, such that $\alpha = \frac{1}{2}$, $\mathbb{E}(q_b^{max}) = b$ at $b \approx 1.21$ and the probability is greater than $\frac{1}{2}$ that the outcome of the mapping is below b at action q_b^{max}.

Expanding the action space is a public good. This is because the larger is the action space, the more likely it is that the equilibrium outcome is b. A counter-intuitive feature of the first-point equilibrium is that both players are better off when their action-preferences are misaligned than when they are aligned. Of course, should this happen too frequently, a breaking point will be reached and the first-point equilibrium will fail. However, within the bound of q_b^{max}, the larger the action space the better.

Sender's bias: The sender's bias has a different impact on utility in complex relative to simple environments. In the simple environment of CS, the sender's bias is a public bad. The larger the bias, the more inefficient is communication, and this hurts both players. In complex environments, in contrast, the sender is better off the larger is her bias, conditional on the first-point equilibrium still existing, whereas the receiver is worse off.

Corollary 2 *In the first-point equilibrium, receiver utility strictly decreases and sender utility strictly increases in b.*

The difference in complex environments is that the first-point equilibrium is efficient and sender optimal. Thus, larger bias does not bring the efficiency cost that it does in simple environments. Instead, the impact is distributional. Because the first-point equilibrium is sender-optimal, larger bias hurts the receiver because the sender's ideal outcome is then further from his own. It is more surprising that the sender benefits as she is already obtaining her best action. The reason is not because of misalignment with the receiver *per se*, but because the sender is better off the closer b is to $\psi(0)$. The larger is her bias, the more likely is the path to cross b and the more likely she obtains her ideal outcome rather than an outcome above it. If instead $\psi(0)$ and b were to increase in parallel, the sender's utility would be unchanged.

Complexity of the environment: An increase in the complexity of the environment impacts welfare differently depending on the size of the action space. On an unbounded space, the path crosses b with probability one and the outcome is b almost surely, regardless of the complexity.

On a bounded space, the path does not cross b with probability one. In fact, the probability is not everywhere monotonic in σ.[20] The limit behavior is much clearer. As σ grows large, the probability that the path hits b goes to one. Given the first-point equilibrium exists for $b \leq \alpha$ and that α increases without bound as σ gets large, we have the following result.

Corollary 3 *In the first-point equilibrium, the expected outcome approaches b as $\sigma \to \infty$.*

This implies that even on the narrowest of action spaces, as long as the environment is complex enough, the sender will obtain her ideal outcome with high likelihood.

At the other extreme, the threshold α approaches zero as complexity approaches zero. This means that efficient cheap talk is possible on an unbounded action space only for vanishingly small bias. In the limit, efficient cheap talk is possible only if bias is zero and the interests of the players are perfectly aligned. This result provides a bridge to the equilibria of CS. CS show in simple environments that the same limit is approached by the most informative partition equilibrium as bias approaches zero.

4.5 Other Efficient Equilibria

We have so far described only a single efficient equilibrium that is sender-optimal. Characterizing more equilibria is difficult when the state space is so large. Nevertheless, it is possible to make some progress on what is not an equilibrium.

For an equilibrium to be efficient, the set of equilibrium actions must have full support (except for a measure zero subset of $\mathcal{A}$). If not, then for some state the omitted action produces outcome b and the outcome of all other actions are strictly greater than b, such that the equilibrium outcome is Pareto inefficient. However, given full support, it follows that the sender must recommend her most preferred action. Thus, it is only when the sender is getting her preferred action that her incentive compatibility constraint is satisfied. For an equilibrium to be efficient, it must be sender-optimal.

Proposition 2 *The only efficient equilibria are sender optimal.*

---

[20] To see the difficulty, suppose that the state is the underlying Wiener process and that the mapping ψ is a transformation of the state that is linear in σ. An increase in σ then changes the mapping for each underlying state. For each state, if the minimum is below the drift line an increase in σ

It does not follow from Proposition 2 that the receiver-optimal equilibrium is inefficient. It may be that information is used inefficiently in every other equilibrium to such a degree that the receiver prefers an efficient sender-optimal equilibrium.

It also does not follow from Proposition 2 that the first-point equilibrium is unique as, given the potential multiplicity of the sender's preferred action, we cannot rule out equilibria in which the sender recommends one of her other preferred actions. Nevertheless, the efficient equilibria that do exist are outcome equivalent.

Sender-optimal strategies all share the same direct informational spillover from a recommendation, but the indirect informational spillover varies. For example, the last-point strategy—that reveals the largest action that is optimal for the sender—flips the logic of the first-point strategy. In the last-point strategy the indirect informational spillover in Event =b is contained to the right of the recommendation, leaving beliefs to the left neutral. This encourages the receiver to override the recommendation as he infers in Event =b that actions to the right are no worse than b, making equilibrium harder to sustain. Other efficient strategies lie somewhere between the extremes of the first- and last-point strategies.Intuition suggests, therefore, that the first-point equilibrium is the easiest to satisfy and exists for the broadest set of actions among all efficient equilibria. A proof of this claim requires formulae for the k-th hitting time of the Brownian motion and the receiver's non-neutral beliefs in Event =b. These are unavailable to us and we leave this as a conjecture for future work.

4.6 Refining Equilibria

In cheap talk, messages take on meaning only with respect to the equilibrium itself. Thus, they are interchangeable. In an influential paper, Farrell (1993) argues that the existence of a "natural language" common to the players allows for communication beyond the equilibrium specification. These *neologisms* provide an additional constraint on equilibrium. Farrell (1993) shows how this refines the set of cheap talk equilibria, often but not always eliminating babbling equilibria as well as influ-

---

moves that minimum closer to b, whereas if the minimum is above the drift line an increase in σ moves it further from b, including moving a state that crosses b for lower σ to not crossing b for higher σ. The corollary relies on the fact that with probability one the path crosses below the drift line for some action. Then, for large enough σ the outcome of this action will hit b.

ential equilibria. In Crawford and Sobel (1982) *all* equilibria are eliminated and a neologism-proof equilibrium fails to exist.

Matthews et al. (1991) develops this concept further, arguing that the credibility of out-of-equilibrium statements should be determined for the set of possible neologisms and not for neologisms individually, as in Farrell (1993). In its strong form, this expands the set of credible neologisms such that the set of *strongly announcement-proof* equilibria is a subset of neologism-proof equilibria. Matthews et al. (1991) argue for a weaker form that expands and reduces the set of credible neologisms relative to Farrell (1993) such that the set of *announcement-proof* equilibria is a superset of strong announcement-proof equilibria but not directly comparable to the set of neologism-proof equilibria.

The first-point equilibrium is strongly announcement-proof. Thus, it is also neologism-proof and announcement-proof. Moreover, all equilibria that are announcement-proof are outcome equivalent to the first-point equilibrium. This implies that all announcement-proof equilibria must be sender-optimal.

Proposition 3 *The first-point equilibrium is strongly announcement-proof. Moreover, all announcement-proof equilibria are outcome-equivalent to it.*

This result motivates our focus on the first-point equilibrium. It is robust to these demanding refinements and represents the unique outcome that satisfies announcement-proofness. It follows that the receiver-optimal equilibrium is either sender-optimal or fails announcement-proofness.

To see why the first-point equilibrium survives these refinements whereas many equilibria do not, we return to the notions of expert power and influence. The existence of a natural language allows the sender to communicate beyond the constraints of an equilibrium. This allows the sender to escape from inefficient strategies by revealing her type to the receiver in the natural language when they both can be made better off. As we saw in Proposition 2, an equilibrium that is not sender-optimal must be inefficient. Thus, any equilibrium other than the first-point, or outcome-equivalent to it, fails announcement-proofness. The first-point equilibrium is neologism- and announcement-proof precisely because it is efficient.

We do not know if other neologism-proof equilibria exist beyond the first-point equilibrium. The richness of the state space makes it difficult to identify all possible neologisms as defined by Farrell (1993) and, thus, prove an equilibrium is not neologism-proof. It is unclear what form a non-sender-optimal neologism-proof equilibrium would take if one did exist. Proving that an equilibrium fails announcement-proofness is easier. Indeed, the first-point strategy itself provides the set of credible neologisms to prove any non-sender-optimal strategy is not announcement-proof.

5 Beyond the Brownian Motion

The Brownian motion is an example of a complex environment that provides particularly clear insight into the mechanism underlying sender power and efficiency in cheap talk. It is not necessary, however, for these properties to emerge. In this section, we expand beyond the Brownian motion to provide deeper insight into cheap talk in complex environments. We have two objectives. First, to extract the key properties of the Brownian motion that support sender power and to contrast those with models in the literature. Second, to use these conditions to construct additional examples of complex environments that support sender power and efficiency in cheap talk.

5.1 Ingredients for Sender Power & Efficiency in Cheap Talk

The necessary and sufficient condition for sender power in cheap talk, trivially, is that it is in the interests of the receiver to follow the sender's recommendation of the sender's most-preferred action. We state this condition below in Definition 4. On its own, however, receiver incentive compatibility is not particularly illuminating. To better understand why it is not satisfied in simple environments, we develop two weaker but still necessary conditions—*partial invertibility* and *response uncertainty*—that help shed light on the nature of the complex environments that support sender power and efficiency in cheap talk.

Partial invertibility has played an important role throughout our analysis of the Brownian motion. It is the most basic condition necessary for efficient cheap talk. Partial invertibility requires that the receiver learns something from the sender's recommendation but not everything. Without partial invertibility, the sender cannot use her information efficiently while keeping some of it private.

Definition 2 *For the sender strategy $m : \Psi \to \mathcal{M}$, recommendation r is partially* *invertible under* $m(\cdot)$ *if* $|m^{-1}(r)| > 1$ *and* $m^{-1}(r) \subsetneq \Psi$*. The strategy* $m(\cdot)$ *is partially-invertible if all the recommendations in the range of* $m(\cdot)$ *are partially-invertible under* $m(\cdot)$*.*

Partial invertibility is the basic condition that efficient strategies fail in the simple environment of CS. If the sender reveals her most-preferred action, the receiver learns the true state precisely and chooses his best action rather than the sender's recommendation.

Partial invertibility is necessary but not sufficient for efficient cheap talk. For efficient cheap talk to emerge, it must be that the receiver is not only unsure of the state, but that he is unsure of his best response given this uncertainty. This requires that the sender pools states into a single message and that the receiver prefers different actions in at least two of the possible states. This is the notion of *response uncertainty* that we introduced earlier and formalize here. Define $\hat{a}(\psi) = \arg\max_{a \in \mathcal{A}} u^R(a, \psi)$ as the receiver's optimal action given state ψ, and, slightly abusing notation, $\hat{a}(\hat{\Psi})$ as the set of actions that are optimal for some state in the set of states $\hat{\Psi}$.

Definition 3 *A strategy* $m(\cdot)$ *satisfies response uncertainty if* $\bigcap_{\psi' \in m^{-1}(r)} \hat{a}(\psi') = \emptyset$ *for every recommendation* r *in the range of* $m(\cdot)$*.*

Response uncertainty is more demanding than partial invertibility, yet it too is insufficient to support sender power in equilibrium. This is illustrated by Morgan and Stocken's (2003) model of unknown bias. The efficient strategy in their setting satisfies partial invertibility *and* response uncertainty, but cannot support an equilibrium. Should the sender recommend her most-preferred action r^*, the receiver does not know if the sender's bias is 0 such that r^* is also his most-preferred action, or whether bias is b and his best choice is $r^* - b$.[21] Despite the receiver's response uncertainty, this is not an equilibrium. The receiver's best response is not to follow the recommendation itself, but rather to choose a compromise action between r^* and $r^* - b$.

The failure of efficient cheap talk in the model of Morgan and Stocken (2003) emphasizes the deeper challenge for the sender to hold power in cheap talk. As noted in Section 4, the sender must convince the receiver that his best response is actually

---

[21] Thus, Morgan and Stocken (2003) impose exogenously an alignment of outcome as well as action preferences with positive probability. In the Brownian environment the possible alignment of action preferences emerges endogenously despite the misalignment of outcome preferences.

the sender's ideal action, despite the receiver knowing that the sender is biased. For the sender to hold power, the receiver cannot compromise or adjust the sender's recommendation even a little bit. This is demanding requirement of receiver incentive compatibility. Denote by $a(r)$ the receiver's best response to a recommendation r.[22]

Definition 4 *A strategy m satisfies receiver incentive compatibility if for every r in the domain of $m^{-1}(\cdot)$, it holds that the receiver best response $a(r) = r$.*

The receiver's incentive compatibility can be satisfied in complex environments, as we saw for the Brownian motion.[23] As we will see below, it may hold even when the recommendation is not itself an optimal response to any individual state and, thus, the players never share a common preference over actions as they do in the Brownian motion. In what follows we will use receiver incentive compatibility, along with the two weaker requirements, to construct environments that support efficient cheap talk and illuminate why it is possible and when it is not.[24]

5.2 More Complex Environments

We present several environments that support sender power and efficiency in cheap talk. We focus on environments that differ substantively from the Brownian motion. We proceed informally here and largely via example. Formal details are in the appendix.

Discontinuous Mappings: The continuity of the Brownian path implies that nearby actions produce nearby outcomes. Our techniques extend immediately to Levy processes with positive jumps. This ensures that, as with the Brownian motion, the sender's recommendation produces an outcome either at or above b. That the outcome path may jump upwards adds variance to the expected outcome of all other actions, making deviations less profitable, and efficient cheap talk easier to sustain.

---

[22] Thus, $a(r) = \arg\max_{a \in \mathcal{A}} \mathbb{E}[u^R(a, \psi) \mid \psi \in m^{-1}(r)]$.

[23] For the Brownian motion, the sufficient condition for receiver IC to be satisfied is a joint restriction on the complexity of the environment, $\frac{\sigma^2}{|\mu|}$, and the size of the action space, q. This can readily extend to provide a sufficient condition for general stochastic processes based on the growth rate of the variance relative to the mean.

[24] The reader will have noted that we state Definitions 2-4 for arbitrary strategies and not just efficient strategies. Indeed, they define the requirements for any cheap talk equilibrium with

Minimal Complexity: In the Brownian environment the sender knows a continuum of information that the receiver does not and complexity is parameterized by the correlation across actions (σ relative to μ). Complexity can also be parameterized by the number of distinct pieces of information the sender knows that the receiver does not. In CS the gap is one. The following example extends this minimally to two pieces of information.

Consider an environment like CS with affine mappings but in which the receiver does not know the intercept as well as the slope. Specifically, suppose that for each $a \in \mathcal{A}$, there are two possible states with slope ± 1 that satisfy $\psi(a) = \psi'(a) = b$.

The sender-optimal strategy is unique: recommend the action that delivers outcome b. The receiver learns a lot from the recommendation, narrowing the set of possible states from a continuum to two. Nevertheless, the strategy satisfies partial invertibility and response uncertainty. For each recommendation r^*, the receiver is unsure whether the slope is $+1$ or -1 and, thus, whether his best response is $r^* + b$ or $r^* - b$.

To satisfy receiver incentive compatilibity, however, it must be that the receiver's best response is the recommendation r^* itself. Quadratic utility implies that for this to hold the receiver's belief about the two possible states must be perfectly balanced. The receiver would prefer a different compromise action if he assigned even a small amount of extra belief to one of the states. This is a stringent condition and, thus, whilst efficient cheap talk is possible in equilibrium, it is fragile to even the smallest perturbation.

The fragility of this equilibrium resonates with the results in models of unknown bias. As noted earlier, the sender's informational advantage in those models is also two pieces of information. The complex environment described here shows that efficient cheap talk is possible in such settings but that the conditions required are demanding.

---

positive sender bias and not just efficient equilibria. Framed this way, the deep insight of CS was to show how an equilibrium can be constructed even in simple environments. Their partition strategies obtain partial invertibility by pooling states and in simple environments this ensures response uncertainty. They then show that, given this strategy, the sender's best recommendation corresponds to the receiver's optimal compromise action. (This alignment relies on the receiver facing directional uncertainty. We develop this notion further momentarily.)

A striking feature of this example is that the sender and receiver never align on the preferred action, yet the receiver accepts the recommendation. He does so because he faces not only response uncertainty, but *directional uncertainty* as well. Directional uncertainty fills the role of the non-monotonicities in the Brownian environment. It creates the necessary uncertainty over outcomes should the receiver override the recommendation, despite the true mapping being monotonic.

Sender-Receiver Misalignment without Directional Uncertainty: In the following example the sender's advantage is again two pieces of information, yet it differs from the preceding example in two key respects. First, it shows that efficient cheap talk is possible even when the receiver faces strict directional *certainty.* Second, the equilibrium is not knife-edged despite the sender's minimal informational advantage. This shows that efficient cheap talk relies not only on how much more the sender knows than the receiver, but the nature of that information.

Consider an action space that is the set of positive integers where, for each integer $n \in \mathbb{Z}^+$, there are exactly two states such that $\psi(n) = \psi'(n) = b$. In one state $\psi(n+1) = 0$, and in the other $\psi(n+2) = 0$. All other actions produce a much worse outcome, say $\psi(a) = \psi'(a) = 100b$ for all $a \neq n$ and either $n+1$ or $n+2$, respectively.

The sender again has a unique optimal action and, as before, the receiver infers from recommendation r^* that the outcome will be exactly b. He knows for sure that his ideal action is different from the sender's, and he knows this action is strictly to the right of the recommendation. In fact, he knows that it is either $r^* + 1$ or $r^* + 2$. However, he doesn't know which and the cost of choosing the wrong one outweighs the benefit of getting it right. Thus, even though the players never have aligned action preferences and the receiver faces no directional uncertainty, he still finds it in his interests to accept the sender's recommendation.

Local Uncertainty: The nature of a sender's informational advantage also matters when that advantage is a continuum. In the Brownian environment, efficient cheap talk requires either a small bias or a bounded space. In this example we show that the same underlying degree of uncertainty can support efficient cheap talk more broadly when that uncertainty takes a different structure.

To see this, suppose that the mapping from actions to outcomes is the realized path of an Orstein-Uhlenbeck (OU) process with mean $\psi(0)$ and scale σ. The sender's advantage remains a continuum of information—indeed, the OU process is simply a different rescaling of the same underlying Wiener process as the Brownian motion. Yet because the OU process generates different beliefs for the receiver, efficient cheap talk is easier to sustain.

The OU process differs from the Brownian motion in that the process is mean-reverting. Thus, the receiver expects actions to the right of a recommendation r^* to deliver outcomes closer to $\psi(0)$ rather than to 0 as did the Brownian motion.Information about the outcome path is, in a sense, localized in the OU process, whereas it is persistent in the Brownian motion.

Cheap talk in the Ornstein-Uhlenbeck environment differs in several important respects from the Brownian environment. Even on an unbounded space, the receiver does not know whether the outcome is at or above b. This distinction is immaterial here, however, as in either event the receiver wants to follow the sender's recommendation. This implies that the first-point equilibrium exists for all biases between 0 and $\psi(0)$ whether the space is bounded or unbounded.

6 Conclusion

In this paper we have shown how expert power can derive from the complexity of the underlying environment and that this leads to communication that is more efficient. These results open up new questions and cause ostensibly settled questions to be reexamined. One example is the design of institutions. Gilligan and Krehbiel (1987) show how for simple environments institutions can rebalance power away from the receiver to the sender to open up a role for expertise. Our results suggest the opposite is required in complex environments. That the receiver would want to design an institution to weaken the sender's grip and rebalance power toward himself. Exploring this implication, as well as the many applications that CS has informed, offers the promise of a deeper understanding of the role of expertise in decision making throughout society.

References

Aghion, Philippe and Jean Tirole (1997): "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105 (1), 1–29.

Arrow, Kenneth J (1963): "Uncertainty and the Welfare Economics of Medical Care," *American Economic Review*, 53 (5), 941–973.

Bagnoli, Mark and Ted Bergstrom (2006): "Log-concave probability and its applications," in *Rationality and Equilibrium: A Symposium in Honor of Marcel K. Richter*, Springer, 217–241.

Bardhi, Arjada (2022): "Attributes: Selective learning and influence," Working Paper.

Bardhi, Arjada and Nina Bobkova (2023): "Local evidence and diversity in minipublics," *Journal of Political Economy*, 131 (9), 2451–2508.

Callander, Steven (2008): "A Theory of Policy Expertise," *Quarterly Journal of Political Science*, 3 (2), 123–140.

——— (2011): "Searching and Learning by Trial and Error," *The American Economic Review*, 101 (6), 2277–2308.

Callander, Steven, Nicolas Lambert, and Niko Matouschek (2021): "The Power of Referential Advice," *Journal of Political Economy*, 129 (11), 3073–3140.

Carnehl, Christoph and Johannes Schneider (2021): "A Quest for Knowledge," Working paper.

Cetemen, Doruk, Can Urgun, and Leeat Yariv (2023): "Collective progress: Dynamics of exit waves," *Journal of Political Economy*, 131 (9), 2402–2450.

Chakraborty, Archishman and Rick Harbaugh (2007): "Comparative Cheap Talk," *Journal of Economic Theory*, 132 (1), 70–94.

——— (2010): "Persuasion by Cheap Talk," *The American Economic Review*, 100 (5), 2361–2382.

Crawford, Vincent P and Joel Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50 (6), 1431–1451.

Farrell, Joseph (1993): "Meaning and credibility in cheap-talk games," *Games and Economic Behavior*, 5 (4), 514–531.

Garfagnini, Umberto and Bruno Strulovici (2016): "Social Experimentation with Interdependent and Expanding Technologies," *Review of Economic Studies*, 83 (4), 1579–1613.

Gilligan, Thomas and Keith Krehbiel (1987): "Collective Decision Making and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures," *Journal of Law, Economics, and Organization*, 3 (2), 287–335.

Gordon, Sidartha (2010): "On infinite cheap talk equilibria," *manuscript, University of Montreal.*

Gruber, Jonathan and Maria Owings (1996): "Physician Financial Incentives and Cesarean Section Delivery," *The RAND Journal of Economics*, 27 (1), 99–123.

Harrison, J Michael (2013): *Brownian models of performance and control*, Cambridge University Press.

Kamenica, Emir and Matthew Gentzkow (2011): "Bayesian Persuasion," *The American Economic Review*, 101 (6), 2590–2615.

Karatzas, Ioannis and Steven Shreve (2012): *Brownian motion and stochastic calculus*, vol. 113, Springer Science & Business Media.

Levitt, Steven D and Chad Syverson (2008): "Market Distortions when Agents are Better Informed: The Value of Information in Real Estate Transactions," *The Review of Economics and Statistics*, 90 (4), 599–611.

Li, Ming and Kristóf Madarász (2008): "When mandatory disclosure hurts: Expert advice and conflicting interests," *Journal of Economic Theory*, 139 (1), 47–74.

Lipnowski, Elliot and Doron Ravid (2020): "Cheap talk with transparent motives," *Econometrica*, 88 (4), 1631–1660.

Matthews, Steven A, Masahiro Okuno-Fujiwara, and Andrew Postlewaite (1991): "Refining cheap-talk equilibria," *Journal of Economic Theory*, 55 (2), 247–273.

Milgrom, Paul and John Roberts (1988): "Economic Theories of the Firm: Past, Present, and Future," *The Canadian Journal of Economics*, 21 (3), 444–458.

Milgrom, Paul R (1981): "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 12 (2), 380–391.

Morgan, John and Phillip C Stocken (2003): "An analysis of stock recommendations," *RAND Journal of Economics*, 183–203.

Shepp, Lawrence A (1979): “The joint density of the maximum and its location for a Wiener process with drift,” *Journal of Applied probability*, 16 (2), 423–427.

Sobel, Joel (2010): “Giving and Receiving Advice,” in *Advances in Economics and Econometrics: Tenth World Congress*, ed. by Daron Acemoglu, Manuel Arellano, and Eddie Dekel, Princeton University Press, 305–341.

——— (2012): “Complexity versus conflict in Communication,” in *2012 46th Annual Conference on Information Sciences and Systems (CISS)*, 1–6.

Urgun, Can and Leeat Yariv (2021a): “Constrained retrospective search,” *AEA Papers and Proceedings*, 111, 549–53.

——— (2021b): “Retrospective search: Exploration and ambition on uncharted terrain,” Tech. rep., National Bureau of Economic Research.

Weber, Max (1958): *From Max Weber: Essays in Sociology*, Oxford University Press: New York.

A Formal Details and Proofs

A.1 Formal Details of the Environment

States and Beliefs: The state of the world $\psi(\cdot)$ is a transformation of the Wiener process $W(\cdot)$ with parameters $\psi_0, \mu \in \mathbb{R}$ and $\sigma^2 \in \mathbb{R}_+$ given by $\psi(a) = \psi_0 + \mu a + \sigma W(a)$.Realization of $W(\cdot)$, and thus $\psi(\cdot)$, are the private information of the sender. The receiver has a prior belief $\omega(\cdot)$ over $W(\cdot)$ given by the Wiener measure on $(\mathcal{W}, \mathcal{B}(\mathcal{W}))$.[25] As the Wiener process $W(\cdot)$ only affects the payoffs through the outcome mapping $\psi(\cdot)$, we will refer to the induced beliefs about $\psi(\cdot)$ instead of $W(\cdot)$.

Equilibrium: We denote a Perfect Bayesian Equilibrium by $\mathcal{E} = (\omega(\cdot \mid \cdot), a(\cdot), m(\cdot))$ where $m : \Psi \to \mathcal{M}$ is the sender's strategy, $a : \mathcal{M} \to \mathcal{A}$ is the receiver's strategy, and a family of probability measures $\omega(\cdot \mid r \in m(\psi)) : \mathcal{B}(\mathcal{C}[0,1]) \times \mathcal{M} \to [0,1]$.[26] Equilibrium requires that the following hold: (i) $\omega(\psi \mid r \in m(\psi))$ is obtained from the prior using Bayes's rule whenever possible, (ii) Receiver's Incentive Compatibility: $a(r) \in \arg\max_{a' \in \mathcal{A}} \mathbb{E}[u_R(a', \psi) \mid \omega(\psi \mid r \in m(\psi))]$ for every $r \in \mathcal{M}$, and (iii) Sender's

---

[25] $\mathcal{B}$ denotes the Borel sigma algebra. The reader is referred to Karatzas and Shreve (2012) for a detailed discussion of the Wiener measure $\omega(\cdot)$.

[26] Formally, $\omega(\psi \mid r \in m(\psi)) = \mathbb{E}[1_{\psi \in \Psi} \mid \psi \in m^{-1}(r)]$.

Incentive Compatibility: $m(\psi) \in \arg\max_{r' \in \mathcal{M}} u_S(a(r'), \psi)$ for every $\psi \in \Psi$.

In our analysis, we utilize two versions of the process $\psi(\cdot)$. The first is the drifting Brownian Motion (BM), expressed as $X(a') = \mu a' + \sigma W(a')$, which shares the same distribution as the a' increment of the outcome map $\psi(a)$. The second is the Brownian Meander (we often plainly refer to it as Meander), denoted by $M(a,q) := X(a) \ \mid \ X(a') \geq 0 \ \forall a' \in [0,q]$.

Our analysis also employ random variables associated with $\psi(\cdot)$. $\tau(x) := \inf a \in \mathbb{R} \mid \psi(a) = x$ represents the first hitting action (time) of the outcome $x \in \mathbb{R}$. We also use the infimum over the interval $[0,q]$, denoted $\iota(q) := \inf \psi(a) \mid a \in [0,q]$. Finally, $\tau_\iota(q) := \tau(\iota(q))$ denotes the first hitting action (time) of the minimum over $[0,q]$.

We restate the first-point strategy in terms of these random variables by partitioning the recommendation into the two events, $\psi(a) = b$ and $\psi(a) > b$:

$$m^*(\psi) = \begin{cases} \min \left\{ a \in [0,q] : \psi(a) = b \right\} & \text{if } \exists a \in [0,q] \ \psi(a) = b \\ \min \left\{ a' \in [0,q] : \psi(a') = \iota(q) \right\} & \text{if } \forall a \in [0,q] \ \psi(a) > b \end{cases} = \begin{cases} \tau(b) & \text{if } \tau(b) \leq q \\ \tau_\iota(q) & \text{if } \iota(q) > b \end{cases}$$

A.2 Proofs for Results in the Text

Throughout the proofs, we simplify notation by omitting the ψ argument from $u_R(a,\psi)$ and $u_S(a,\psi)$. In some proofs, we fix all parameters of the game except for one and change the remaining parameter. Whenever this is the case, we subscript the strategy with the changing parameter, e.g. $m^*_q(\psi)$ when changing q and fixing other parameters. At several points we call upon technical properties of stochastic processes and closed form expressions of certain distributions. The proofs for these properties and the derivation of expressions are detailed separately in the online appendix. The online appendix also contains the proofs of Corollaries 1-3.

Proof of Lemma 1. By the mean-variance representation of quadratic utility, the receiver's expected utility is: $\mathbb{E}[u_R(a)] = -\left[\psi(0) + \mu a\right]^2 - \sigma^2 a$. The first and second order conditions for optimality are:

$$\frac{d\mathbb{E}[u_R(a)]}{da} = -2\mu\left[\psi(0) + \mu a\right] - \sigma^2, \qquad \frac{d^2\mathbb{E}[u_R(a)]}{da^2} = -2\mu^2 \leq 0.$$

The result follows from the first order condition. ■

Proof of Lemma 2. It is a well-known mathematical fact that $\mathbb{P}(\tau(b) < \infty) = 1$, i.e. almost every path eventually (in finite time) hits b. Thus, for every message realization r^* of the first-point strategy $m^*(\psi)$, we have that $\mathbb{P}(\psi(r^*) = b \mid m^*(\psi) = r^*) = 1$ whenever $q = \infty$. Then, by Lemma 1, there are no profitable deviations to $\hat{a} \in \mathbb{R}_+$ if and only if $b \leq \alpha$. ■

Proof of Lemma 3. Suppose that a first-point equilibrium exists for the game with action space $\mathcal{A} = [0, q]$ for some $q \in \mathbb{R}_{++}$, and fixed $\psi_0 > b > 0$, μ and σ. We denote the corresponding first-point strategy of the sender by $m_q^*(\cdot)$. The receiver's incentive compatibility implies that for a recommendation $r^* \in [0, q]$, the deviation to action $\hat{a} \in [0, q]$ is not profitable: $0 \geq \mathbb{E}[u_R(\hat{a}) - u_R(r^*) \mid m_q^*(\psi) = r^*]$. By the law of total probability, this implies:

$$\begin{aligned}
0 &\geq \mathbb{P}(\tau(b) \in dr^* \mid m_q^*(\psi) = r^*)\mathbb{E}[u_R(\hat{a}) - u_R(r^*) \mid \psi(r^*) = b] \\
&\quad + \mathbb{P}(\tau_\iota(q) \in dr^*, \iota(q) > b \mid m_q^*(\psi) = r^*)\mathbb{E}[u_R(\hat{a}) - u_R(r^*) \mid \tau_\iota(q) = r^*, \iota(q) > b] \\
&= \frac{\mathbb{P}(\tau(b) \in dr^*)}{\mathbb{P}(\tau(b) \in dr^*) + \mathbb{P}(\tau_\iota(q) \in dr^*, \iota(q) > b)}\mathbb{E}[u_R(\hat{a}) - u_R(r^*) \mid \psi(r^*) = b] \\
&\quad + \frac{\mathbb{P}(\tau_\iota(q) \in dr^*, \iota(q) > b)}{\mathbb{P}(\tau(b) \in dr^*) + \mathbb{P}(\tau_\iota(q) \in dr^*, \iota(q) > b)}\mathbb{E}[u_R(\hat{a}) - u_R(r^*) \mid \tau_\iota(q) = r^*, \iota(q) > b]
\end{aligned}$$

Now consider the game on the action space $\mathcal{A}' = [0, q']$ where $q' < q$ and the corresponding first-point strategy is $m_{q'}^*(\cdot)$. Again, we have that $m_{q'}^*(\psi) = r^*$ if and only (i) $\tau(b) = r^*$ or (ii) $\tau_\iota(q') = r^*$ with $\iota(q') > b$.

The second set of paths, $\{\psi \in \Psi \mid \tau_\iota(q') = r^*, \iota(q') > b\}$, can be partitioned into two: Paths that satisfy $\tau_\iota(q) = r^*$ i.e. $\{\psi \in \Psi \mid \tau_\iota(q) = r^*, \iota(q) > b\}$, and those that do not, $\{\psi \in \Psi \mid \tau_\iota(q') = r^*, \iota(q') > b, \tau_\iota(q) > q'\}$. We can write the expected change in payoff for action $\hat{a} \in [0, q']$ when the recommendation is $r^* \in [0, q']$ using the law of total probability $\mathbb{E}[u_R(\hat{a}) - u_R(r^*) \mid m_{q'}^*(\psi) = r^*]$ can be described as:

$$\begin{aligned}
&\frac{\mathbb{P}(\tau(b) \in dr^*)}{\mathbb{P}(\tau(b) \in dr^*) + \mathbb{P}(\tau_\iota(q') \in dr^*, \iota(q') > b)}\mathbb{E}[u_R(\hat{a}) - u_R(r^*) \mid \psi(r^*) = b] \\
+&\frac{\mathbb{P}(\tau_\iota(q') \in dr^*, \iota(q') > b))}{\mathbb{P}(\tau(b) \in dr^*) + \mathbb{P}(\tau_\iota(q') \in dr^*, \iota(q') > b)}\mathbb{E}[u_R(\hat{a}) - u_R(r^*) \mid \tau_\iota(q') = r^*, \iota(q') > b]
\end{aligned}$$

Which is equivalent to:

$$
\begin{aligned}
=&\frac{\mathbb{P}(\tau(b)\in dr^*)\mathbb{E}[u_R(\hat{a})-u_R(r^*)\mid \psi(r^*)=b]}{\mathbb{P}(\tau(b)\in dr^*)+\mathbb{P}(\tau_\iota(q)\in dr^*,\iota(q)>b)+\mathbb{P}(\tau_\iota(q')\in dr^*,\iota(q')>b,\tau_\iota(q)>q')}\\
+&\frac{\mathbb{P}(\tau_\iota(q)\in dr^*,\iota(q)>b)\mathbb{E}[u_R(\hat{a})-u_R(r^*)\mid \tau_\iota(q)=r^*,\iota(q)>b]}{\mathbb{P}(\tau(b)\in dr^*)+\mathbb{P}(\tau_\iota(q)\in dr^*,\iota(q)>b)+\mathbb{P}(\tau_\iota(q')\in dr^*,\iota(q')>b,\tau_\iota(q)>q')}\\
+&\frac{\mathbb{P}(\tau_\iota(q')\in dr^*,\iota(q')>b,\tau_\iota(q)>q')\mathbb{E}[u_R(\hat{a})-u_R(r^*)\mid \tau_\iota(q')=r^*,\iota(q')>b,\tau_\iota(q)>q']}{\mathbb{P}(\tau(b)\in dr^*)+\mathbb{P}(\tau_\iota(q)\in dr^*,\iota(q)>b)+\mathbb{P}(\tau_\iota(q')\in dr^*,\iota(q')>b,\tau_\iota(q)>q')}
\end{aligned}
$$

The expectation in the last expression is conditional on $\tau_\iota(q')=r^*$, hence it directly follows that it is negative. The remaining part is proportional to receiver incentive compatibility condition for the game with action space $[0,q]$, adjusted with probability weights, and it is negative by assumption. Thus, we can rewrite the above expression:

$$
\begin{aligned}
&\mathbb{E}[u_R(\hat{a})-u_R(r^*)\mid m^*_{q'}(\psi)=r^*]\\
=&\frac{\mathbb{P}(\tau(b)\in dr^*)+\mathbb{P}(\tau_\iota(q')\in dr^*,\iota(q')>b)}{\mathbb{P}(\tau(b)\in dr^*)+\mathbb{P}(\tau_\iota(q)\in dr^*,\iota(q)>b)}\overbrace{\mathbb{E}[u_R(\hat{a})-u_R(r^*)\mid m^*_q(\psi)=r^*]}^{\leq 0}\\
+&\frac{\mathbb{P}(\tau_\iota(q')\in dr^*,\iota(q')>b,\tau_\iota(q)>q')\overbrace{\mathbb{E}[u_R(\hat{a})-u_R(r^*)\mid \tau_\iota(q')=r^*,\iota(q')>b,\tau_\iota(q)>q']}^{\leq 0}}{\mathbb{P}(\tau(b)\in dr^*)+\mathbb{P}(\tau_\iota(q)\in dr^*,\iota(q)>b)+\mathbb{P}(\tau_\iota(q')\in dr^*,\iota(q')>b,\tau_\iota(q)>q')}\leq 0
\end{aligned}
$$

where the expectation in the second term is negative, since the expectation is conditional on $r^*=\tau_\iota(q')$ and $\hat{a}<q'$ combined with $u_R(a,\psi)$ being weakly concave in $\psi(a)$ and maximized at $\psi(0)$. Thus, if the first-point equilibrium exists for $\mathcal{A}=[0,q]$, then it exists for $[0,q']$. ■

Proof of Lemma 4. Consider the game with $\mathcal{A}=[0,q]$, for some $q\in\mathbb{R}_{++}$. The first-point equilibrium has full support over $\mathcal{A}$. Let any off-path recommendation $r'\notin\mathcal{A}$ be interpreted as an on-path message, say $r''=0$. Thus, it is sufficient to show there are no on-path deviations to establish the equilibrium.

The sender's incentive compatibility is immediate as the recommendation implements her best action. Consider the receiver's utility upon seeing message $m^*_q(\psi)=r^*$ and taking action $\hat{a}$. It is straightforward to observe that any action $\hat{a}<r^*$ is strictly dominated by r^* by the construction of first-point strategy. Now consider a deviation to action $\hat{a}=r^*+a'$ for some $a'>0$. For every concave utility function $u_R(\cdot)$ that is uniquely maximized at 0, we have $\mathbb{E}[u_R(a'+r^*)-u_R(r^*)\mid m^*_q(\psi)=r^*]\leq 0$ whenever the following two conditions hold: (i) $\mathrm{Var}[\psi(a'+r^*) \mid m_q^*(\psi)=r^*] \geq \mathrm{Var}[\psi(r^*) \mid m_q^*(\psi)=r^*]$, and (ii) $\mathbb{E}[\psi(a'+r^*) \mid m_q^*(\psi)=r^*] \geq \mathbb{E}[\psi(r^*) \mid m_q^*(\psi)=r^*] > 0$.

Recall that $X(\cdot)$ denotes the BM with initial point 0, drift μ and scale σ, and $M(\cdot,k)$ is the corresponding Meander of length k. By the stationary independent-increments property of BM, it follows that the random variable $\psi(a'+r^*)$, conditional on $m_q^*(\psi)=r^*$ and the realization of $\psi(r^*) \in [b,\psi_0]$, is equal to the random variable: $\psi(r^*) + \mathbb{1}_{\{\psi(r^*)>b\}} M(a', q-r^*) + \mathbb{1}_{\{\psi(r^*)=b\}} X(a')$ in probability law.

Thus, it directly follows that condition (i) holds. By the law of total probability, the LHS of condition (ii) is given by:

$$\begin{aligned}
&\mathbb{E}[\psi(a'+r^*) - \psi(r^*) \mid m_q^*(\psi)=r^*] \\
&= \mathbb{P}(\tau(b) \in dr^* \mid m_q^*(\psi)=r^*)\mathbb{E}[\psi(a'+r^*) - \psi(r^*) \mid \psi(r^*)=b] \\
&+ \mathbb{P}(\tau_\iota(q) \in dr^*, \iota(q) > b \mid m_q^*(\psi)=r^*)\mathbb{E}[\psi(a'+r^*) - \psi(r^*) \mid \tau_\iota(q)=r^*, \iota(q)>b]
\end{aligned}$$

The stationary independent increments property implies that: $\mathbb{E}[\psi(a'+r^*) - \psi(r^*) \mid \psi(r^*)=b] = \mathbb{E}[X(\cdot)] = \mu a'$, and moreover $\mathbb{E}[\psi(a'+r^*) - \psi(r^*) \mid \tau_\iota(q)=r^*, \iota(q) > b] = \mathbb{E}[X(a') \mid \min\{X(a'') : a'' \leq q-r^*\} > 0] = \mathbb{E}[M(a', q-r^*)]$.

We utilize two technical properties of these processes that we prove in the online appendix. In the online appendix Lemma B.1, we show that: $\lim_{r^* \to 0} \mathbb{P}(\tau(b) \in dr^* \mid m_q^*(\psi)=r^*) = 0$. Using the definition of a limit, this implies that $\forall \varepsilon > 0$ there exists a $\delta_\varepsilon > 0$ such that $\mathbb{P}(\tau(b) \in dr^* \mid m_q^*(\psi)=r^*) \leq \varepsilon$ whenever $r^* \leq \delta_\varepsilon$. Similarly, in the online appendix Corollary B.1 we show that: $\lim_{a' \to 0} \frac{\partial}{\partial a'} \mathbb{E}[M(a', q-r^*)] = \infty$. Thus, for every $N > 0$ there exists a $\delta_N \in \mathbb{R}_+$ such that $\mathbb{E}[M(a', q-r^*)] > Na'$ whenever $a' < \delta_N$.

Now let ε and N such that $\varepsilon\mu + (1-\varepsilon)N \geq 0$, and let q be such that $q < \min\{\delta_\varepsilon, \delta_N\}$. We have that $r^* < q = \min\{\delta_\varepsilon, \delta_N\}$ and $a' < q = \min\{\delta_\varepsilon, \delta_N\}$. So, it follows that for every r^* and a' such that $r^* + a' \leq q$, we have that $\mathbb{E}[\psi(a'+r^*) - \psi(r^*) \mid m_q^*(\psi)=r^*]$ is given by:

$$\underbrace{\mathbb{P}(\tau(b) \in dr^* \mid m_q^*(\psi)=r^*)}_{\leq \varepsilon} \mu a' + \underbrace{\mathbb{P}(\tau_\iota(q) \in dr^*, \iota(q) > b \mid m_q^*(\psi)=r^*)}_{\geq 1-\varepsilon} \underbrace{\mathbb{E}[\psi(a'+r^*) - \psi(r^*) \mid \tau_\iota(q)=r^*, \iota(q)>b]}_{=\mathbb{E}[M(a',q-r^*)] > Na'}$$

$$\geq \varepsilon(\mu a') + (1-\varepsilon)Na' = a' \underbrace{(\varepsilon\mu + (1-\varepsilon)N)}_{\geq 0} \geq 0$$

Hence, it follows that a first-point equilibrium exists whenever $q < \min\{\delta_\varepsilon, \delta_N\}$. ■

Proof of Theorem 1. Theorem 1 directly follows from Lemmata 2, 3 and 4. More precisely, consider the game with $\mathcal{A} = [0, q]$. Lemma 4 shows that an equilibrium exists for some $q \in \mathbb{R}_{++}$, and Lemma 3 shows if an equilibrium exists for such q, it exists for every $q' < q$. By Lemma 2, there exists an equilibrium with $q = \infty$ if and only if $b \leq \alpha$. Hence $q_b^{\max} = \infty$ if and only if $b \leq \alpha$, and a finite number otherwise. ■

Proof of Proposition 1. Suppose that $\psi(0) > \alpha$ and denote the corresponding first-point strategy for a given b by $m_b^*(\psi)$. For $q = q_b^{\max}$, the following holds:

$$0 \geq \mathbb{E}\left[u_R(a) - u_R\left(r^*\right) \mid m_b^*(\psi) = r^*\right] \leq 0 \quad \forall a, r^* \in \left[0, q_b^{\max}\right],$$

and there exists some $\tilde{a}, \tilde{r} \in \left[0, q_{\max}^b\right]$ with $\tilde{a} = \tilde{r} + a'$ and $a' > 0$ such that this holds with equality by the maximality of $q_b^{\max}$.[27] We can write this as:

$$\begin{aligned}
0 &= \mathbb{E}\left[u_R\left(a' + \tilde{r}\right) - u_R\left(\tilde{r}\right) \mid m_b^*(\psi) = \tilde{r}\right] \\
&= \mathbb{P}\left(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r}\right) \mathbb{E}\left[u_R\left(\psi(a' + \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \psi\left(\tilde{r}\right) = b\right] \\
&+ \mathbb{P}\left(\tau_\iota(q) \in d\tilde{r}, \iota(q) > b \mid m_b^*(\psi) = \tilde{r}\right) \mathbb{E}\big[u_R\left(\psi(a' + \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\big] \\
0 &= \mathbb{P}\left(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r}\right) \mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] \qquad (4) \\
&+ \mathbb{P}\left(\tau_\iota(q) \in d\tilde{r}, \iota(q) > b \mid m_b^*(\psi) = \tilde{r}\right) \mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q - \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right].
\end{aligned}$$

The last line follows from rewriting in terms of increments given by $X(\cdot)$ and $M(\cdot, \cdot)$.

In order to show that, $\tilde{a}, \tilde{r}$ constitutes a profitable deviation for $q' > q$, it is sufficient to show that this indifference condition has a strictly positive derivative with respect to b. Suppose that $u_R(\cdot)$ satisfies the condition:

$$\frac{\partial}{\partial b} \log \mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] \geq \frac{\partial}{\partial b} \log \mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q - \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right] \tag{5}$$

Proposition B.1 in the the online appendix shows that this condition is satisfied by the quadratic utility. The derivative of the indifference condition (4) is given by:

$$\left(\frac{\partial}{\partial b} \mathbb{P}\left(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r}\right)\right) \mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] \tag{6}$$

$$+ \left(\frac{\partial}{\partial b} \mathbb{P}\left(\tau_\iota(q) \in d\tilde{r}, \iota(q) > b \mid m_b^*(\psi) = \tilde{r}\right)\right) \mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q - \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right] \tag{7}$$

$$+ \mathbb{P}\left(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r}\right) \left(\frac{\partial}{\partial b} \mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right]\right) \tag{8}$$

$$+ \mathbb{P}\left(\tau_\iota(q) \in d\tilde{r}, \iota(q) > b \mid m_b^*(\psi) = \tilde{r}\right) \left(\frac{\partial}{\partial b} \mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q - \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right]\right) \tag{9}$$

---

[27] It follows that $\tilde{a} > \tilde{r}$. For any $\tilde{a}$ smaller the value is negative. Thus, $\tilde{a} = \tilde{r} + a'$ for some $a' > 0$.

In the online appendix Lemma B.3, we show that $\mathbb{P}(\tau(b) \in d\tilde{r})$ is log-concave in b, and we conclude that $\mathbb{P}(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r})$ is increasing in b, by Bagnoli and Bergstrom (2006). Thus, we conclude that:

$$\frac{\partial}{\partial b}\mathbb{P}\left(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r}\right) > 0 > \frac{\partial}{\partial b}\mathbb{P}\left(\tau_\iota(q) \in d\tilde{r}, \iota(q) > b \mid m_b^*(\psi) = \tilde{r}\right).$$

By the properties of $u_R(\cdot)$ we have the following inequalities.[28]

$$\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] > 0 > \mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q - \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right]$$

Thus, the terms (6) and (7) are positive for any weakly concave utility function that is uniquely maximized at 0. Similarly, it directly follows that the sum of (8) and (9) are non-negative if and only if the following holds.[29]

$$\frac{\mathbb{P}\left(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r}\right)}{\mathbb{P}\left(\tau_\iota(q) \in d\tilde{r}, \iota(q) > b \mid m_b^*(\psi) = \tilde{r}\right)} \geq -\frac{\frac{\partial}{\partial b}\mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q - \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right]}{\frac{\partial}{\partial b}\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right]} \tag{10}$$

However, rearranging the indifference condition arising from the definition of $q_b^{\max}$ given by (4), we have that:

$$\frac{\mathbb{P}\left(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r}\right)}{\mathbb{P}\left(\tau_\iota(q) \in d\tilde{r}, \iota(q) > b \mid m_b^*(\psi) = \tilde{r}\right)} = -\frac{\mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q - \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right]}{\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right]}. \tag{11}$$

Using (11), the condition (10) reduces to the equation (5), which is assumed to hold. Thus, under condition (5), the first-point equilibrium does not exist for any bias $b' > b$ and action space of length $q_b^{\max}$.[30] By Lemmata 3 and 4, we conclude that $q_b^{\max} > q_{b'}^{\max}$.

To study the limits, we denote the best deviation by the receiver, conditional on the event $\tau(b) = r^*$, by $a'(r^*)$. Applying Lemma 1, $a'(r^*)$ is given by $\mathbb{E}[\psi(a') \mid \psi(r^*) = b] = \psi(r^*) + \mu(a' - r^*) = \alpha$. Let $b \to \alpha$, then it follows that $a'(r^*) \to r^*$ and $\psi(a') \to b$.

Moreover, using the online appendix Corollary B.1, we show that $\lim_{a' \to r^*} \frac{\partial}{\partial a'}\mathbb{E}[\psi(a') -$

---

[28] $M(a', \cdot) > 0$ in every realization. Since, $u_R(\cdot)$ is weakly concave and uniquely maximized at 0 this implies that $0 > \mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q - \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right]$. Equation (4) necessitates that $\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] > 0$.

[29] Whenever the denominator is 0, the condition is violated unless the numerator is 0. For simplicity, we can take $\frac{0}{0} = 0$ for the RHS (without loss), and state this in terms of ratios.

[30] $\mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q - \tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right] < 0$ and $\frac{\partial}{\partial b}\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] > 0$ follows from weak-concavity and unique maximum at 0,. Rearranging RHS of (10) we get (5).

$\psi(r^*) \mid \tau_\iota(q) = r^*, \iota(q) > b] = \infty$. It is immediate to conclude that $\lim_{a' \to r^*} \frac{\partial}{\partial a'} \mathbb{E}[u_R(a') - u_R(r^*) \mid \tau_\iota(q) = r^*, \iota(q) > b] = -\infty$, as discussed in Lemma 4.

Finally, for any $q \in \mathbb{R}_+$, we have that $\tau_\iota(q) = r^*$, $\iota(q) > b$ has strictly positive probability for every $r^* \in [0, q]$ and $b < \psi(0)$. Thus, for any finite q as $b \to \alpha$, the expected payoff of a deviation from first-point equilibrium has a strictly negative payoff. We conclude $q_b^{\max} \to \infty$.

Letting $b \to \psi(0)$, we have that $\frac{\mathbb{P}(\tau(b) \in dr^*)}{\mathbb{P}(\tau_\iota(q) \in dr^*)} \to \infty$ for every $q \in \mathbb{R}_{++}$. Thus, for every action space $\mathcal{A} = [0, q]$ with $q > 0$ and corresponding first-point strategy $m_q^*(\psi)$, we have that $\mathbb{P}(\tau(b) \in r^* \mid m_q^*(\psi) = r^*) = 1$. So, for every $q > 0$ there exists a profitable deviation by Lemma 1. We conclude that $q_b^{\max} \to 0$ as $b \to \psi_0$. ■

Proof of Proposition 2. Suppose that $(m(\cdot), a(\cdot), \omega(\cdot \mid \psi(\cdot) \in m^{-1}(\cdot))$ is an equilibrium, and it is ex-post Pareto efficient. A necessary condition for efficient equilibrium is that $a(\cdot)$ takes all the values in $[0, q]$. Suppose not, let $\hat{a} \in \mathcal{A}$ and $\hat{a} \notin a(m(\Psi))$. There exists a path realization $\hat{\psi}(\cdot)$ with $\hat{\psi}(\hat{a}) = b$, and $\hat{\psi}(a') > b$ $\forall a' \in [0, q] \setminus \{\hat{a}\}$. This implies that both players can be made strictly better of with action $\hat{a}$ and the equilibrium is not Pareto efficient. Thus, $a(m(\Psi)) = \mathcal{A}$.

Then consider an efficient equilibrium $(m^*(\psi), a(\cdot), \omega(\cdot \mid \psi \in m^{*-1}(\cdot))$ with $a(m^*(\Psi)) = \mathcal{A}$. For any receiver strategy $a(\cdot)$, to satisfy sender incentive compatibility, the equilibrium recommendation $r^* = m^*(\psi)$ must satisfy $a(r^*) \in \arg\max_{a \in \mathcal{A}}[-(\psi(a) - b)^2]$, and the equilibrium is sender-optimal. ■

Proof of Proposition 3. An announcement strategy (n, Θ) specifies the set Θ of all types that deviate from the equilibrium , and their communication strategy $n : \Theta \to \Delta(\mathcal{M})$. A credible announcement strategy (n, Θ) satisfies conditions (A1)-(A4). In Section F of the Online Appendix, we formalize the details for announcement-proof equilibria. For specifics regarding the notation and formal statements of conditions (A1)-(A4), readers are directed there. Below, we briefly outline these conditions.

(A1) and (A2) state that the equilibrium is stable to deviations by a credible announcement strategy. (A1) ensures that types deviating with an announcement, $\psi \in \Theta$, prefer the lowest possible payoff from the announcement strategy over the highest possible payoff from the equilibrium. (A2) ensures non-deviating types, $\psi' \in \Psi \setminus \Theta$, prefer the highest possible payoff from the equilibrium over the lowest possible payoff from the announcement strategy.

(A3) states that the announcement strategy is internally consistent: Each deviating sender type $\psi \in \Theta$ prefers to send the message prescribed by the announcement strategy $s \in n(\psi)$ over deviations to other messages $s' \in n(\Theta)$ in the support of the announcement strategy.

Finally, (A4) states that there is no other announcement strategy (n', Θ') for which a deviating type $\psi \in \Theta \cap \Theta'$ is strictly better off under (n', Θ') compared to (n, Θ).

We proceed with the proof of our claim, which follows analogously to Proposition 6.1 in Matthews, Okuno-Fujiwara, and Postlewaite (1991).

Suppose that (m, a, ω) is an expert-optimal equilibrium, thus $u^S(m, a \mid \psi) = \max\{u^S(a' \mid \psi) \mid a' \in \Delta(A)\}$. It is immediate that for any announcement (s, n, Θ) the deviation is not profitable as the equilibrium payoff is weakly greater, and condition (A1) is violated. Thus, no announcement is weakly credible to (m, a, ω), making it a strongly announcement-proof equilibrium.

Suppose another equilibria (m', a', ω') exists, and it induces a different equilibrium payoff. Then, by expert optimally of (m, a, ω), there exists a sender type ψ who strictly prefers to (m', a', ω'). Moreover, all other sender types weakly prefer (m, a, ω) over (m', a', ω').

We will show that (m, Ψ) forms a credible announcement strategy against the equilibrium (m', a', ω'). Formally, we show that announcement strategy (m, Ψ) and equilibrium (m', a', ω') satisfy conditions (A1) to (A4).

- (A1) is satisfied because all types $\psi' \in \Psi$ weakly prefer (m, a, ω) to (m', a', ω'), and type ψ strictly prefers (m, a, ω) to (m', a', ω').
- (A2) is satisfied vacuously because the set of non-deviant types is an empty set.
- (A3) is satisfied because all types get the highest possible payoff, and (m, a, ω) constitutes an equilibrium.
- (A4) is satisfied because expert-optimality implies that no type $\psi' \in \Psi$ strictly prefers another announcement strategy relative to (m, Ψ)

Thus, we conclude that (m', a', ω') is not an announcement-proof equilibrium. ■