Efficient Cheap Talk in Complex Environments*

Yunus C. Aybas[†] Steven Callander[‡]

November 1, 2023

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, so knowledge of the sender's preferred action may not reveal the receiver's preferred action. We identify an equilibrium in which the action is exactly what the sender would choose if she held full decision making authority. This contrasts with the inefficient equilibria of the canonical model of Crawford and Sobel (1982) in which the decision making environment is simpler. Thus, strategic communication is not only more favorable to the expert when the environment is complex, it is also more effective. We explore the implications of this result on the size and structure of the choice set, the decision making mechanism, and how these vary in the complexity of the decision making problem.

*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, or from car repair to business investment, experts are able to steer decisions toward their own interests even when those conflict with the decision makers' interests.[1] 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'."

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, if the decision maker knew the expert's preferred 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 real world. In practice, a doctor not only knows much more than her patient, but a single diagnosis, no matter how precise, may not reveal to the patient what her 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 acquires no leverage and would be better off if she could simply transfer her expertise to the decision maker for free.

In this paper we develop a model of expertise in complex environments in which the decision maker cannot infer his own optimal decision from the expert's preferred choice. Our expert's recommendation may not be "invertible" because there are many things she knows that the decision maker does not. Revisiting the classic sender-receiver game of Crawford and Sobel (1982) for such complex environments, we identify an equilibrium that delivers the outcome the expert would obtain were she to hold full decision making authority herself. This result shows how complex decision environments enhance the power of experts.

We illustrate and develop these ideas through a representation of one particular complex environment in which the possible mappings from actions to outcomes are paths of a Brownian motion, which form an infinite dimensional space. The correlation of the payoffs for different pairs of actions depends on the scale (or variance) of the Brownian motion, which

[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.

we use together with the drift (or expectation) of the Brownian motion to parameterize the complexity of the underlying decision environment. The expert knows the outcome of every action, observing the entire path, whereas the decision maker knows only the distribution of outcomes for each action and the correlation structure among them (the drift and scale parameters, and the outcome of a status quo action).

Experts are empowered in complex environments because an expert can make a recommendation that reveals precisely her most preferred action, yet does not reveal all of her information. In the Brownian model, if the expert recommends her most preferred action, that provides just one dimension of information about the infinitely many uncertain dimensions of the payoff function. Such precise but imperfect communication leaves the decision maker unsure as to his best response and more willing to accept the expert's recommendation.

A patient may know, for example, that his doctor is overly cautious, and that a prescription reflects the doctor's preferred dosage and not his own; yet he cannot infer from this the dosage that is best for him. We show that it is the expert's ability to use her information fully while simultaneously keeping some of it private that is the foundation of her power.[2] When revealing the expert's preferred action reveals all of her information, her power to affect outcomes is very much reduced.

Although this logic suggests that the expert will maximize her influence by revealing as little information as possible beyond her recommendation, we show this is not the case. We identify a strategy that allows the expert to shape beliefs beyond her recommendation in a way that systematically favors the recommendation itself. She is able to do this because in choosing one action to recommend, she is implicitly not choosing other actions, and this generates a spillover of information about the other actions that influences the decision maker's choice. Through this informational spillover, the expert is able to dissuade the decision maker from taking other actions and, in turn, this persuades him to accept a recommendation that he knows is the expert's ideal action and not his own, even when the expert's bias is large. Importantly, the logic of this result does not rely on risk aversion of the decision maker. The expert's power is such that she can convince the decision maker that her ideal action is in expectation also his ideal action.

Behavior in the equilibrium matches the many situations in practice where a decision maker acquiesces to an expert's recommendation. It captures a board of directors accepting unchanged a CEO's recommended strategy, a patient adopting a doctor's recommended treatment, or a homeowner following a realtor's recommendation to accept an offer. Although such behavior may appear as deference, or even a decision maker rubber-stamping 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 that the recommen-

[2] As Arrow (1963, p.965) pointed out, as a 'consequence of information inequality' between a physician and her patient, 'the patient must delegate to the physician much of his freedom of choice.'

dation serves the expert's interests and not his own. This result is also relevant to settings, such as politics, in which the formal delegation of authority plays a central role in models of expertise, yet is tenuous in practice (Gilligan and Krehbiel, 1987). Our result shows that the formal delegation of decision rights is not necessary for an expert to obtain Weber's (1958) "overtowering" position when the decision environment is complex.

The Brownian motion is but one example of a complex environment that allows us to illustrate the mechanics of efficient cheap talk. Building on the logic of the Brownian example, we identify a variety of natural settings in which efficient cheap talk is possible. The common element across all settings is that by communicating precisely yet imperfectly, a biased expert is able to induce the decision maker to accept a recommendation even though he knows that the recommendation serves the interests of the expert.

The equilibrium we identify establishes a possibility result for efficient cheap talk. We also prove that any equilibrium that is better for the receiver must necessarily be inefficient. Whether such an equilibrium exists, or whether, instead, the sender-optimal equilibrium is also receiver-optimal, remains an open question. In establishing our results, we use and adapt results from stochastic processes and provide several new analytic characterizations of distributions and moments for the Brownian motion and the Brownian meander. These tools may prove useful for a more general development of strategic communication in complex environments, including finding an answer to the above question.

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 illuminating to note 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 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.

The size of the state space relative to the action space distinguishes our approach from models of multidimensional cheap talk. Chakraborty and Harbaugh (2007, 2010) and Levy and Razin (2007) consider higher dimensional state spaces although they increase the action space commensurately and retain the assumption that expertise is perfectly invertible, precluding efficient cheap talk. Levy and Razin (2007) introduce correlation across the dimensions of choice and show how the spillover of one dimension of choice to the other can further constrain cheap talk, foreshadowing the importance of informational spillovers in our

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

model of complex environments.[4]

It is well known that a discrete action space can support efficient cheap talk, but only if the actions are sufficiently far apart. Indeed, CS's insight is to show how a continuous action space can be discretized endogenously to support cheap talk equilibria. CS's equilibria come at the cost of wasted information and the elimination of expert power. Our result shows that cheap talk equilibria are possible in continuous action spaces with expert power and without wasted information if the environment is sufficiently complex.

The invertibility of expertise can also be relaxed in other ways, for example, by assuming the decision maker is unsure of the expert's bias. Morgan and Stocken (2003) allow for the possibility that the interests of the expert and decision maker are aligned and show that communication remains inefficient if there is any positive probability of misalignment.[5] 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). These models minimally complexify the simple environment of CS as the expert now knows two pieces of information the decision maker does not. This empowers the expert but only to a limited degree. Equilibria remain of the partitional form and communication is inefficient. In Section 6 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 in complex environments and that this advantage can robustly support efficient cheap talk.

We differ from the burgeoning Bayesian persuasion literature in that we do not assume any commitment power for the sender (Kamenica and Gentzkow, 2011). A striking feature of our result is that in complex environments the expert can obtain her first-best even without commitment.

The Brownian motion has been used to represent the action-outcome mapping in a variety of applications.[6] Callander, Lambert and Matouschek (2021) analyze a model of verifiable information and show how the expert can obtain leverage by providing information in addition to a recommendation (the "referential" advice of their title). The verifiability of information is essential to their result. The information in the expert's recommendation persuades the decision maker and the separate referential information dissuades the decision maker from

[4] Deimen and Szalay (2019) study the acquisition of expertise in a one dimensional action space with two unknowns, focusing on which of the two pieces the sender acquires and whether a conflict of interest emerges.

[5] Aghion and Tirole (1997) consider a similar setting with a discrete action space. They obtain efficient cheap talk, although, as noted above, this derives from the discreteness of the action space rather than an unknown bias.

[6] Applications include search and experimentation (Callander, 2011; Garfagnini and Strulovici, 2016; Urgun and Yariv, 2021; Cetemen, Urgun and Yariv, 2021), "attributes" problems (Bardhi, 2022; Bardhi and Bobkova, 2021; Carnehl and Schneider, 2021), and industrial organization (Callander and Matouschek, 2022).

taking any other action. The surprising contribution of the present paper is to show that a similar logic can emerge using only cheap talk. Callander (2008) studies a model of cheap talk although he identifies only the polar and unique case in which equilibrium exists without dissuasion. The equilibrium he identifies is efficient and sender-optimal, although it exists only for small bias and relies on risk aversion of the decision maker. We show how an efficient equilibrium that does not rely on risk aversion can exist for any bias up to a known status quo outcome.

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 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, 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*

[7] Formally, the state space is the set of all continious functions with domain $\mathcal{A}$ and range $\mathbb{R}$.
[8] We point out the results that are special to the quadratic form where relevant.

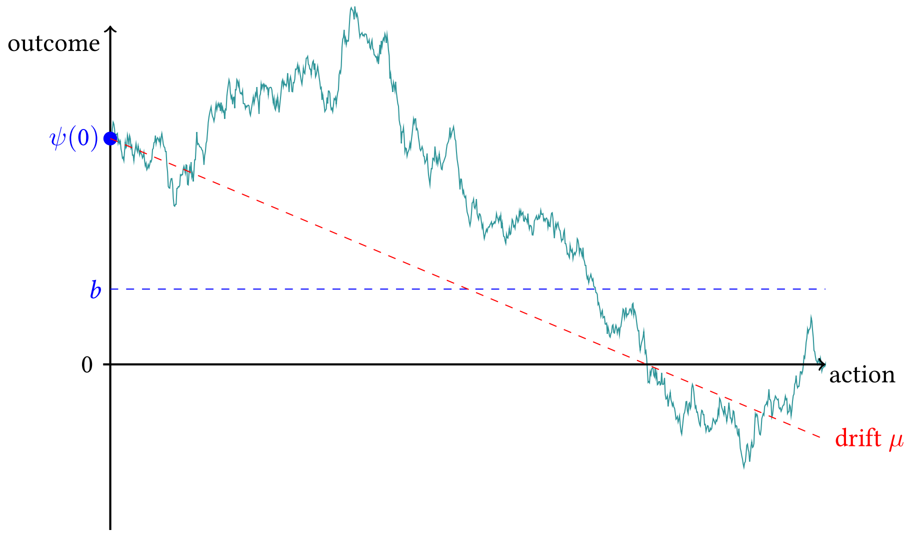

Figure 1: The Mapping From Actions to Outcomes

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. It 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. 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

[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. A possibility that the status quo may be more attractive to the receiver than the sender's ideal outcome adds a contingency to the receiver's decision making, although the logic of our equilibrium should still apply.

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

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 6.2 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 exercises of 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. The size of the action space is itself economically meaningful as in many applications it is in the control of the players, such as in models of delegation. We discuss connections to the delegation literature in Section 5.

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

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

The expected outcome is determined by the drift line, which by assumption is negative. This is depicted in Figure 2 by the red dashed line (see also Figure 1). Variance is increasing in the distance an action is from the status quo, capturing the idea that uncertainty is increasing

[11] Such a representation implies that a decision maker becomes completely informed—an expert—after observing 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.

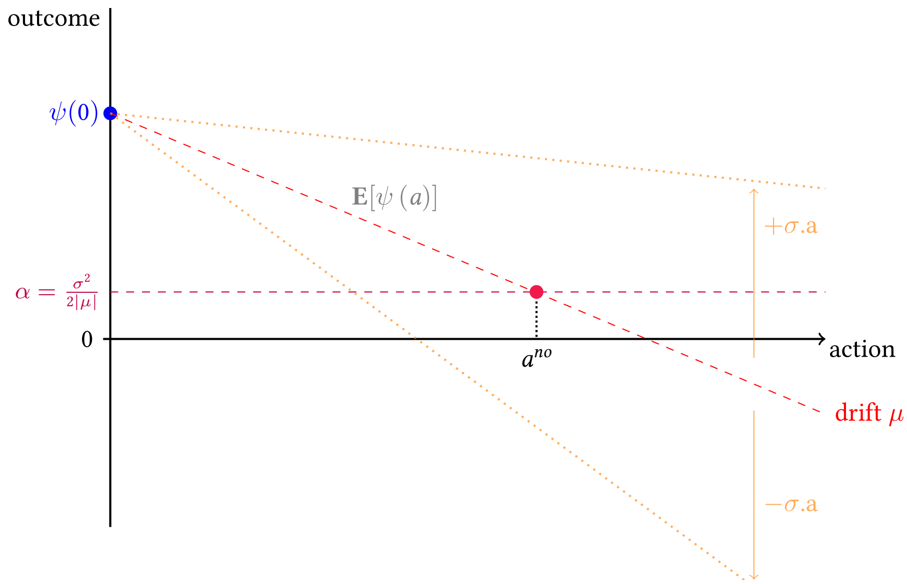

Figure 2: Beliefs in the Absence of an Expert

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 $\alpha = \frac{\sigma^2}{2|\mu|}$. We then have the following.

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

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

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

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

[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 matters at one other point; see footnote 24 on the comparative static of Proposition 1.

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 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 2. 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:*

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

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 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. In complex environments a distinction arises between direct and indirect informational spillover.[13] *Direct* spillover is what the receiver learns from the inference that the mapping passes through the point (r^*, b). This knowledge 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 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 the receiver being risk averse.

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

[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 the end of Section 4.5.

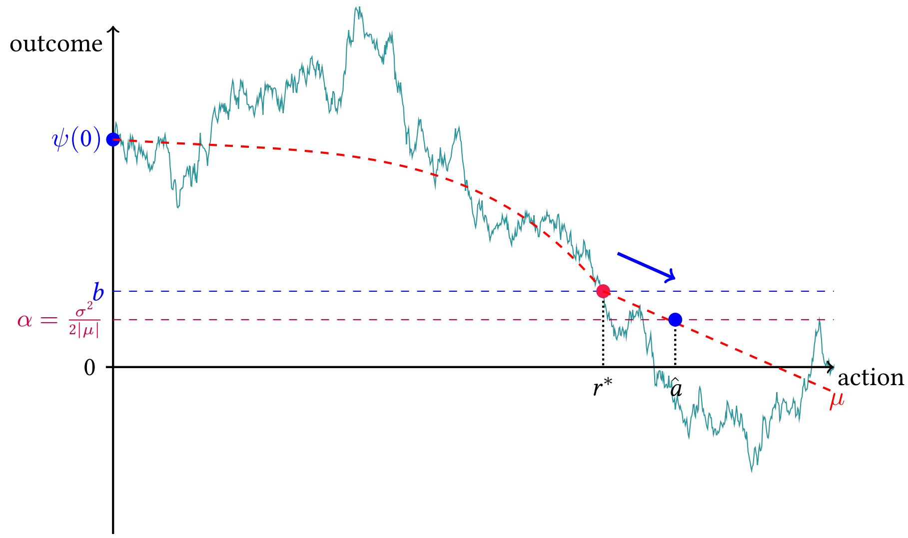

Figure 3: First-Point Strategy: Recommendation r^* and Receiver's Optimal Response $\hat{a}$.

inference from the recommendation, 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 informational 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 3 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.[15]

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

[15]This ability may hurt the receiver as an outcome of b with certainty dominates experimenting on his own if $b - \alpha$ is small, and may be better than any inefficient cheap talk equilibria that are possible.

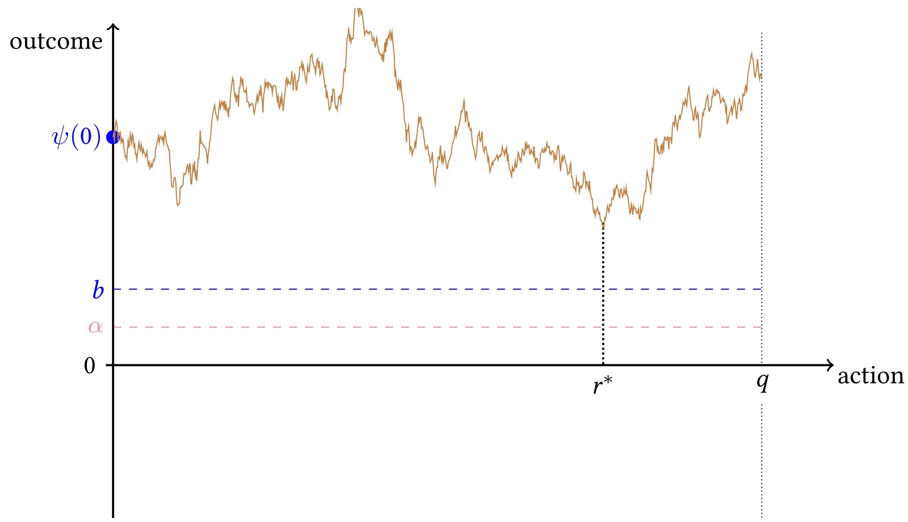

Figure 4: First-Point Strategy & Recommendation r^*: Outcome Above b

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 situation is depicted in Figure 4.

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.[16] 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 low levels of bias, 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 $\hat{a} = r^* + \frac{b-\alpha}{|\mu|}$ (by Lemma 1), and with complementary probability, the outcome is above b and the receiver's interests are aligned with the sender on the recommendation r^*.[17]

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 that he must not choose an intermediate action even though all are available.[18] 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, therefore, 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:*

[16] Indeed, here and in Section 6.2 we show this is neither necessary nor sufficient for efficient cheap talk.

[17] Where $\hat{a}$ is itself how the receiver resolves response uncertainty when the outcome of the recommendation is b and his beliefs are neutral.

[18] It is a special property of the Brownian environment that one of the receiver's possible best responses is also the sender's preferred action. In Section 6.2 we provide examples where this is not the case. We show also it is not necessary that the receiver is unsure of the outcome of the recommendation, as is the case with the Brownian motion.

(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 $q = 0$.[19] Interestingly, the upper bound on bias is independent of the complexity of the underlying process.[20] 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, as depicted in Figures 3 and 4, respectively.

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

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

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

formalized as the probability that the Brownian motion first hits b at action r^*.[21] 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).

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:

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

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

[21] 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 minimum location of a Brownian Motion. To highlight that these are probability densities in the formulas, we employ the notation "$\in dx$" rather than "= x."

the receiver's belief in Event =b conditional on recommendation r^* is:

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

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 5. 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 5 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.

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.

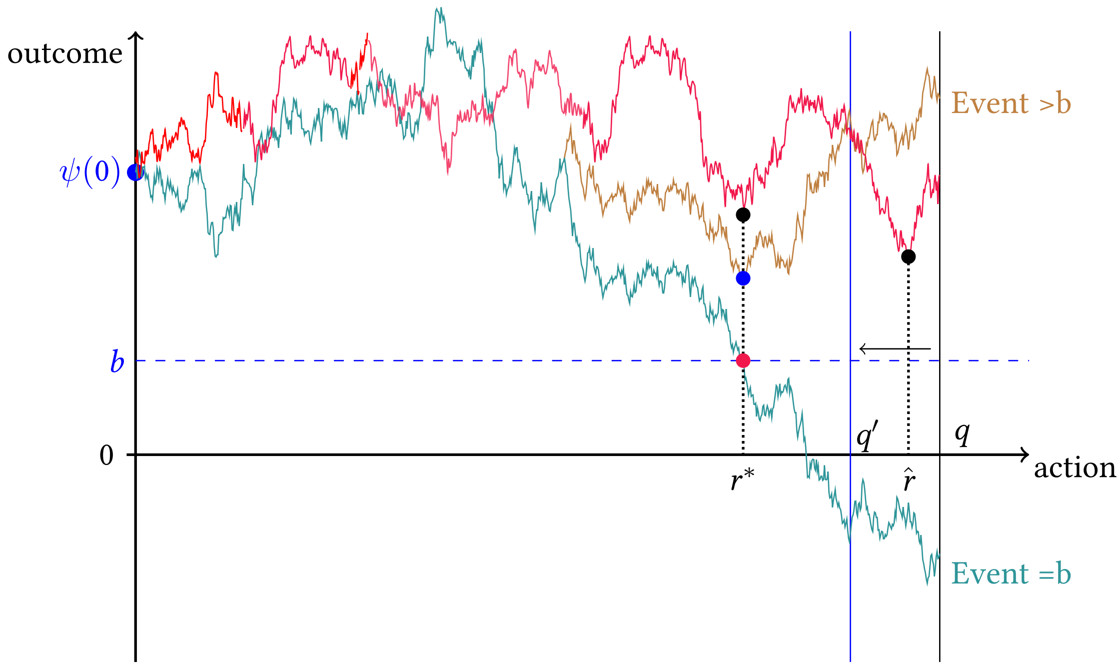

Figure 5: Paths that Induce r^* as $q \to q'$.

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.[22] 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

[22] The Brownian meander is generally studied with $\mu = 0$ and $\sigma = 1$. Recent work has extended the characterization to general μ for the distribution (Iafrate and Orsingher, 2020) and moments (Riedel, 2021). We extend both results to general μ, σ and q in the online appendix.

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 $[0, q_b^{max}]$. This creates the same potential common interest depicted in Figure 4. 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.[23] 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.[24]

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.*

Figure 6 plots $q_b^{\max}$ as a function of b for parameter values $\mu = -1$, $\sigma^2 = 1$, and $\psi(0) = 2$, such that $\alpha = \frac{1}{2}$. 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 that the path crosses b—and the sender and

[23] This does not imply that the sender needs to know the entire path. For example, the expert can merely monitor the running minimum of the path or observe the realization of the random variable $m^*(\psi)$. In the latter case, the expertise is of the same dimension as the action space i.e., the state space is isomorphic to the action space.

[24] 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.

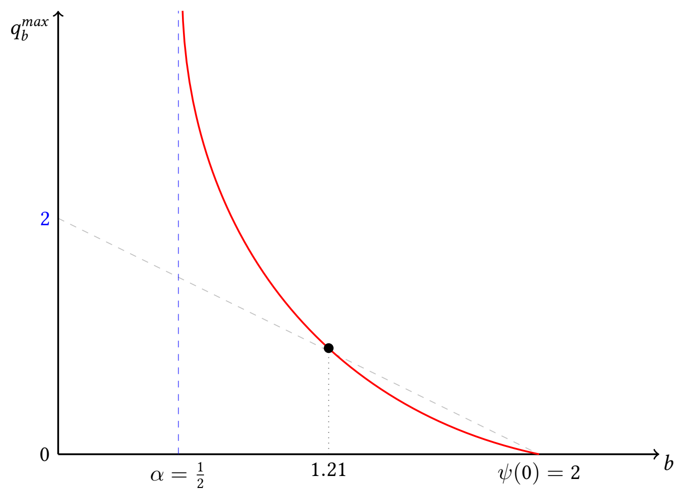

Figure 6: q_b^{max} for $\mu = -1, \sigma^2 = 1, \psi(0) = 2$, and $\alpha = \frac{\sigma^2}{2|\mu|} = \frac{1}{2}$.

receiver have opposing interests—with substantial probability.[25]

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.*

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

[25] At $b \approx 1.21$ the expected outcome of q_b^{max} equals b itself (the red dashed line). Thus, for bias less than this, the probability is greater than $\frac{1}{2}$ that the outcome of the mapping is below b at action q_b^{max}. For larger bias, q_b^{max} decreases at a slow rate such that this probability remains substantial.

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.[26] 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.

[26] 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 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.

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.*

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.[27] 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.

[27] Consider, for example, the kth-point strategy, for $k > 1$, and situations in which the path crosses b less than k times. To satisfy incentive compatibility, the sender must recommend an earlier crossing action. This possibility means that, in Event =b, the receiver's beliefs to the right of the recommendation are not neutral, with positive weight given to the possibility that the path is thereafter bounded above by b, as is the case for the last-point strategy.

5 Discussion

Actions to the Left of the Status Quo. Allowing actions to the left of the status quo opens a new possibility for the sender and makes the equilibrium easier to sustain. To accommodate this possibility, amend the first-point strategy so that it reads the mapping to left of the status quo and continues to the right only if the mapping has not crossed b.[28]

A recommendation to the left has a different impact on the receiver's beliefs. To the left of the recommendation, beliefs are neutral but because $\mu < 0$, the expected outcome is worse than the recommendation rather than better. Between the status quo and the recommendation, outcomes are bounded below by the outcome of the recommendation. Thus, even in Event =b, the receiver expects his interests to be aligned with the sender on this part of the action space and he does not want to override the recommendation locally. His only alternative is to choose an action to the right of the status quo, but this involves more variance than had the recommendation itself been to the right of the status quo.

Negative Bias. The direction of the expert's bias is immaterial in CS. This is not the case in complex environments with a known status quo outcome. Negative bias generates many of the same intuitions as does positive bias, although with an additional subtlety.

To see this, suppose bias is small and negative and the sender uses the first-point strategy. On an unbounded space, the outcome is almost surely b and close to the receiver's ideal. The most attractive deviations are now to the left of the recommendation in the region where the receiver knows the path must cross 0. The risk in overriding the recommendation depends on the size of the recommendation itself as this determines the slope of the Brownian bridge that forms between the recommendation and action 0. For a recommendation close to the status quo, the risk is small and the receiver will override the recommendation.[29] The equilibrium can be restored on a bounded space as then there is the possibility that the path doesn't cross 0 (and, therefore, is above b), such that the interests of the players are aligned.

Practical versus Theoretical Knowledge. Until now we have presumed the receiver's lack of knowledge is purely practical—he lacks knowledge of the mapping. It is also possible that the receiver lacks deeper *theoretical* knowledge of how the mapping was generated. In many situations the receiver lacks the underlying theoretical knowledge that generated the environment.[30] In the Brownian environment, this corresponds to knowledge of the drift and variance parameters. An expert who holds both a theoretical as well as a practical advantage over a decision maker more easily supports an efficient cheap talk equilibrium as

[28]Formally, we define the new strategy as $m^*(\psi) = \min\{\ |a| : a \in \arg\min_{a' \in \mathcal{A}} u^S(a', \psi)\}$.

[29]Reversing the sign of the drift and the bias does not restore the base model, even if the action space is expanded to the left. The key impact of negative bias is that the receiver's best response now lies on a Brownian bridge.

[30]This may correspond to Kuhn's (1962) distinction between technological and scientific knowledge.

the receiver, in his ignorance, is even less inclined to override the expert's recommendation. This opens up the interesting possibility of differences in the type of expertise, with some experts holding practical knowledge about the mapping and other experts possessing theoretical knowledge.The decision maker must then decide not only whether to consult an expert, but which type of expert he should listen to.

An Endogenous Action Space. In many contexts the action space is itself a choice. In drafting policy, for example, legislators often describe in broad terms the boundaries that bureaucratic rules must conform to. Kolotilin, Li and Li (2013) study "limited commitment" in which the receiver can commit ex ante to a limited set of actions and show in the simple environment of CS that this can improve the quality of communication and make both players better off. This result does not hold in complex environments if $q < q_b^{\max}$ as any further constraint on the action space only lowers the probability that the path crosses b, making at least the sender worse off.

An interesting possibility emerges when the first-point equilibrium does not exist. Then a restriction of the action space to $q < q_b^{\max}$ creates an environment in which efficient cheap talk can occur, potentially making both players better off. This complements the delegation literature as it implies the common practice of restricting an action space to an interval may occur even when commitment is "limited," and not only when all decision rights are handed over to an agent, as assumed in the delegation literature (Holmstrom, 1977, 1984; Melumad and Shibano, 1991; Alonso and Matouschek, 2008). This new rationale for interval restrictions resonates with those applications—such as legislative policymaking—where a principal's ability to commit to the delegation of decision rights is unclear.

Delegation versus Communication. This discussion leads us naturally to Dessein's (2002) question: Is the receiver better off delegating or communicating? In the simple environment of CS, the answer to this question involves a trade-off between the loss of control versus the loss of information, and Dessein (2002) shows this generally favors delegation. The answer is different in complex environments. Communication is efficient in the first-point equilibrium and no information is lost. This favors communication. Countering this benefit is that the sender obtains leverage and the receiver loses control even when communicating. When the first-point equilibrium exists on an unbounded space, these forces exactly balance such that delegation and communication are equivalent.

A different conclusion emerges for larger bias when the first-point equilibrium requires a bounded action space. As we noted earlier, a striking feature of the first-point equilibrium is that both players receive a *worse* outcome when their action preferences align than when they are opposed. Counter-intuitively, alignment is created by making both players worse off. Thus, when delegating, the receiver is better off not bounding the action space and instead giving the sender full freedom of choice.[31] Delegating dominates communicating for larger bias (via the first-point equilibrium) despite the fact that the receiver's loss of control increases in the sender's bias.[32]

6 Beyond the Brownian Motion

The Brownian motion is an example of a complex environment that provides particularly clear insight into the mechanism underlying efficient cheap talk. It is not the only environment that can support efficient cheap talk, however. To understand what is required, we develop several increasingly demanding requirements on the receiver's response. The most demanding is the receiver's incentive compatibility condition that he accepts the recommendation. This condition is both necessary and sufficient for equilibrium (given the sender's incentive compatibility is satisfied trivially whenever she is recommending her optimal action). By itself, however, the receiver's incentive compatibility condition is not particularly illuminating. To better understand why many environments violate it, 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 efficient cheap talk equilibria.

6.1 Ingredients for Efficient Cheap Talk

Throughout the paper we have emphasized the notion of partial invertibility. This is the most basic necessary condition 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 $m(\cdot)$.*[33]

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 he chooses his best action rather than the sender's recommendation.

[31] This presumes restrictions on the action space of the form $[0, q]$.

[32] Other communication protocols may give the receiver more control and facilitate communication when the first-point equilibrium does not exist. One interesting possibility is to apply to complex environments the mediation and arbitration protocols that Goltsman et al. (2009) study in the simple environments of CS.

[33] The definitions of perfectly-invertible and non-invertible follow naturally by negating the relevant premise i.e. $m^{-1}(r) = \Psi$ for non-invertible and $|m^{-1}(r)| = 1$ for perfectly invertible.

Partial invertibility is necessary but not sufficient for efficient cheap talk when the sender is biased. 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 he 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 an efficient equilibrium. This is illustrated by Morgan and Stocken's (2003) model of unknown bias. If the sender recommends her most-preferred action r^*, the receiver does not know if the sender's bias is 0 and r^* is also his most-preferred action, or whether bias is b and his best choice is $r^* - b$.[34] The efficient strategy satisfies partial invertibility and also response uncertainty, but it is not an equilibrium. That is because the receiver's best response to this uncertainty is not the recommendation itself but rather a compromise action between r^* and $r^* - b$. This leads Morgan and Stocken (2003) to the conclusion that efficient cheap talk fails in their model whenever there is any probability of positive bias. The challenge of efficient cheap talk, as noted in Section 4, is that the receiver's optimal compromise action must be the recommendation itself. The receiver's incentive compatibility condition is equivalent, therefore, to a willingness to accept the sender's recommendation. For sender-optimal strategies, the receiver's incentive compatibility is necessary and sufficient for equilibrium. Denote by $a(r)$ the receiver's best response to a recommendation r.[35]

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 $a(r) = r$.*

Incentive compatibility for the receiver is a demanding requirement. As we saw in the Brownian motion, it can be satisfied in complex environments even with positive sender bias. 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.[36]

[34]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.

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

[36]The reader will have noted that we state Definitions 2-4 for arbitrary strategies and not just efficient

6.2 Other Complex Environments

We present several environments that satisfy the three requirements and support efficient 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.[37]

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 situation is depicted in Figure 7 for recommendation r^*.

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, 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.

strategies. Indeed, they define the requirements for any cheap talk equilibrium with 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.)

[37] Jumps downward open the possibility that a recommendation delivers an outcome below b. This complicate the analysis though efficient cheap talk can still be sustained.

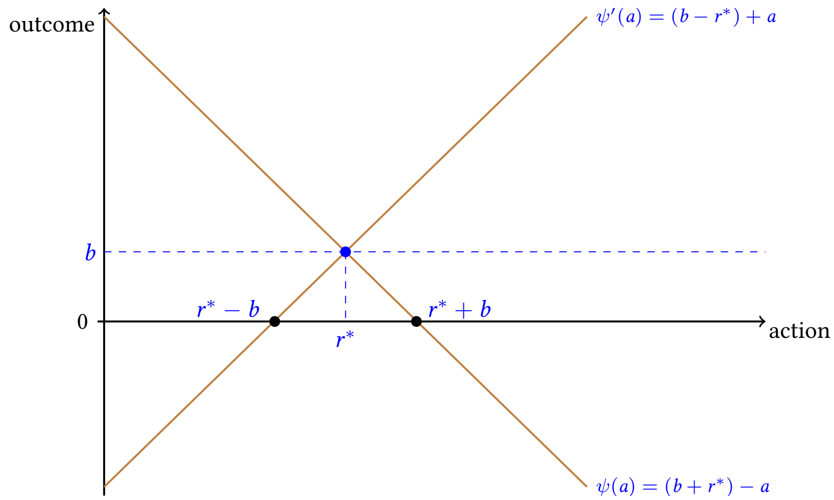

Figure 7: Directional Uncertainty

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.[38]

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.[39] 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 $n \in \mathbb{Z}^+$,

[38]Similar examples can be constructed with unknown bias although an additional difficulty emerges in that formulation. In those models the slope is known to be +1 and with bias either $\pm b$, the receiver's best response to an efficient recommendation is either $r^* \pm b$. The additional complication is that on a bounded state space, there will be recommendations made by a sender with bias $+b$ that are not made by a sender with bias $-b$, and vice versa, and partial invertibility fails. This property motivates bias that depends on the state and the 'globally outward' bias condition of Gordon (2010).

[39]Formally, there exists $a, a' \in \hat{a}({m^*}^{-1}(r^*))$ such that $a < r^* < a'$.

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.

Figure 8 depicts one possible OU path. r^* denotes the recommendation from the sender using the first-point strategy. The OU process differs from the Brownian motion in that the process is mean-reverting. Thus, the receiver expects actions to the right of r^* to deliver outcomes closer to $\psi(0)$ rather than to 0 as did the Brownian motion, as depicted by the red dashed line in the figure. 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.

More Knowledge About the World: We have assumed that the receiver begins with knowledge of only a single point in the mapping. In many situations the receiver may know more, whether from the history of play or his own research. To see how this can matter, suppose the receiver knows the outcome of a second action, which for simplicity we suppose is the right-most action q on the bounded action space $[0, q]$, and that $\psi(q) > b$.

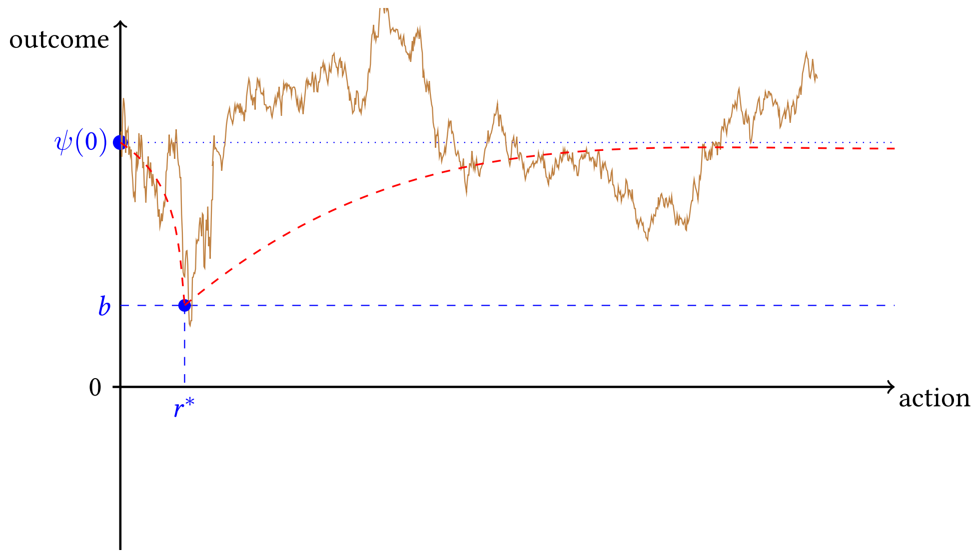

Figure 8: Mean Reverting Ornstein-Uhlenbeck Process With Mean $\psi(0)$

This environment closely resembles the OU process above. Following the first-point strategy, a recommendation r^* creates a Brownian bridge between r^* and q. Although the receiver believes with probability one that a better action for him lies somewhere to the right, he expects the outcome to increase in either direction from the recommendation, and his best response is to follow the recommendation regardless of whether the outcome is at b or above. As it did in the OU environment, the first-point equilibrium exists in this environment for all biases between 0 and the minimum of $\psi(0)$ and $\psi(q)$.

7 Conclusion

In this paper we have shown how expert power can derive from the complexity of the underlying environment. This power emerges not despite the complexity but because of it. Complexity allows the expert to communicate precisely yet imperfectly, obviating the decision maker's ability to appropriate her expertise for his own ends. Communication in complex environments takes a particularly simple form: the expert recommends her most preferred action and the receiver rubber-stamps it. Thus, strategic communication not only favors the expert more when the environment is complex, it can also be more efficient.

Our analysis has limitations but also opens up new questions. The primary limitation is that we characterize only a single equilibrium. Identifying additional equilibria is an obvious direction for future work. A more complete understanding of the set of equilibria may help create a bridge between the efficient equilibria in complex environments and the inefficient equilibria of CS in simple environments. A productive step is to add uncertainty over the status quo to the Brownian environment (see Remark 1). The simple environment of CS would then correspond to the limit of the sequence of environments in which the scale of the Brownian motion approaches zero. Identifying the set of equilibria along the continuum between simple and complex environments offers the promise of even deeper insight into the nature of strategic communication.

Complex environments open up new questions about the role of institutions. In simple environments, Gilligan and Krehbiel (1987) show how institutions can rebalance power away from the receiver to the sender. In complex environments, the opposite incentive takes hold. As the sender is better able to protect her information, the receiver may design an institution to weaken that grip and rebalance power toward himself. He may, for example, structure the institution to make it difficult for the sender to learn the full outcome mapping, or that restricts the sender's ability to communicate in some way.

These ideas only scratch the surface of the questions that open up in a world of informational richness and complex expertise. Exploring the possibilities, and embedding complex expertise into 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.

Alonso, Ricardo and Niko Matouschek. 2008. "Optimal delegation." *The Review of Economic Studies* 75(1):259–293.

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 pp. 217–241.

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

Bardhi, Arjada and Nina Bobkova. 2021. "Local Evidence and Diversity in Minipublics." Working Paper.

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

Callander, Steven. 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.

Callander, Steven and Niko Matouschek. 2022. "The Novelty of Innovation: Competition, Disruption, and Antitrust Policy." *Management Science* 68(1):37–51.

Carnehl, Christoph and Johannes Schneider. 2021. "A Quest for Knowledge." *arXiv preprint arXiv:2102.13434* .

Cetemen, Doruk, Can Urgun and Leeat Yariv. 2021. "Collective Progress: Dynamics of Exit Waves." *Journal of Political Economy* p. forthcoming.

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

Chakraborty, Archishman and Rick Harbaugh. 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.

Deimen, Inga and Dezső Szalay. 2019. "Delegated expertise, authority, and communication." *American Economic Review* 109(4):1349–1374.

Dessein, Wouter. 2002. "Authority and Communication in Organizations." *Review of Economic Studies* 69:811–838.

Durrett, Richard T, Donald L Iglehart and Douglas R Miller. 1977. "Weak convergence to Brownian meander and Brownian excursion." *The Annals of Probability* pp. 117–129.

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.

Goltsman, Maria, Johannes Hörner, Gregory Pavlov and Francesco Squintani. 2009. "Mediation, arbitration and negotiation." *Journal of Economic Theory* 144(4):1397–1420.

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.

Holmstrom, Bengt. 1977. "On Incentives and Control in Organizations." Ph.D. dissertation, Graduate School of Business, Stanford University.

Holmstrom, Bengt. 1984. On the Theory of Delegation. In *Bayesian Models in Economic Theory*, ed. M. Boyer and R. Kihlstrom. Elsevier.

Iafrate, Francesco and Enzo Orsingher. 2020. "Some results on the Brownian meander with drift." *Journal of Theoretical Probability* 33(2):1034–1060.

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.

Kolotilin, Anton, Hao Li and Wei Li. 2013. "Optimal Limited Authority for Principal." *Journal of Economic Theory* 148:2344–2382.

Kuhn, Thomas. 1962. Comment. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton, NJ: Princeton University Press pp. 379–384.

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.

Levy, Gilat and Ronny Razin. 2007. "On the limits of communication in multidimensional cheap talk: a comment." *Econometrica* 75(3):885–893.

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

Melumad, Nahum D and Toshiyuki Shibano. 1991. "Communication in settings with no transfers." *The RAND Journal of Economics* pp. 173–198.

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* pp. 183–203.

Riedel, Kurt. 2021. "Mean and variance of Brownian motion with given final value, maximum and argmax." *Stochastic Models* 37(4):679–698.

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.

Shreve, Steven E. 2004. *Stochastic calculus for finance II: Continuous-time models*. Vol. 11 Springer.

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

Sobel, Joel. 2012. Complexity versus conflict in Communication. In *2012 46th Annual Conference on Information Sciences and Systems (CISS)*. pp. 1–6.

Urgun, Can and Leeat Yariv. 2021. "Constrained retrospective search." *AEA Papers and Proceedings* 111:549–53.

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

Appendix A

A.1 Formal Details of the Environment

We begin by completing formal details of the environment that were omitted from the text.

States and Beliefs: The state of the world $\psi(\cdot)$ is a transformation of the Wiener process $W(\cdot)$ with parameters $\mu, \psi_0 \in \mathbb{R}$ and $\sigma^2 \in \mathbb{R}_+$ given by $\psi(a) = \psi_0 + \mu a + \sigma W(a)$. We denote the space of all outcomes as Ψ. 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 $(\mathcal{W}, \mathcal{B}(\mathcal{W}))$.[40] 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 $\omega(\cdot \mid r \in m(\psi)) : \mathcal{B}(\mathcal{C}[0,1]) \times \mathcal{M} \to [0,1]$.[41] Equilibrium requires that the following hold:

1. $\omega(\psi \mid r \in m(\psi))$ is obtained from the prior using Bayes's rule whenever possible,

2. $a(r) \in \arg\max_{a' \in \mathcal{A}} \mathbb{E}[u_R(a', \psi) \mid \omega(\psi \mid r \in m(\psi))]$ for every $r \in \mathcal{M}$,

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

The receiver's beliefs in equilibrium are conditional distributions of the drifting Brownian motion $\psi(\cdot)$ conditional on $\psi \in m^{-1}(r)$. Details about the relevant conditional distributions and their derivation are provided in the proofs and online appendix as required.

A.2 Properties of the Mapping

Throughout the paper, we use results on random variables of the drifting Brownian motion $\psi(a)$. The mapping $\psi(a)$ is uniquely characterized by three properties:

1. $\psi(0) = \psi_0$ with probability 1.

2. The increments $\psi(a + a') - \psi(a)$ are normally distributed with mean $\psi_0 + \mu a'$, and variance $\sigma^2 a'$ for every $a, a' \in \mathbb{R}_+$ and constants $\mu \in \mathbb{R}_+$ and $\sigma^2 \in \mathbb{R}_+$.

3. The distributions of the increments are stationary and independent.

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

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

We use the third property frequently and refer to it as the *stationary-independent increments property*. Throughout the proofs, we use two variants of the process $\psi(\cdot)$ in our analysis. First is the drifting Brownian motion $X(a') = \mu a' + \sigma W(a')$.[42] Note that, this process is identically distributed as the a' increment of the outcome map $\psi(a)$ i.e. $\psi(a+a') - \psi(a)$. The second variant is the Brownian meander.[43] This is the process in which the Brownian motion is conditioned to remain above its starting value over an interval of length q. Formally, it is $M(a, q) := \{X(a) \mid X(a') \geq 0 \ \ \forall a' \in [0, q]\}$.

We use the well studied random variable *hitting time* of a Brownian Motion, and its variants frequently in our analysis.

1. First hitting action (time) of outcome $x \in \mathbb{R}$: $\tau(x) := \inf\{a \in \mathbb{R} \mid \psi(a) = x\}$.
2. Infimum over the interval $[0, q] : \iota(q) := \inf\{\psi(a) \mid a \in [0, q]\}$.
3. First hitting action (time) of the minimum over $[0, q]$: $\tau_\iota(q) := \tau(\iota(q))$.

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

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

This formulation allows us to use properties of these random variables to study the receiver's conditional beliefs when facing the first-point strategy.

A.3 Proofs for Results in the Text

Throughout the proofs, we drop the ψ argument from $u_R(a, \psi)$ and $u_S(a, \psi)$ for conciseness, and write it as $u_R(a)$ and $u_S(a)$. 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

[42] For a more detailed discussion, the reader is referred to: Chapter 1 of Harrison (2013) and Chapter 3 of Shreve (2004).

[43] Durrett, Iglehart and Miller (1977), Iafrate and Orsingher (2020) and Riedel (2021) provide the fundamental results for our application of the Brownian meander. We provide extensions of their results in the online appendix.

expressions of certain distributions. The proofs of these properties and the derivation of the expressions are provided separately in the online appendix.

Proof of Lemma 1. By the mean-variance representation of quadratic utility, the receiver's expected utility is:

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

The first and second order conditions for optimality are:

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

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: Those 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

$\{\psi \in \Psi \mid \tau_\iota(q') = r^*, \iota(q') > b, \tau_\iota(q) > q'\}$.[44] Thus, we can write the expected change in payoff for the receiver taking action $\hat{a} \in [0, q']$ when the recommendation is $r^* \in [0, q']$, again by the law of total probability:

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

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 as:

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

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

[44] Note that, whenever $\psi(\cdot)$ attains a minimum greater than b over $[0, q]$ at r^*, the same path also attains a minimum greater than b over $[0, q']$ at r^*.

equilibrium has full support over $\mathcal{A}$. Let any off-path recommendation $r' \notin \mathcal{A}$ be interpreted as an on-path message, say $r'' = 0$, and generating the same beliefs. 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 both hold:

i. $\mathrm{Var}[\psi(a' + r^*) \mid m_q^*(\psi) = r^*] \geq \mathrm{Var}[\psi(r^*) \mid m_q^*(\psi) = r^*]$.

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 (standard) Brownian motion with initial point 0, drift μ and scale σ, and $M(\cdot, k)$ is the corresponding Brownian meander of length k. By the stationary independent-increments property of the Brownian motion, 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:

1. $\mathbb{E}[\psi(a' + r^*) - \psi(r^*) \mid \psi(r^*) = b] = \mathbb{E}[X(\cdot)] = \mu a'$.

2. $\mathbb{E}[\psi(a' + r^*) - \psi(r^*) \mid \tau_\iota(q) = r^*, \iota(q) > b] = \mathbb{E}[X(a') \mid \min\{X(a'') : a'' \leq q - r^*\} > 0]$ $= \mathbb{E}[M(a', q - r^*)]$.

We now 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 $r^* \leq \delta_\varepsilon$.[45] Similarly, in the online appendix

[45] This result can be also shown by using the property that Lèvy processes are continuous in probability, so we have that $\forall \varepsilon > 0$ it holds that $\lim_{a \to 0} P(|\psi(a) - \psi(0)| > \varepsilon) = \lim_{a \to 0} P(|\psi(a)| > \varepsilon) = 0$. ψ is continuous in probability if for any $\varepsilon > 0$ and $a \geq 0$ it holds that $\lim_{h \to 0} P(|\psi(a + h) - \psi(a)| > \varepsilon) = 0$.

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:

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

and there exists some $\tilde{a},\tilde{r}\in\left[0,q^b_{\max}\right]$ with $\tilde{a}=\tilde{r}+a'$ and $a'>0$ such that this holds with equality by the maximality of $q_b^{\max}$.[46] As usual, we can write this as:

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

Identically to Lemma 4, we can write this in terms of increments given by Brownian motion $X(\cdot)$ and Brownian meander $M(\cdot,\cdot)$. Thus, we have that:

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

[46] It directly follows that $\tilde{a}>\tilde{r}$, as for any $\tilde{a}$ smaller the value is obviously negative, as discussed before. Thus, we say $\tilde{a}=\tilde{r}+a'$ for some $a'>0$.

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. In order to prove this claim, we suppose that $u_R(\cdot)$ satisfies the condition:

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

This condition can be interpreted as: For the recommendation and deviation pair $(\tilde{r}, \tilde{a})$ where the receiver is indifferent, the percentage increase in the benefit of deviation is higher than the percentage increase in the cost of deviation.[47] 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:

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

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

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

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

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.[48]

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

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

[47] This condition is not necessary. The necessary and sufficient condition can be obtained by incorporating terms (6) and (7) to the inequality. We focus on the sufficient condition for a concise statement for our result.

[48] More precisely, by the definition of a Brownian meander, $M(a', \cdot) > 0$ 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 $\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] > 0$.

and only if the following holds.[49]

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

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 $q_b^{\max}$.[50] 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 directly follows that $a'(r^*) \to r^*$ and $\psi(a') \to b$.

Moreover, using the online appendix Corollary B.1, we show that $\lim_{a' \to r*} \frac{\partial}{\partial a'}\mathbb{E}[\psi(a') - \psi(r^*) \mid \tau_\iota(q) = r^*, \iota(q) > b] = \infty$. 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 that $q_b^{\max} \to \infty$.

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, $a(m(\Psi)) = \mathcal{A}$.

[49]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.

[50]Note that, by weak-concavity and unique maximization at 0, we have that $\mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q-\tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right] < 0$ and $\frac{\partial}{\partial b}\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] > 0$. Thus, rearranging replacing the RHS of (10) and rearranging we get (5).

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 Corollary 1. Consider the first-point strategies $m_q^*(\cdot)$ and $m_{q'}^*(\cdot)$ for games with action spaces $\mathcal{A} = [0, q]$ and $\mathcal{A}' = [0, q']$ with $q < q' \leq q_{\max}$. By definition of a first-point strategy, for every $\psi \in \Psi$ it holds that $\psi(m_q^*(\psi)) \geq \psi(m_{q'}^*(\psi)) \geq b$, where the set of paths that this holds with equality has measure zero under $\omega(\cdot)$. The conclusion follows immediately. ■

Proof of Corollary 2. Similarly, take any realized path $\psi(\cdot)$. Consider the first-point strategies $m_b^*(\cdot)$ and $m_{b'}^*(\cdot)$ for games with bias $b < b'$ such that the first-point equilibrium exists. By definition of a first-point strategy and continuity of the Brownian path:

1. If $\psi(m_b^*(\psi)) = b$, then $\psi(m_{b'}^*(\psi)) = b'$.
2. If $b' > \psi(m_b^*(\psi)) > b$, then $\psi(m_{b'}^*(\psi)) = b'$.
3. If $\psi(m_b^*(\psi)) > b'$, then $\psi(m_{b'}^*(\psi)) = \psi(m_{b'}^*(\psi)) > b'$.

Thus, for all paths $\psi(m_{b'}^*(\psi)) - b' < \psi(m_b^*(\psi)) - b$ with the inequality is strict for a measurable set of paths. The opposite is true relative to the receiver's ideal outcome of 0. The conclusion follows. ■

Proof of Corollary 3. By construction, we have that:

$$\mathbb{E}[\psi(r^*) \mid m^*(\psi) = r^*] = \mathbb{P}(\iota(q) > b)\ \mathbb{E}[\iota(q) \mid \iota(q) > b] + \mathbb{P}(\iota(q) \leq b)\ b$$

We have that $\psi(a) = \mu a + \sigma W(a)$ where $W(a)$ is the Wiener Process. It is a well-established result that $W(a) < 0$ for some a with probability one. Thus, with probability one there is an a such that $\psi(a) < \mu a$. As $\sigma \to \infty$ we have that $\psi(a) < b$ for some $a \in [0, q]$ with probability 1. It follows that $\sigma \to \infty$, $\mathbb{P}(\iota(q) \leq b) \to 1$ and the result holds. ■

Appendix B Other Complex Environments

In this section we provide the formal details for the environments discussed in Section 6.2. Unless amended otherwise, the details of the model are as in the main text.

Minimal Complexity:
Let the state space be $\Psi = \mathbb{R} \times \{-1, 1\}$ such that for $(w, z) \in \Psi$:

$$\psi(a \mid w, z) = b + z(a - w).$$

The receiver has prior belief given by $\omega((w,z))$ over Ψ. Note there is no known status quo point. The sender follows the first-point strategy (the optimal action is now unique and the 'first' modifier moot) $m^* : \Psi \to \mathbb{R}$.

For the receiver, the set of states consistent with an r^* are:

$$m^{*-1}(r^*) = \{(r^*, -1), (r^*, 1)\}.$$

$m^{*-1}(r^*)$ is not single valued for any $r^* \in \mathbb{R}$, thus $m^*(\cdot)$ satisfies partial invertibility. Sender strategy $m^*(\cdot)$ satisfies response uncertainty, as the set of optimal responses to the states that are consistent with message r^* are given by:

$$a(m^{*-1}(r^*)) = \{r^* - b, r^* + b\}$$

To check receiver incentive compatibility, consider a deviation $a' \in \mathcal{A}$ given recommendation r^*. The receiver's conditional beliefs are $\omega((r^*, 1) \mid m(\psi) = r^*)$ and $\omega(r^*, -1 \mid m^*(\psi) = r^*)$ where $\omega(r^*, 1 \mid m^*(\psi) = r^*) + \omega(m^*, -1 \mid m^*(\psi) = r^*) = 1$.

$$\begin{aligned}
\mathbb{E}(a \mid m^*(\psi) = r^*) &= -\omega(m^*, 1 \mid m^*(\psi) = r^*)(b + (a - r^*))^2 - \omega(r^*, -1 \mid m^*(\psi) = r^*)(b - (a - r^*))^2 \\
&= -b^2 - (a - r^*)^2 - 2b(a - r^*)(\omega(r^*, 1 \mid r^*) - 1) \\
\frac{d}{da}\mathbb{E}(a \mid m^*(\psi) = r^*) &= -2 \mid a - r^* \mid -2b(2\omega(r^*, 1 \mid m^*(\psi) = r^*) - 1) \\
\frac{d^2}{da^2}\mathbb{E}(a \mid m^*(\psi) = r^*) &= -2
\end{aligned}$$

The first order condition is satisfied if and only if $\omega(m^*, 1 \mid m^*(\psi) = r^*) = \frac{1}{2}$.

Sender-Receiver Misalignment without Directional Uncertainty:
Let the action and message spaces be the set of positive integers, $\mathbb{Z}_+$. The state space is $\Psi = \mathbb{Z}_+ \cup \{1, 2\}$ such that for $(w, z) \in \Psi$:

$$\psi(a \mid w, z) = \begin{cases} b & \text{if } a = w \\ 0 & \text{if } a = w + z \\ 100b & \text{if } a \notin \{w, w + z\} \end{cases}$$

The receiver has beliefs prior belief given by $\omega((w,z))$ over Ψ. Note there is no known status quo point. The sender follows the first-point strategy (the optimal action is now unique and the 'first' modifier moot). $m^* : \Psi \to \mathbb{R}$.

For the receiver, the set of states consistent with a reccomendation r^* are:

$$m^{*-1}(r^*) = \{(r^*, 1), (r^*, 2)\}.$$

$m^{*-1}(r^*)$ is not single valued for any $r^* \in \mathbb{R}$. Thus $m^*(\cdot)$ satisfies partial invertibility. Sender strategy $m^*(\cdot)$ satisfies response uncertainty, as the set of optimal responses to the states that are consistent with recommendation r^* are given by:

$$a(m^{*-1}(r^*)) = \{r^* + 1, r^* + 2\}$$

To check receiver incentive compatibility, consider a deviation $a' \in \mathcal{A}$ given recommendation r^*. The receiver's conditional beliefs are $\omega((r^*, 1) \mid m^*(\psi) = r^*)$ and $\omega((r^*, 2) \mid m^*(\psi) = r^*)$ where $\omega(r^*, 1 \mid m^*(\psi) = r^*) + \omega(r^*, 2 \mid m^*(\psi) = r^*) = 1$.

$$\mathbb{E}(a \mid m^*) = \begin{cases} -b^2 & \text{if } a = r^* \\ -\omega(r^*, 2 \mid m^*(\psi) = r^*)10000b^2 & \text{if } a = r^* + 1 \\ -\omega(r^*, 1 \mid m^*(\psi) = r^*)10000b^2 & \text{if } a = r^* + 2 \\ -10000b^2 & \text{if } a \notin \{r^*, r^* + 1, r^* + 2\} \end{cases}$$

It is optimal for the receiver to follow the recommendation as long as $\omega(r^*, 1 \mid r^*(\psi) = r^*)$ is not too close to 0 or 1.

Local Uncertainty:

The space of outcome maps Ψ is the paths of Ornstein-Uhlenbeck process. Formally, Ψ is the set of solutions to the following stochastic differential equation where $W(a)$ is the Wiener process:

$$d\psi(a) = -\kappa\left(\psi(0) - \psi(a)\right) da + \sigma dW(a)$$

For this process κ is the mean-reversion coefficient, and σ is the constant volatility term. This environment has the same state space as the Brownian environment, differing only in how the states are translated into outcomes via the outcome mappings.

Partial invertibility is satisfied under the first-point strategy as, just like the Brownian Motion, there are infinitely many paths $\psi(\cdot)$ of the Ornstein-Uhlenbeck process consistent with the message $m^*(\psi) = r^*$. Moreover, for every action $a \in \mathbb{R}_{++}$ there exists a realization of ψ such that $\psi(a) \in \arg\max_a -\psi(a)^2$, and the response uncertainty is also satisfied.

Consider the recommendation r^* and a deviation $a \in \mathbb{R}$. Deviations to $a < r^*$ are worse for the receiver as, by the first-point strategy and the continuity of OU process, $\psi(a) > b$ for every $a < r^*$ with certainty. For deviations $a > r^*$, the expected outcome and variance are, recalling that $\psi(0)$ is the mean of the process:

$$\mathbb{E}[\psi(a) \mid m^*(\psi) = r^*] = \psi(0) - (\psi(0) - \psi(r^*))\exp(-\kappa(a - r^*))$$
$$\text{Var}(\psi(a) \mid m^*(\psi) = r^*) = \frac{\sigma^2}{2\kappa}(1 - \exp[-2\kappa(a - r^*)])$$

As $\exp(-\kappa(a - r^*)) < 1$, the expected outcome is weakly greater than $m^*(\psi)$ for $a > r^*$ whenever $\psi(r^*) \leq \psi(0)$, which must be true given the first-point strategy. As variance is positive, it is optimal for the receiver to accept the recommendation. (Note that this argument holds even if $\psi(r^*) \in (b, \psi(0)]$).

More Knowledge About the World:
Set Ψ is the set of all paths of a Brownian Motion with diffusion σ, where $\psi(\cdot)$ is conditioned to satisfy fixed $\psi(0) > b$ and $\psi(q) > b$. The underlying state space is the same as in the Brownian motion (and OU) environments. For the first-point strategy and recommendation r^*, all actions $a < r^*$ are dominated by the recommendation itself. For $a > r^*$, the receiver's beliefs are given by:

$$\begin{aligned}\mathbb{E}[\psi(a)] &= \psi(r^*) + \frac{\psi(q) - \psi(r^*)}{q - r^*}(a - r^*) \\ &> \psi(r^*) \text{ as } \psi(q) > \psi(r^*) \\ \text{Var}(\psi(a)) &= \sigma^2 \frac{(q-a)(a-r^*)}{q - r^*}.\end{aligned}$$

The first-point strategy satisfies partial invertibility, response uncertainty, and receiver incentive compatibility are analogously to the Brownian motion and Ornstein-Uhlenbeck environments.