Efficient Cheap Talk in Complex Environments*

Yunus C. Aybas† Steven Callander‡

May 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 the classic sender-receiver game when the environment is complex. We identify an efficient cheap talk equilibrium that is sender-optimal. In fact, the equilibrium action is exactly what the sender would choose were she to hold 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 not only favors the expert more when the environment is complex, it can also be more efficient.

*For helpful comments, we thank Anton Kolotilin, Johannes Schneider, Juan Ortner, Avi Acharya, Dave Baron, Dana Foarta, Spencer Pantoja, Takuo Sugaya, Paul Milgrom, 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, Emiel Awad, Gilat Levy, Ricardo Alonso, Leeat Yariv, Nicolas Lambert, João Ramos, Wouter Dessein, Joel Sobel, seminar audiences, Christoph Carnehl for a close reading of an earlier draft, and especially Mike Harrison.

†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 delivers power to those who hold it. In domains from medical care to real estate, from car repair to business investment, experts are able to leverage decisions toward their own interests and against those of the decision makers themselves.[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'."

This has led to a large literature on the role of experts in decision making. At the foundation of all such models is a conception of expertise that specifies what the expert knows that a layperson does not. In canonical formulations, the informational gap is presumed to be minimal, with the expert's advantage being only a single piece of information and perfectly invertible (Crawford and Sobel, 1982; Milgrom, 1981). This parsimonious approach has yielded enormous insight into the structure of expert advice and provided the foundation for innumerable studies of expertise in markets and institutions.

In simplifying expertise to such a degree, however, these models generate several stark properties that do not resonate with how expertise is understood in practice. The canonical approach implies that a single recommendation by an expert can render a layperson an expert. In practice, a doctor not only knows much more than her patient, but a single diagnosis, no matter how precise, does not make her patient a medical expert. Moreover, in the classic model of Crawford and Sobel (1982), the expert's minimal informational advantage compels her to communicate inefficiently, and the information that she does convey leads to a decision that is optimal for the decision maker and not the expert herself. Rather than "overtowering" the decision maker, the expert retains no leverage and, in fact, 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. The expert knows many things the decision maker does not and her expertise is only partially invertible. We study the classic sender-receiver game of Crawford and Sobel (1982) for complex environments and identify an equilibrium that is sender optimal and Pareto efficient. In fact, the equilibrium delivers the outcome the expert would obtain were she to hold decision making authority herself. This result shows how the power of experts that is observed in practice can derive from the complexity of the decision making environment itself.

We develop these ideas through a particular representation of complex environments. Specifically, we represent the mapping from actions to outcomes as the realized path of a

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

Brownian motion. Thus, the state consists of a continuum of correlated variables, where each variable corresponds to the outcome of a particular action. The correlation across actions is given by the scale (or variance) of the Brownian motion, which, along with the drift of the motion, we use to parameterize the complexity of the underlying decision environment. The expert knows the entire path, whereas the decision maker knows only the parameters of the motion (the drift and scale) and the outcome of a status quo action.

Experts are empowered in complex environments because they are able to communicate precisely yet imperfectly. Because expertise is only partially invertible, an expert can make a recommendation that reveals precisely her most preferred action yet does not reveal all of her information. Precise but imperfect communication leaves the decision maker unsure as to his best response and more willing to accept the 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 ability to use her information while simultaneously keeping some of it private that is the foundation of expert power. Such a duality is not possible with simple expertise. When the expert's advantage is only a single piece of information, advice about one action is necessarily advice about all actions. In using her information in simple environments, the expert cannot also keep some of it private.

The expert's ability to control her information depends not only on the complexity of the environment, but also on the strategy she employs. In choosing one action to recommend, the expert is implicitly not choosing other actions, and this indirect information is valuable to the decision maker. We study a strategy in which the expert uses her recommendation to shape the indirect spillover of information in a way that is favorable to her.

In the Brownian environment, the expert's ability to shape indirect informational spillover depends on the size of the action space. On an unbounded action space, the expert is able to make her recommendation while containing the indirect spillover to a segment of the action space that is not otherwise attractive to the decision maker. In the region that is attractive to the decision maker, the expert's strategy keeps her information as private as possible, revealing nothing to the decision maker beyond the recommendation itself.

The logic is different in a bounded space. The sender's strategy cannot avoid indirect informational spillover to all actions. Although intuition suggests the expert is better off keeping her information as private as possible, we show that in this case the expert has more leverage precisely because the informational spillover is broader. The expert accrues more power because the additional spillover of information systematically favors the recommendation. Thus, the expert's power is maximized not because she minimizes informational spillover, but because she controls it in her favor. We show that the efficient equilibrium exists on an unbounded space only for small bias, whereas on a bounded space, the equilibrium exists for small and large bias.

Behavior in the equilibrium matches the many situation 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 recommendation serves the expert's interests and not his own. The equilibrium action also corresponds to that when the decision maker delegates authority to the expert. This correspondence is of particular interest to settings, such as politics, where the formal ability to delegate authority is tenuous yet plays a vital role in models of expertise (Gilligan and Krehbiel, 1987). Our result shows that when the environment is complex, the formal delegation of decision rights is not necessary for an expert to obtain Weber's (1958) "overtowering" position.

The Brownian motion representation allows us to illustrate the mechanics of efficient cheap talk, though it is not necessary for it to occur. Building on the logic of the Brownian representation, we identify a variety of natural settings in which efficient cheap talk is possible. In some environments the informational spillover is similar to the unbounded action space with the Brownian motion and in others it is similar to that in the bounded space. The common element across all settings is that by communicating precisely yet imperfectly, the expert is able to induce the decision maker to accept a recommendation despite knowing that it serves the interests of the biased 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)[2] Our contribution is to observe that in complex environments a gap emerges between efficient and perfect communication

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

and that the former does not require the latter. We show how this gap can be leveraged by the expert into favorable outcomes even when her interests do not coincide with those of the decision maker.

The logic of our equilibrium relies on CS's insight that informative communication with a biased expert requires pooling, or partitioning, of states. In the simple environment of CS, where each state corresponds to a different optimal action for the receiver, pooling is necessarily inefficient. In complex environments, the larger state space means that the expert can partition according to her optimal action. Therefore, the receiver can only partially invert the state, even from a message that perfectly reveals the agent's optimal action. In complex environments, precise communication can also be imperfect.

The size of the state space relative to the action space also distinguishes our approach from models of multidimensional cheap talk.[3] 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 model of complex environments.[4]

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

[3]It is well known that cheap talk can be efficient if the action space is reduced in size but only if the actions are discretized and far enough apart. By exogenously restricting the receiver's ability to adjust a recommendation, enough common interest is created between the players that the sender can obtain her optimal action. Indeed, CS's insight is to show for continuous action spaces how the action space can endogenously be discretized, but only if information is wasted and only if the sender gives up all of her leverage. Our result shows that the expert can communicate efficiently with continuous support and gain leverage over the decision if the state space is expanded and sufficiently complex.

[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.

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 not assuming 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, including communication games with cheap talk (Callander, 2008) and verifiable information (Callander, Lambert and Matouschek, 2021).[6] Callander (2008) identifies the efficient equilibrium when the expert's bias is sufficiently small and the action space unbounded. He leaves open the question of equilibrium for larger bias, noting only that the efficient equilibrium fails on an unbounded action space. We characterize an efficient equilibrium for larger biases and show how its existence depends on the size of the action space. The equilibrium logic for larger bias and a bounded action space is different to the small bias case and the mechanism for efficient cheap talk that we provide is new.[7]

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, $\psi(a) \in \mathbb{R}$. The status quo is action 0 with outcome $\psi(0) > 0$. The mapping is 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 space is the set of all such paths, which we denote by Ψ. 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 $\mathcal{A}$).

[6] Other applications include search and experimentation (Callander, 2011; Garfagnini and Strulovici, 2016; Urgun and Yariv, 2021; Cetemen, Urgun and Yariv, 2021), "attributes" problems (Callander and Clark, 2017; Carnehl and Schneider, 2021; Bardhi, 2022; Bardhi and Bobkova, 2021), and industrial organization (Callander and Matouschek, 2022).

[7] The "referential advice" of Callander, Lambert and Matouschek (2021) refers to information revelation beyond the recommendation itself, a possibility with verifiable information but not with cheap talk.

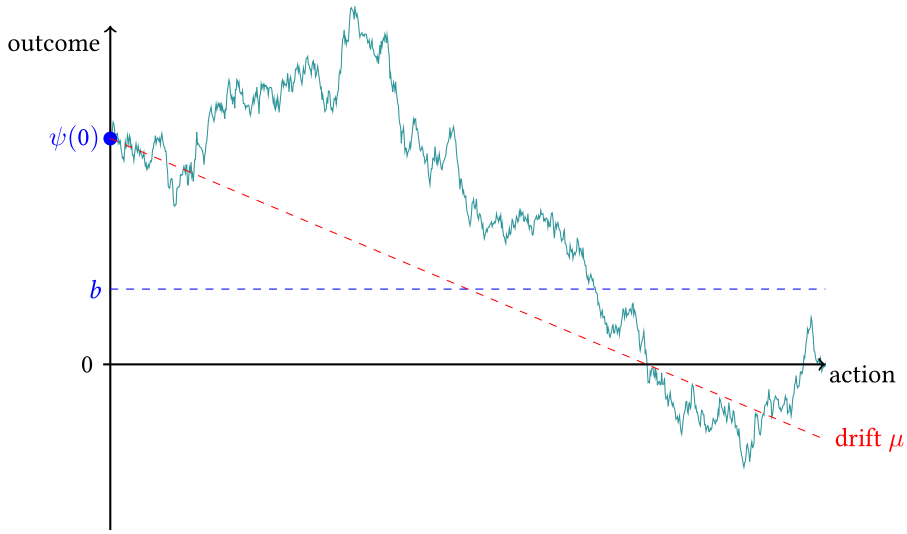

Figure 1: The Mapping From Actions to Outcomes

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 utility functions that are weakly concave in outcomes and uniquely maximized at outcomes 0 and b.[8] We assume $b < \psi(0)$ and address the case of larger bias separately after our main result.

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

Remark 1: In CS actions map directly to utility. In most applications actions map to an outcome, from which agents draw utility. Formalizing this intermediate step allows a clearer view of the decision making environment. Viewed through this lens, the state space in CS in the fixed bias case is equivalent to a mapping $\hat{\psi}(a) = \theta - a$, where $\theta \in [0, 1]$ is the expert's single piece of private information.

This formulation highlights another difference between the models. We assume, for convenience, that the status quo outcome is known. Such an assumption in CS would undermine the sender's expertise as it would fully reveal the mapping. In our model, knowledge of a

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

single point amongst a continuum of unknowns is immaterial (and can easily be relaxed).

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

With the Brownian motion, the sender's advantage is a continuum of information and complexity is the correlation across that information. An alternative representation of complexity is by the number of discrete pieces of information a sender knows that a receiver does not.[10] In Section 6.2 we present several environments that extend CS in this way and which support efficient cheap talk.

Remark 3: 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:

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

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

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

[10]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.

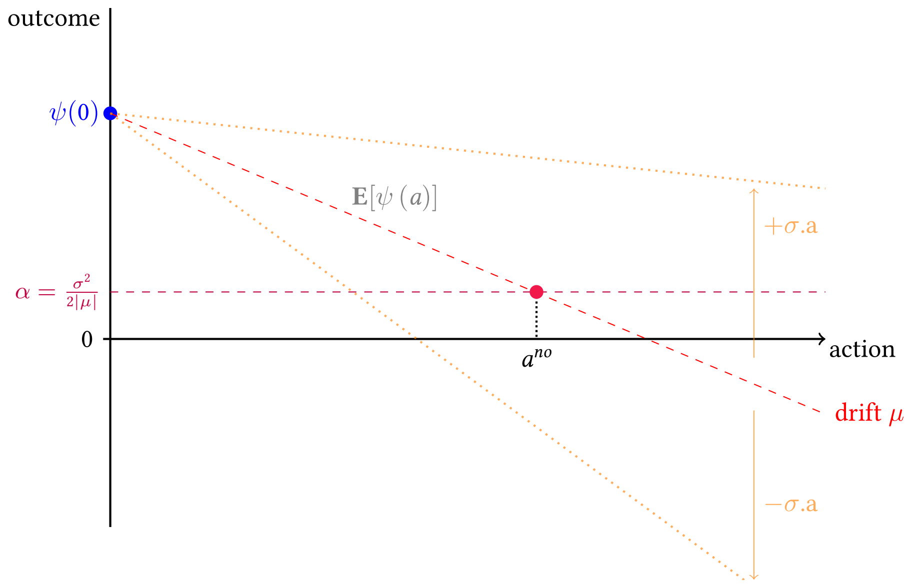

Figure 2: Beliefs in the Absence of an Expert

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.[11]

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

[11]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 17 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 $\psi(0) \leq 0$ as his utility is simply $-\psi(0)^2$. For $\psi(0) > 0$ 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 depends on 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 power of the first point strategy is how it compartmentalizes the spillover of information from the recommendation. In complex environments a distinction arises between direct and indirect informational spillover.[12] *Direct* spillover is what the receiver learns from the recommendation itself. In recommending action r^*, the sender reveals that the mapping passes through the point (r^*, b) and this shapes the receiver's beliefs about all other actions.

Indirect informational spillover comes instead from the strategy. It is what the receiver infers from the fact that r^* was the recommendation and not some other action.[13] The first-point strategy is effective because it contains the indirect informational spillover into 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 there is a better action for him with probability one, but he doesn't know what it is. He faces what we refer to as *response uncertainty*. The receiver's best response could be any action to the right of the recommendation, and this uncertainty creates enough risk in choosing any particular action that the receiver acquiesces to the recommendation. The logic of this equilibrium does not depend on the quadratic utility form, though it does depend on the receiver being risk averse.

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

[13]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.

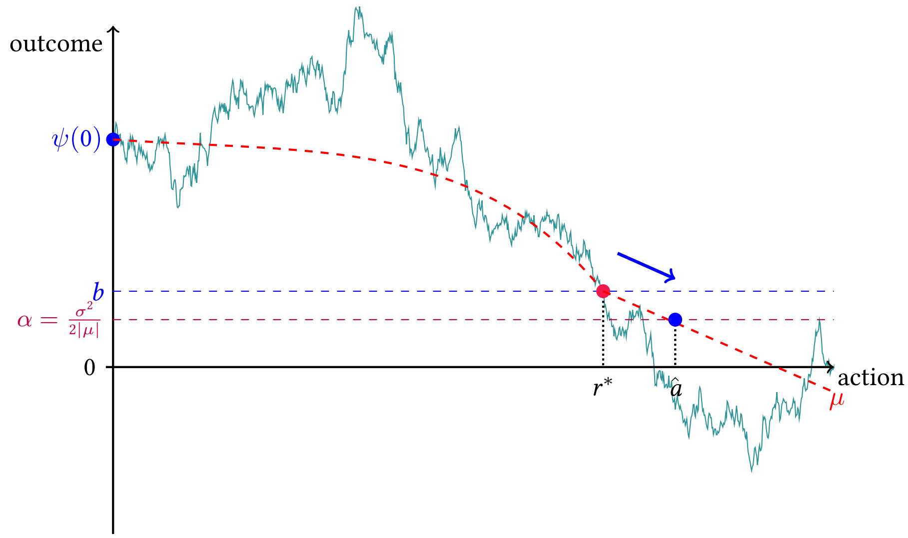

Figure 3: First-Point Strategy: Recommendation r^* and Receiver's Optimal Response $\hat{a}$.

The equilibrium does not hold for larger bias as the same inference would lead the receiver to override the recommendation. It is important to note that in this situation the receiver would not ignore the information contained in the recommendation. He would use the knowledge that r^* maps into outcome b to guide his response. Figure 3 depicts the situation in which the receiver overrides recommendation r^* with action $\hat{a}$. Knowledge from the recommendation means the receiver obtains an expected outcome of α with lower variance than had the expert not made her recommendation.[14]

This example gets to the heart of the sender's challenge. In giving advice, the sender must use her information, but by using her information, she makes it possible for the receiver to repurpose that information to his own ends. In simple environments informational spillover is complete and this undermines efficient cheap talk. In complex environments, the sender's challenge does not go away although it is ameliorated. It is intuitive that the sender will 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 more information spills over from the first-point strategy.

4.2 Bounded Action Space

A bounded action space brings several changes to the decision problem. It implies that the expert knows fewer things, and that the receiver is constrained in his choice should he override a recommendation. The most important difference, however, is that with positive probability

[14]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.

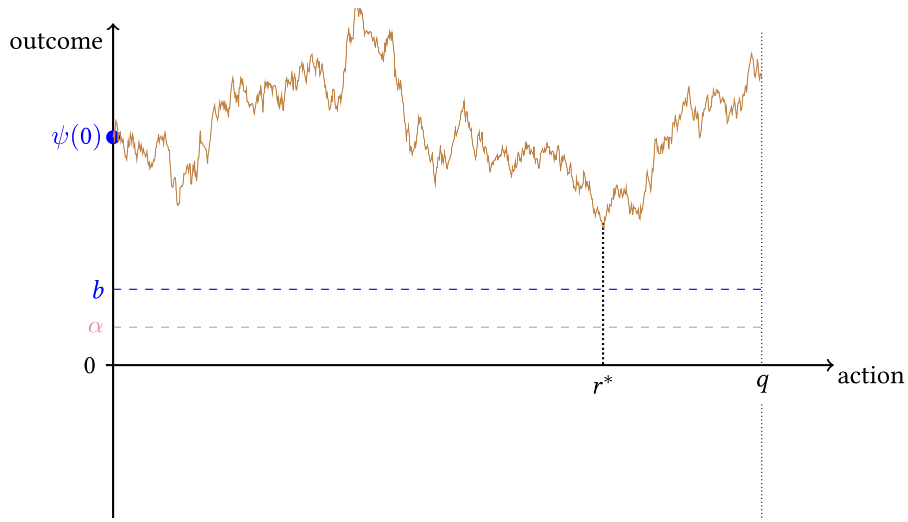

Figure 4: First-Point Strategy & Recommendation r^*: Outcome Above b

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. Despite misaligned preferences over outcomes, the players now have aligned preferences over actions, where this alignment emerges endogenously from the complexity of the environment. This situation is depicted in Figure 4.

This possibility implies that on a bounded space the sender reveals even less information about the recommendation. Using 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.[15] 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. For an outcome of the recommendation 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. How much his beliefs are skewed is critical to supporting equilibrium.

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

For low levels of bias, this uncertainty 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 the complementary probability, the outcome is above b and the receiver's interests are aligned with the sender on the recommendation r^*.[16]

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. The essential requirement for efficient cheap talk, regardless of the set of possible best responses, is that the receiver resolves his response uncertainty by choosing the recommendation itself. In the Brownian environment, this means that the receiver chooses one particular possible best response and not the other, and not an intermediate action even though they are available. Were the receiver to deviate from the recommendation to any degree, the sender, anticipating this response, would shade her recommendation to the left, and efficient cheap talk would unravel as it does in simple environments.

For the equilibrium to hold, 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 wide.

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

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

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

That the sender revealing more information 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 limits the informational spillover and contains it in one part of the action space. Increasing informational spillover improves communication on a bounded space because the first-point strategy shapes that spillover in a way that favors 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.

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

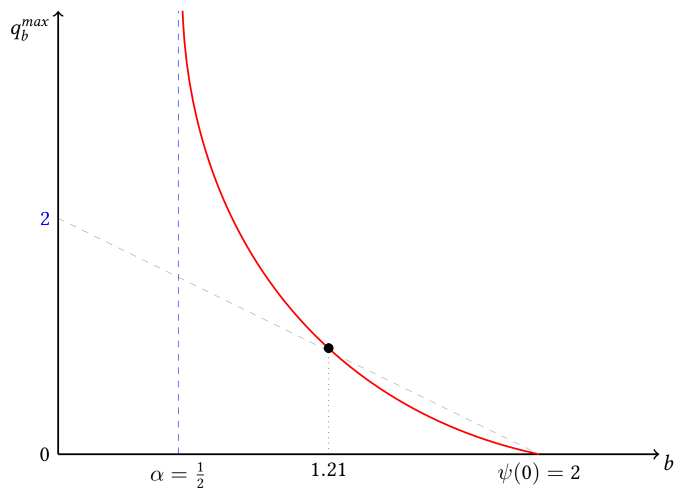

Figure 5: q_b^{max} for $\mu = -1, \sigma^2 = 1, \psi(0) = 2$, and $\alpha = \frac{\sigma^2}{2|\mu|} = \frac{1}{2}$.

Theorem 1 establishes that the first-point equilibrium exists for biases up to the level of the status quo outcome $\psi(0)$. For each level of bias, there is a maximum size of the action space, q_b^{max}, that supports the equilibrium, where q_b^{max} is finite for bias greater than α. The theorem establishes that the equilibrium exists for all action spaces smaller than the maximum. This represents a strengthening of Lemma 2 for small bias as it implies that the first-point equilibrium exists for an action space of any size.

The upper bound on the size of the action space decreases in the 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 the bias, giving the receiver a greater incentive to override the recommendation. To maintain equilibrium, therefore, the action space must contract in b.[17]

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 5 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 receiver have opposing interests—with substantial probability.[18]

For bias beyond $\psi(0)$, the interests of the players relative to the status quo are directly

[17] 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.

[18] 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.

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. 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 logic of Theorem 1 in more detail. 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 $q > 0$. We begin with 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 formalized as the probability that the Brownian motion first hits b at action r^*. Defining the first hitting action for outcome y as:

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

We have:

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

In the appendix, we provide a closed form expression for this probability 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

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

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:

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

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

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

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

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

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

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 6. 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 6 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

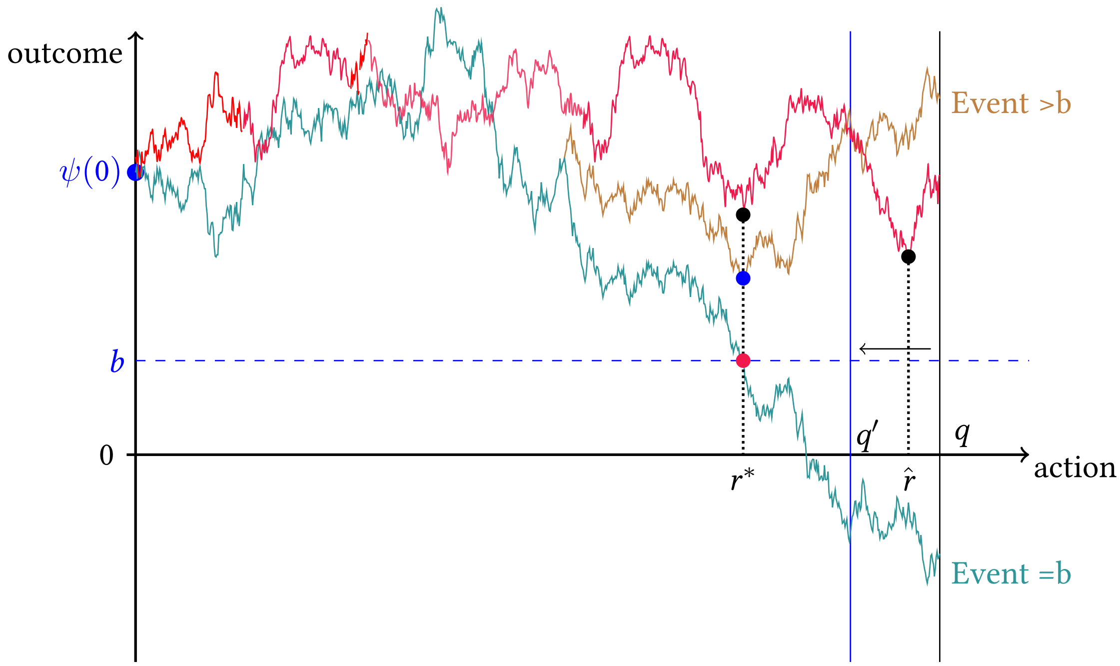

Figure 6: Paths that Induce r^* as $q \to q'$.

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.

As noted earlier, the receiver faces two potential best responses, either r^* or $r^* + \frac{b-\alpha}{|\mu|}$, and the challenge of efficient cheap talk is that the receiver must resolve this uncertainty by choosing r^* exactly and not 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.[20] We obtain expressions for the expected value of a Brownian meander with a known terminal value c at $a = q$. This value is continuous in a and c and we show that the derivative at the recommendation r^* is infinite for any $c > b$. Thus, by the law of iterated expectations, the marginal cost of deviating from the recommendation in Event >b is infinite.

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

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

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

Lemmas 3 and 4, along with the earlier Lemma 2, deliver Theorem 1. It is worth noting that the equilibrium existence argument does not rely on risk aversion for the receiver (although risk aversion does change the exact domain of existence). Deviating to the right in Event >b leaves the receiver worse off with certainty and not just in expectation. With sufficiently high likelihood of Event >b, therefore, even a risk neutral receiver would accept the recommendation. This is not the case for the equilibrium on an unbounded action space.

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

[20]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 and σ in the online appendix.

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.[21] How the receiver's beliefs are shaped by the recommendation is as important as the extent of the sender's knowledge.

4.4 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 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.

It also 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

[21] This does not imply that the sender needs to know the entire path.

a degree that the receiver prefers an efficient sender-optimal equilibrium. Identifying inefficient equilibria in complex environments, if they exist beyond the babbling equilibrium, is an important question for future work.

4.5 Welfare and Comparative Statics

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, for $q < q_b^{max}$, 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 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, for $q < q_b^{max}$, 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 herself is better off as she is already obtaining her best action. Her utility increases as 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.[22]

[22] The greater likelihood of obtaining outcome b is relative to the fixed status quo. If $\psi(0)$ and b increase in parallel, the sender's utility is unchanged.

Complexity of the environment: The complexity of the environment has a different impact depending on the size of an action space. On an unbounded space, the path crosses b with probability one and, conditional on the first-point equilibrium existing, the sender obtains her ideal outcome, regardless of how complex the environment is. The receiver accepts the recommendation because the risk of overriding is not worth the benefit.

On a bounded space, the path does not cross b with probability one. Indeed, this is essential to supporting the first-point equilibrium for bias larger than α. In numerical simulations q_b^{max} increases in σ for $b > \alpha$, such that the size of the action space that can support equilibrium increases in complexity. For a fixed action space, $[0, q]$, the probability that the mapping crosses b approaches one as σ becomes large. This does not undermine the first-point equilibrium as α increases in σ, and once $\alpha > b$ the receiver will not want to override the recommendation even in Event =b. Thus, even on a bounded space, highly complex issues imply that the sender obtains her ideal outcome with high likelihood.[23]

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

At the other extreme, the threshold α approaches zero as complexity approaches zero. On an unbounded action space, this means that efficient cheap talk is possible only for vanishingly small bias. In the limit, it 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. They show that, in simple environments, the same limit is approached by the most informative partition equilibrium as bias approaches zero.

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.[24]

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

[23]It is difficult to establish monotonicity properties of the equilibrium in σ analytically as it is not immediate that the probability the path hits b is monotonic for all values of σ. To see why, it is easiest to view the state as the underlying Wiener process and the mapping ψ as 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 then an increase in σ moves that minimum closer to b, whereas if the minimum is above the drift line then 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.

[24]Formally, set $T(a) = \arg\min_{a \in \mathcal{A}} \mid b - \psi(a) \mid$, then $m^*(\psi) = \min_{T(a)}[a \mid a \geq a' \ \forall a' < 0]$.

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.[25]

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.[26] The equilibrium can be restored on a bounded space as then there is the possibility that the path doesn't reach b and the outcome of the recommendation is above 0, Event >b dominates Event =b for small recommendations.

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. 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 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

[25] See Callander and Hummel (2014) for a related effect in a search model with symmetric information.

[26] 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.

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.[27] 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.

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

6 Beyond the Brownian Motion

The Brownian motion provides one environment that permits efficient cheap talk, yet it is not the only such environment. In this section we develop the key ingredients necessary for efficient cheap talk and construct complex environments that support it.

6.1 Ingredients for Efficient Cheap Talk

The mechanics of efficient cheap talk can be understood through several increasingly demanding requirements on the receiver's response. The most demanding is the equilibrium condition that the receiver accept the recommendation. We refer to this below as *recommendation acceptance*. Two weaker requirements—*partial invertibility* and *response uncertainty*—help shed light on the nature of uncertainty that underlies equilibrium.

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 $\mid m^{-1}(r) \mid > 1$ 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 $m(\cdot)$.*[28]

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.

Partial invertibility is necessary but not sufficient for efficient cheap talk with a biased sender. 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 $\hat{a}(\psi) = \arg\max_a u^R(a, \psi)$ 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 $\mid \hat{a}(m^{-1}(r)) \mid > 1$ 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

[28] The definitions of perfectly-invertible and non-invertible follow naturally.

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 $r^* - b$.[29] 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 a compromise action between r^* and $r^* - b$ and not the recommendation itself. (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 following condition, which we refer to as *recommendation acceptance*, is sufficient for this and efficient cheap talk to hold. Denote by $a(r)$ the receiver's best response to a recommendation r.[30]

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

Recommendation acceptance 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 the three requirements to construct environments that support efficient cheap talk and illuminate why it is possible and when it is not.[31]

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

[29] 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.

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

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

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.[32]

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

The fragility of this equilibrium resonates with the results in models of unknown bias. As noted earlier, the sender's informational advantage in those models is also two pieces of information. The complex environment described here shows that efficient cheap talk is possible in such settings but that the conditions required are demanding.[33]

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

[32] 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.

[33] 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).

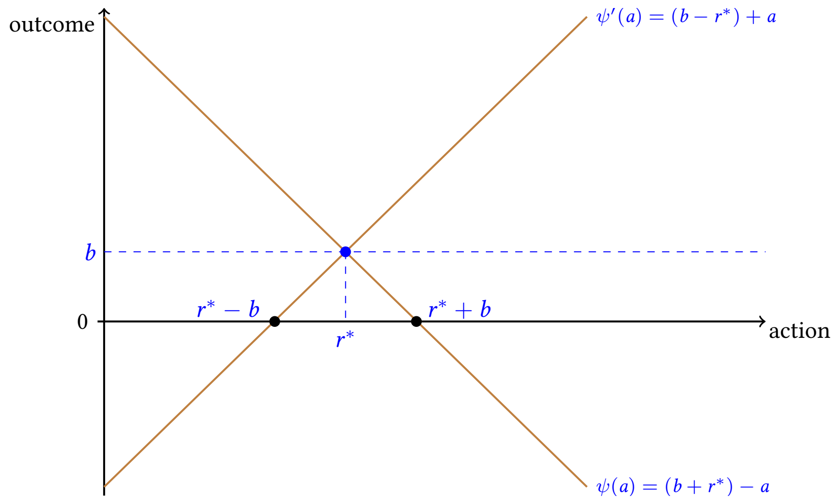

Figure 7: Directional Uncertainty

not only response uncertainty, but *directional uncertainty* as well.[34] Directional uncertainty fills the role of the non-monotonicities in the Brownian environment. It creates the necessary uncertainty over outcomes should the receiver override the recommendation, despite the true mapping being monotonic.

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

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

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

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

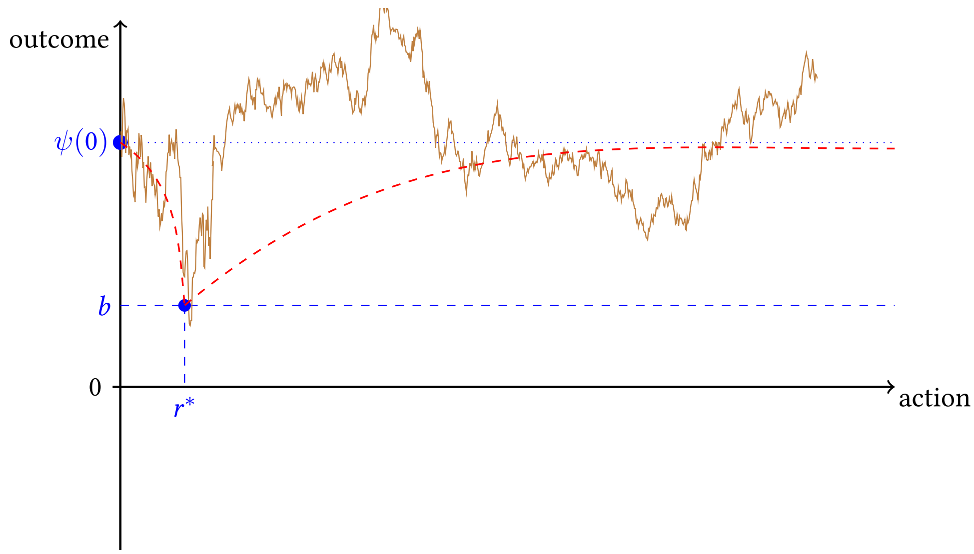

Figure 8: Mean Reverting Ornstein-Uhlenbeck Process With Mean $\psi(0)$

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 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 $\psi(q) > b$.

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. What other equilibria exist is a question of obvious interest. An answer to this question would show how to create a bridge between the efficient equilibria of complex environments and the inefficient equilibria of CS in simple environments. Building such a bridge requires that we relax the assumption of a known status quo, which our analysis can easily accommodate. The simple environment of CS can then be understood as the limit of the sequence of environments in which the scale of the Brownian motion approaches zero. Identifying the set of equilibria along this path is a natural question for future work.

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 restrict her 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.

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.

Callander, Steven and Patrick Hummel. 2014. "Preemptive policy experimentation." *Econometrica* 82(4):1509–1528.

Callander, Steven and Tom S. Clark. 2017. "Precedent and Doctrine in a Complicated World." *American Political Science Review* 111(1):184–203.

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.

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.

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 tranformation 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 $(\mathcal{W}, \mathcal{B}(\mathcal{W}))$.[35] 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 $\omega(\cdot \mid r \in m(\psi)) : \mathcal{B}(\mathcal{C}[0,1]) \times \mathcal{M} \to [0,1]$.[36] 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.

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

[36] 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 $X(a') = \mu a' + \sigma W(a')$.[37] 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.[38] 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

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

[38] 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:

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

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

$\{\psi \in \Psi \mid \tau_\iota(q') = r^*, \iota(q') > b, \tau_\iota(q) > q'\}$.[39] 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 $[0, q']$. ■

[39]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^*.

Proof of Lemma 4. Consider the game with $\mathcal{A} = [0, q]$, for some $q \in \mathbb{R}_{++}$. The first-point equilibrium has full support over $\mathcal{A}$. Let any off-path recommendation $r' \notin \mathcal{A}$ be interpreted as an on-path message, say $r'' = 0$, and generating the same beliefs. Thus, it 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. $\text{Var}[\psi(a' + r^*) \mid m_q^*(\psi) = r^*] \geq \text{Var}[\psi(r^*) \mid m_q^*(\psi) = r^*]$.

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 $r^* \leq \delta_\varepsilon$.[40] Similarly, in the online appendix Corollary B.1 we show that: $\lim_{a' \to 0} \frac{\partial}{\partial a'} \mathbb{E}[M(a', q - r^*)] = \infty$. Thus, for every $N > 0$ there exists a $\delta_N \in \mathbb{R}_+$ such that $\mathbb{E}[M(a', q - r^*)] > Na'$ whenever $a' < \delta_N$.

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

$$\underbrace{\mathbb{P}(\tau(b) \in dr^* \mid m_q^*(\psi) = r^*)}_{\leq \varepsilon} \mu a' + \underbrace{\mathbb{P}(\tau_\iota(q) \in dr^*, \iota(q) > b \mid m_q^*(\psi) = r^*)}_{\geq 1-\varepsilon} \underbrace{\mathbb{E}[\psi(a' + r^*) - \psi(r^*) \mid \tau_\iota(q) = r^*, \iota(q) > b]}_{= \mathbb{E}[M(a', q - r^*)] > Na'}$$
$$\geq \varepsilon(\mu a') + (1-\varepsilon)Na' = a' \underbrace{(\varepsilon\mu + (1-\varepsilon)N)}_{\geq 0} \geq 0$$

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

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

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

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

and there exists some $\tilde{a}, \tilde{r} \in \left[0, q_{\max}^b\right]$ with $\tilde{a} = \tilde{r} + a'$ and $a' > 0$ such that this holds with equality by the maximality of $q_b^{\max}$.[41] As usual, we can write this as:

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

Identically to Lemma 4, we can write this in terms of increments given by Brownian motion

[40] 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 $\lim_{h \to 0} P(|\psi(a+h) - \psi(a)| > \varepsilon) = 0$.

[41] 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 $a' > 0$.

$X(\cdot)$ and Brownian meander $M(\cdot,\cdot)$. Thus, we have that:

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

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.[42] 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:

$$\begin{aligned}&\left(\frac{\partial}{\partial b}\mathbb{P}\left(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r}\right)\right)\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] \quad (6)\\ &+\left(\frac{\partial}{\partial b}\mathbb{P}\left(\tau_\iota(q) \in d\tilde{r}, \iota(q) > b \mid m_b^*(\psi) = \tilde{r}\right)\right)\mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q-\tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right] \quad (7)\\ &+\mathbb{P}\left(\tau(b) \in d\tilde{r} \mid m_b^*(\psi) = \tilde{r}\right)\left(\frac{\partial}{\partial b}\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right]\right) \quad (8)\\ &+\mathbb{P}\left(\tau_\iota(q) \in d\tilde{r}, \iota(q) > b \mid m_b^*(\psi) = \tilde{r}\right)\left(\frac{\partial}{\partial b}\mathbb{E}\left[u_R\left(\psi(\tilde{r}) + M(a', q-\tilde{r})\right) - u_R\left(\psi(\tilde{r})\right) \mid \tau_\iota(q) = \tilde{r}, \iota(q) > b\right]\right) \quad (9)\end{aligned}$$

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.[43]

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

[42] 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.

[43] 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 $\mathbb{E}\left[u_R\left(b + X(a')\right) - u_R\left(b\right)\right] > 0$.

Thus, the terms (6) and (7) are positive for any weakly concave utility function that is uniquely maximized at 0. Similarly, it directly follows that the sum of (8) and (9) are non-negative if and only if the following holds.[44]

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

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

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

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 $q_b^{\max}$.[45] 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

[44] 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.

[45] 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 $\frac{\partial}{\partial b}\mathbb{E}\left[u_R\left(b+X(a')\right)-u_R(b)\right] > 0$. Thus, rearranging replacing the RHS of (10) and rearranging we get (5).

can be made strictly better of with action $\hat{a}$ and the equilibrium is not Pareto efficient. Thus, $a(m(\Psi)) = \mathcal{A}$.

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

Proof of 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 recommendation acceptance, 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 recommendation acceptance, 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 response acceptance analogously to the Brownian motion and Ornstein-Uhlenbeck environments.