The Power of Referential Advice*

Steven Callander† Nicolas S. Lambert‡ Niko Matouschek§

May 11, 2021

Abstract

Expert advice is often rich and broad, going beyond a simple recommendation. In this paper we show that this additional, referential information plays an important strategic role in expert advice and that it is vital to an expert's power. We develop this result in the context of the canonical model of strategic communication with hard information, enriching the model with a notion of expertise that allows for a meaningful distinction between a recommendation and referential information. Referential advice matters because effective communication is a process of both persuasion and dissuasion. When constructed just right, referential advice dissuades the decisionmaker from choosing options other than the recommendation, thereby making the recommendation itself more persuasive. We identify an equilibrium in which, with probability one, the expert is strictly better off by providing referential advice than she is in any equilibrium in which she provides a recommendation alone. The benefit of referential advice to the expert is non-monotonic in the complexity of her expertise, reaching its peak when expertise is moderately complex.

*We thank Arjada Bardhi, Dirk Bergemann, Ben Brooks, Wouter Dessein, Marina Halac, Gilat Levy, Meg Meyer, David Myatt, Motty Perry, Andrea Prat, Luis Rayo, Daniel Seidmann, Joel Sobel, Bruno Strulovici, Ashutosh Thakur, Eray Turkel, participants at various conferences and seminars, an editor and three referees for their comments and suggestions. Lambert thanks Microsoft Research and the Cowles Foundation at Yale University for their hospitality and financial support.

†Stanford University, Stanford Graduate School of Business; sjc@stanford.edu.

‡Massachusetts Institute of Technology, Department of Economics; lambertn@mit.edu.

§Northwestern University, Kellogg School of Management;
n-matouschek@kellogg.northwestern.edu

1 Introduction

Advice takes many forms. A common form is for an expert to simply offer a recommendation: A librarian recommends a book, a travel agent recommends a tour, or a sales assistant recommends a pair of shoes. In many situations, however, an expert does not limit herself to only a recommendation. Instead, the expert provides advice that is more expansive and that conveys information about options beyond the one recommended. This richer, contextual advice—what linguists refer to as *referential* information (Jakobson, 1960)—is particularly relevant when the expert possesses complex knowledge. For instance, in addition to recommending a treatment, a doctor will often discuss alternative treatments and why she does not recommend them. Similarly, a mechanic might detail likely outcomes should a car owner undertake only superficial repairs instead of the more extensive and expensive ones she does recommend.

The role that referential information plays in the supply of expert advice has not previously been examined. As such, it is unclear whether referential information plays a meaningful role in communication or whether it is superfluous or even babbling. The objective of this paper is to offer the first analysis of referential information and address these questions.

In a model of hard (verifiable) information, we show that referential advice fundamentally changes the nature of strategic communication. By supplying referential advice in just the right way, the expert is able to leverage her expertise and systematically sway decisions in her favor. We identify an equilibrium in which the expert provides referential advice and obtains a payoff that is strictly higher than in any equilibrium in which she provides only a recommendation. This equilibrium implies that the mechanic, by providing referential advice, is able to induce the owner to spend more money on repairs, regardless of the true damage to a car, than were she to provide a recommendation alone.

To begin to understand the role of referential advice, we need to examine more deeply what it means to be an expert. In the classical formulations of Milgrom (1981) and Crawford and Sobel (1982), expertise is modeled as a single piece of information that the expert alone possesses. This formulation provides no space for referential advice. Because the same single piece of information affects all options (often in an identical way), advice about one option is necessarily advice about all options. Moreover, this simple view of expertise conflicts with how expertise is commonly understood. Weber (1922) emphasises that the gap in knowledge between an expert and a decisionmaker is large and that it is this gap that provides the source of an expert's power.[1]

Taking inspiration from Weber, we build a model that broadens the conception of what it means to be an expert. Specifically, we suppose that each option is associated with a unique state variable that determines the outcome of that option, and that the states are imperfectly correlated. The expert knows these states but the decisionmaker does not and the gap between them is large.[2] The degree of correlation across what the expert knows measures the *complexity* of her expertise.

The key novelty of this approach is that it allows the expert to communicate precisely yet imperfectly. This matters because it implies that the information that spills over from the recommendation to other options is incomplete. Informational spillover that is incomplete resonates more closely with expertise in practice. It says that in recommending a treatment for migraines, say, a doctor reveals some information about the likely outcome of similar treatments, but does not reveal all of her expertise.

This contrasts with the canonical formulation in which informational spillover is complete. An expert who communicates precisely also communicates perfectly. The expert is unable to use her information to make a precise recommendation without revealing all of her expertise to the receiver and, in effect, rendering the receiver an expert. This leads to the seminal result of the hard information literature—the famous "unravelling" result—that in the unique equilibrium the expert's information advantage unravels and she fully reveals her information to the decisionmaker (Milgrom, 1981; Grossman, 1981).[3] As a result, the expert retains no leverage and the decision that is made aligns fully with the preferences of the decisionmaker.

Our richer notion of expertise provides the expert with greater ability to keep her information private, to use her information and protect it at the same time. We show she is able to do this to the maximal extent. We identify a continuum of equilibria in which the expert reveals the minimum amount of information to influence decision making—i.e., the outcome of a single state variable. The option that is revealed constitutes a recommendation—what linguists refer to as the *conative* function of language (Jakobson, 1960)—and the decisionmaker follows the recommendation.

---

[1]Weber remarked that "the political master finds himself in the position of the 'dilettante' who stands opposite the 'expert'," and that "under normal conditions, the power position of a fully developed bureaucracy is always overtowering" (Weber et al., 1958).

[2]The gap in knowledge may go the other way as well if the receiver has information about his preferences that the expert does not. We do not explore this possibility here.

[3]See also Milgrom and Roberts (1986); Matthews and Postlewaite (1985); Seidmann and Winter (1997).

The existence of recommendation-only *conative* equilibria imply that complex expertise does not necessarily imply complex communication. A doctor can, despite her extensive knowledge, simply recommend a treatment and know that the patient will follow her advice. As these equilibria maximize the expert's ability to shield her information from the decisionmaker, it may be reasoned that they are the expert's preferred equilibria. We show, however, that this is not the case and that she can do strictly better with referential advice.

The expert's leverage is limited in conative equilibria because informational spillovers, while imperfect, still exist. The decisionmaker can take the information in a recommendation and, though unsure about what it implies about other options, use the information it does reveal to find a better alternative. A cautious doctor, for example, has limited ability to persuade because when she recommends four weeks of physiotherapy for a knee injury, the less-cautious patient can infer that his injury is not life-threatening or even permanently disabling, and instead decide that a few sessions are enough. This allows the decisionmaker to extract the surplus from expertise. We show that in the case in which the decisionmaker has quadratic utility, the expected decision in the receiver-optimal conative equilibrium is equal to what the decisionmaker would choose in the absence of expertise. This means that a risk-neutral expert is, at best, indifferent between providing conative advice and not having expertise at all.[4]

Referential advice allows the expert to avoid this fate. By supplementing her recommendation with referential advice, the expert is not so much able to avoid information spillover, but to redirect it in her favor. To understand how, it is important to understand that effective communication is a process of both persuasion and dissuasion. To *persuade* the decisionmaker to follow a recommendation, the expert must simultaneously *dissuade* him from choosing another option. Referential advice allows the expert to separate these tasks. The recommendation persuades and the referential advice dissuades. A recommendation alone cannot carry the weight of both tasks.

Although some referential equilibria benefit the expert, not all do. Full revelation is an equilibrium in our setting, for example, and it reveals the maximal amount of referential advice while delivering to the receiver his optimal action. While full revelation eliminates the expert's leverage, it informs why better referential equilibria do exist. Referential advice creates an expectation in the receiver for extra information, and this expectation ties the hands of

[4]We follow the hard information literature in assuming that the expert cares about the decision made (the cost of repairs to the mechanic) whereas the decisionmaker cares about the outcome.

the sender, limiting her ability to strategically choose the advice she conveys. When done poorly, as with full revelation, this expectation is bad for the sender. When done right, the expectation helps the sender precisely because tying her hands makes her messages more credible and, thus, dissuasion more effective.

The richer notion of expertise we introduce significantly expands the space of types, messages, and beliefs. This reflects the reality of expertise in practice, but it does mean a complete characterization of equilibrium is out of reach. Nevertheless, we are able to establish a general dominance result. We show that one particular referential equilibrium—the "interval" equilibrium—dominates all conative equilibria and dominates all equilibria that also dominate all conative equilibria, and that it is undominated across a broad class of equilibria. The interval equilibrium is intuitive and simple to describe, and captures clearly the power of referential advice.

Related Literature

We are not the first paper to enrich the informational structure of the canonical sender-receiver model, although the focus and intention of those papers is very distant from ours, and we are the first to identify a role for referential information. The knowledge of the sender (the expert) is generalized to multiple dimensions in Glazer and Rubinstein (2004), Shin (2003), and Dziuda (2011). More recently, Hart et al. (2017) generalize further and assume only that knowledge satisfies a partial order (see also Ben-Porath et al., 2019 and Rappoport, 2020). In these settings, the information that is strategically withheld from the receiver (the decisionmaker) is relevant to all options and, in equilibrium, the receiver is uncertain of the outcome he will receive from his choice. In equilibrium in our model, in contrast, the receiver is certain of the outcome he will receive from his decision as all unrevealed information is irrelevant to his choice. This provides a sharp distinction in interpretation. In our setting, the extra information provided is purely referential and aimed at dissuasion (and persuasion indirectly, as explained above), whereas in these other models, the conative and referential functions of language are intermingled and all information that is provided constitutes part of the recommendation and is aimed at persuasion directly.

A separate, prominent strand of the hard information literature, due to Dye (1985), incorporates the possibility that the sender is uninformed and, thus, not an expert at all (see also Dziuda, 2011). This implies that should the receiver not receive some information, he is unsure whether the sender is deliberately withholding the information or whether she doesn't have it at all. This concern is not present in our model. Throughout our analysis, the sender is informed and the receiver knows this with certainty.

In his famous treatise, Jakobson (1960) delineates six purposes of language, and we borrow only two. Although there is a considerable chasm between the setting of Jakobson (1960) and the formalism of models of strategic communication, we find resonance. According to Jakobson (1960), the conative function of language is a call for action, what we take as a recommendation. The referential function of language, in contrast, is to provide context and is the additional, non-prescriptive information that a speaker may offer about the world. This translates in our setting into the information an expert provides beyond her recommendation as, by construction, this information does not affect beliefs about the recommendation itself.[5]

Our approach represents a novel use of stochastic process, in particular the Brownian motion, in strategic communication with hard information. The use of the Brownian motion to represent uncertainty has recently found application in models of search (Callander, 2011; Garfagnini and Strulovici, 2016) and other symmetric information environments. Callander (2008) examines a model of cheap talk communication when the expert cares about the outcome rather than the action, as is the case here. He identifies an informative equilibrium, although referential advice plays no role and communication is purely conative. The equilibria we identify are logically distinct.

Our model is also distinct from the flourishing literature on information design (Kamenica and Gentzkow, 2011; Rayo and Segal, 2010). Our core difference with that literature is the absence of commitment. In our model, neither the receiver nor the expert can commit to any particular course of action. Similarly, communication in our model is without institutional constraint. In political economy, the influential model of Gilligan and Krehbiel (1987) demonstrates how legislatures can organize themselves by committing to formal institutional structure and rules that incentivize and leverage expertise in policymaking. Our model, instead, contributes to our understanding of how and why experts can wield power even in absence of commitment or institutional structure.

[5] We treat the hard information about the recommendation as part of the recommendation itself. An alternative interpretation is to treat that information also as referential. Sobel (2013) was the first to connect Jakobson's typology to games of strategic communication and he offers a more complete translation and interpretation into the language of strategic communication games.

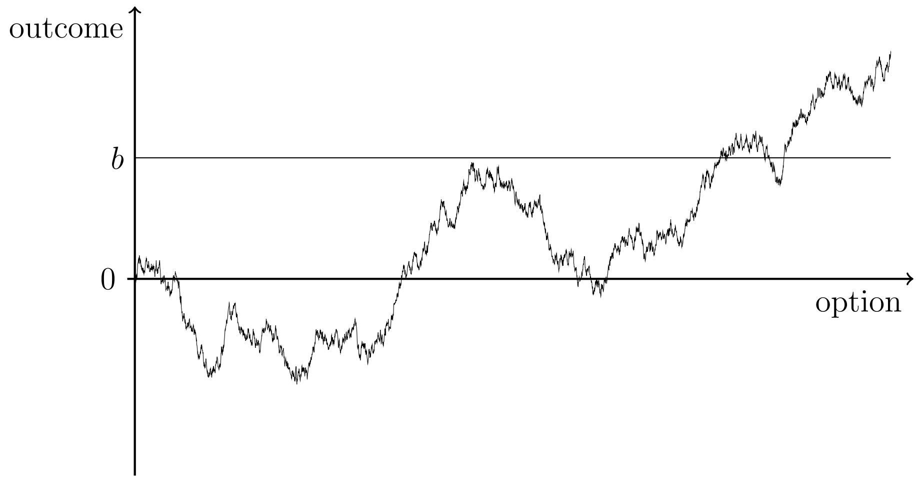

Figure 1: Example of an outcome path.

2 Model

A sender and a receiver play a game of strategic communication with hard information. We introduce a novel notion of rich expertise but otherwise hew as closely as possible to classical models of hard information (Milgrom, 1981; see also Meyer, 2017 and Gibbons et al., 2013). We follow the convention of using the female pronoun for the sender and the male pronoun for the receiver.

Technology: To expand beyond expertise as a single variable, we suppose that expertise is knowledge of an entire function. To capture the reality that this knowledge is multifaceted and complex, we represent the function using a stochastic process that in the limit is given by a Brownian motion. Figure 1 depicts one such path.

To construct this path, we study a discrete stochastic process that converges on the Brownian motion in the limit. The set of options is $\mathcal{D} = \{d_0, d_1, \ldots, d_n\}$, where $n \geq 2$, $d_0 = 0$, and

$$d_i = d_{i-1} + \frac{1}{\sqrt{n}} \quad \text{for} \quad i = 1, 2, \ldots, n.$$

The $n+1$ options, therefore, span the interval $[0, \sqrt{n}]$, with each being equally far from its neighbors. Option d generates outcome $X(d)$ according to the outcome function $X(\cdot)$. The outcome function is the realization of a random walk with drift. Specifically, $X(d_0) = 0$ and

$$X(d_i) = X(d_{i-1}) + \frac{\mu}{\sqrt{n}} + \frac{\sigma}{\sqrt[4]{n}}\theta_i \qquad \text{for} \qquad i = 1, \ldots, n,$$

where each θ_i is independently drawn from the standard normal distribution, $\mu > 0$ measures the expected rate of change from one option to the next, and $\sigma > 0$ scales the variance of each option relative to its neighbors.

In the limit, this outcome function $X(\cdot)$ becomes the realization of a Brownian motion with drift μ and scale σ on the non-negative half line $[0, \infty)$. For simplicity of exposition, we focus on this limit case throughout the paper, and each reference to the limit is in this sense. This enables simpler expressions and clearer intuition for our results, even though our formal results apply to the discrete model.[6]

We denote the state of the world by $\theta = (\theta_1, \ldots, \theta_n)$ and the set of possible states of the world by $\Theta = \mathbf{R}^n$. Since d_0 is the only option whose outcome is fixed, we refer to it as the *default option* and its outcome $X(d_0) = 0$ as the *default outcome*. It is sometimes convenient to make the dependence on θ explicit, so we also write $X(d; \theta)$ to denote the outcome of option d in state θ.

Preferences: As is standard in games of verifiable information, the receiver has preferences over outcomes and the sender over the option chosen. The receiver has an ideal outcome $b > 0$, and for transparency, in the main text, we focus on the functional form $u_{\mathrm{R}}(x) = -(x - b)^2$, where x is the realized outcome. With the sole exception of Corollary 1, all of our results are proven under more general utility functions, and general receiver utility does not qualitatively change our results. We refer to Appendix A for details.

[6] We work with the finite option space to deal with technical issues that emerge at the limit, in particular in the formation of off-path beliefs. It has been known generally since the work of Seidmann and Winter (1997) that the existence of a 'worst option' for the sender is important for the existence of equilibrium in games of strategic communication. Finiteness ensures that such an option always exists, regardless of the message sent, and that is attractive to the receiver.

To see the technical issue, suppose that the option space is the interval $[0, 1]$, that $b = 1$, and consider the existence of a fully revealing equilibrium (for example). The message m that communicates only the hard information "$X(1) = -1$" is off equilibrium path. Suppose that, upon receiving m, the receiver chooses option d. Option 1, the worst possible for the sender, is not as attractive for the receiver as option 0, so we must have $d < 1$. Now consider a state such that the receiver gets his ideal outcome only for options in $(d, 1)$ and that is consistent with message m. In such a state, the sender is better off concealing information and sending m instead, which rules out the existence of a fully revealing equilibrium. This argument is quite general.

The sender's utility is strictly monotonic and declining in the choice of option; an example is the linear form $u_S(d) = -d$. The assumption of decreasing utility is more than a normalization given the default option is fixed. We discuss this issue and explore the opposite case in which the sender prefers larger options in Appendix A.

Information: The sender is an expert. She is privately informed about state θ, and, thus, outcomes $X(d_i)$, for $i = 1, \ldots, n$. We often refer to θ as the sender's type. All other information is public, including the outcome of the default option d_0 and parameters μ and σ.

In accordance with Weber (1922), the informational advantage of the expert is significant. The advantage is n distinct pieces of information, although the significance of the advantage depends on how correlated the information is. We define the *complexity* of expertise by the ratio σ^2/μ, which also reflects the complexity of the underlying decision problem. As $\sigma^2/\mu \to 0$ the receiver can infer much about the environment from his knowledge of the default option, whereas as $\sigma^2/\mu \to \infty$ such inference is increasingly useless.

To ensure that communication is non-trivial, we impose the following minimum bound on the extent of the preference misalignment between the sender and receiver.[7]

$$b > \frac{\sigma^2}{2\mu}.$$

Communication: For each option, the sender has a piece of hard information that verifies its outcome. The sender can hide or reveal any number of these pieces of information but she cannot fake them. Her message is formally described by a mapping $m : \mathcal{D} \to \mathbf{R} \cup \{\emptyset\}$, where $m(d) = \emptyset$ if she hides the outcome of option d, otherwise, she reveals it and $m(d)$ must be the outcome of option d.[8] We denote the set of messages by $\mathcal{M}$.

Timing: First, nature draws the outcome function—the state of the world—and the sender learns the realization. Second, the sender sends her message.

[7]Without this restriction, the sender trivially gets her first best for every n large enough: It is an equilibrium for the sender not to communicate any information and for the receiver to choose the default option (see Section 4.1).

[8]This does rule out vague communication, such as when the sender reveals an outcome is in some range, although this restriction is immaterial to our results. If we allow for a "rich language" in the terminology of Seidmann and Winter (1997), then any equilibrium in our setting continues to be an equilibrium with the general message space, as we prove formally in Appendix A.

Third, the receiver updates his beliefs and chooses one of the options. Finally, the sender and the receiver realize their payoffs and the game ends. Note that neither party has commitment power: The sender cannot commit to a message rule and the receiver cannot commit to a decision rule.

Solution Concept: A strategy for the sender is a mapping $M : \Theta \to \mathcal{M}$. A strategy for the receiver is a mapping $D : \mathcal{M} \to \mathcal{D}$. The receiver's beliefs are described by a belief mapping $B : \mathcal{M} \to \Delta(\Theta)$ that assigns a probability distribution over states to every possible message. The solution concept is perfect Bayesian equilibrium. Throughout, $\mathbb{E}$ denotes the expectation operator under the original state distribution, and $\mathbb{E}^{B(m)}$ the expectation operator when states are distributed according to belief $B(m)$.

Terminology: The focus of our analysis is the amount and the nature of information communicated in equilibrium. To that end, some terminology is helpful. An *interaction* between sender and receiver is the message sent and the option chosen. This is the pair of observable actions, the advice provided by the mechanic about potential repairs, and the repair chosen by the customer. The mechanic makes a *recommendation* if one of the repairs he describes is chosen by the customer. If a recommended repair is all he says, the interaction is *conative* (by convention, we include in this definition the recommendation and choice of the default option). Advice beyond the recommended repair is *referential*. We say that an interaction is *prescriptive* if it contains a recommendation—whether it contains referential advice or not—otherwise the interaction is *non-prescriptive*.[9] An equilibrium is referential/prescriptive/etc. if equilibrium interactions are referential/prescriptive/etc.

Our model has many equilibria, some of which differ only in minor and extraneous details of the sender strategy. To simplify the statement of our results, we say that two equilibrium profiles are *equivalent* if they generate the same mapping from sender type to receiver choice.

So that comparisons of equilibria are not dependent on the functional form of utility, we evaluate welfare properties via dominance. We distinguish between a strict and a weak form of dominance.

Definition 1 *An equilibrium* (M, D, B) weakly (resp. strictly) dominates *an equilibrium* (M', D', B') *in state* θ *if* $D(M(\theta))$ *is no greater than (resp. less than)* $D'(M'(\theta))$.

[9]Thus, a non-prescriptive interaction is necessarily referential other than if no advice is offered.

We then say an equilibrium weakly (strictly) dominates another if it weakly (strictly) dominates on a probability one set of states of the world. When this condition holds, therefore, the sender prefers the dominant equilibrium for *any* (decreasing) utility function.

Finally, to study an equilibrium property with n large, we consider not just one equilibrium, but a sequence of equilibria—one equilibrium for each value of n. Throughout the paper, sequences of equilibria implicitly refer to sequences indexed by n, the dimension of the state space. A property holds "in the limit" when the probability that such a property holds for given n converges to one as n goes to infinity.[10] All convergence statements refer to convergence in probability.

Off-Path Beliefs: The richness of the type and message spaces ensures there are many off-path messages available to the sender, all of which require the specification of beliefs for the receiver.[11] Moreover, for any given equilibrium interaction, there are many off-path beliefs that can support the equilibrium. We refer to these as *suspicious* beliefs.

To define these beliefs, some additional notation is required. We say that a state θ and a message m are *compatible* if the hard information included in m does not rule out θ. For a message m, $\Gamma(m)$ represents the set of all the states that are compatible with m.

Definition 2 *Given a strategy profile* (M, D, B)*, off-path beliefs are* suspicious *when for every* $m \notin M(\Theta)$*,*

$$\max_{\substack{m' \in M(\Theta) \\ \Gamma(m) \cap \Gamma(m') \neq \emptyset}} D(m') \leq \max\left(\arg\max_{d \in \mathcal{D}} \mathbb{E}^{B(m)}[u_{\mathrm{R}}(X(d))] \right).$$

The definition is complicated although its intent is simple: It ensures that all deviations off the equilibrium path are potentially unprofitable. It stipulates that for every off-path message m, beliefs are such that at least one optimal response is possible by the receiver that makes the sender weakly worse off, regardless of which on-path message m' the receiver was expected to play in equilibrium.

[10] For example, if we discuss an equilibrium defined for every n and say that equilibrium interaction is conative in the limit as n goes to infinity, the corresponding fully formal statement is that the probability that equilibrium interaction is conative converges to one as n goes to infinity.

[11] We use the following terminology: A message is "on path" if it is communicated in some state, otherwise it is "off path." The set of on-path messages associated with sender strategy M is thus $M(\Theta)$.

Suspicious beliefs imply the deviation may be unprofitable. To ensure it is unprofitable it must be that the receiver chooses one of the options that leaves the receiver weakly worse off. When the receiver does so, we say that his off-path decisions are *suspicious*. Formally, off-path decisions are suspicious if for all $m \notin M(\Theta)$ and all states $\theta \in \Gamma(m)$, $D(m) \geq D(M(\theta))$. Combining these two requirements, we have the following immediate result.[12]

Lemma 1 *Off-path beliefs and decisions are suspicious in all equilibria.*

Beliefs are suspicious in the sense that the receiver presumes that the sender deviates to avoid revealing unfavorable information. A receiver is, of course, entitled to more credulous beliefs but credulous beliefs will not support equilibrium. Milgrom (1981) shows in his canonical model that the unique equilibrium requires maximally suspicious beliefs, what he labels skeptical beliefs. The equilibria here do not always require such extreme beliefs, and for many off-path messages they can involve little or no suspicion at all. In Appendix B, we give general conditions on the existence of suspicious off-path beliefs.

This richness means that, unlike Milgrom (1981), the equilibrium interactions we identify can be supported by many different off-path suspicious beliefs. In fact, any interaction that can be supported as an equilibrium for less suspicious beliefs can also support an equilibrium for more suspicious beliefs. To avoid equilibrium multiplicity that is not action or outcome relevant, whenever possible we refer to the class of equilibria that differ only in off-path suspicious beliefs and decisions as a single equilibrium.[13]

3 Preliminaries

3.1 Informational Spillovers and On-Path Beliefs

In choosing her message, the sender must consider not only what she reveals but how that information will spillover into the receiver's beliefs about other options. The nature of the spillover is important to the existence of equilibrium. To understand the forms it can take, begin by considering the following example.

Example 1 *The sender's strategy is reveal $X(d^a)$ if $X(d^a) < b/2$, otherwise reveal $X(d^b)$, for some $d^b > d^a > 0$.*

[12] All omitted proofs can be found in Appendix C.

[13] Multiplicity may also arise due to ties as then the receiver's best response is not unique. As ties occur for prescriptive interactions with zero probability, we assume without loss of generality that when indifferent, the receiver chooses the option preferred by the sender.

The spillover of information in this example is both *direct* and *indirect*. The direct component is unavoidable, and comes purely from knowledge of a particular point. In revealing the outcome of option d^b, the sender conveys that the outcome function passes through that point and this knowledge shapes beliefs in either direction. The indirect component is what the receiver can infer in addition to that knowledge. In the example, revealing the outcome of d^b also reveals information about the option d^a, namely that the outcome is larger than $b/2$, and this indirect knowledge also informs beliefs.

Indirect informational spillovers complicate the receiver's inference problem. He must ask himself: What information did the sender reveal and *why* did she reveal it? The answer depends on the strategy used by the sender and leads to complicated beliefs that generally do not permit closed form representations.

The receiver's inference problem is considerably simpler if informational spillover is only direct. In this case the receiver need not ask the *why* question. He can take the information at face value as it does not depend on the strategy used to reveal it. Fortunately, strategies exist that have only direct spillovers and the class of such strategies play an important role in our analysis. The following is an example.

Example 2 *The sender's strategy is reveal $X(d^a)$, for some $d^a > 0$.*

The spillover from this strategy is only direct as the decision to reveal the outcome of d^a is independent of not only the other options but of the outcome of d^a itself. Without indirect spillover, the receiver's beliefs condition only on the hard information. When this holds we say that beliefs are *neutral*.

Definition 3 *Given a message m and an associated belief $B(m)$, we say that $B(m)$ is* neutral *when $B(m)$ is the original distribution over states θ conditional on $\theta \in \Gamma(m)$.*

Thus, $B(m)$ is neutral when it is equal to the state distribution conditional on the hard information in m. In the context of Gaussian outcome paths, neutral beliefs yield particularly simple and tractable functional forms. If d^r is the right-most known option, beliefs for options $d > d^r$ are normally distributed with mean

$$\mathbb{E}[X(d) \mid X(d^r)] = X(d^r) + \mu|d - d^r|, \tag{1}$$

and variance

$$\mathrm{Var}[X(d) \mid X(d^r)] = \sigma^2|d - d^r|. \tag{2}$$

These beliefs follow immediately from the law of motion that defines the distribution of $X(d)$, and represent an intuitive extrapolation from what is

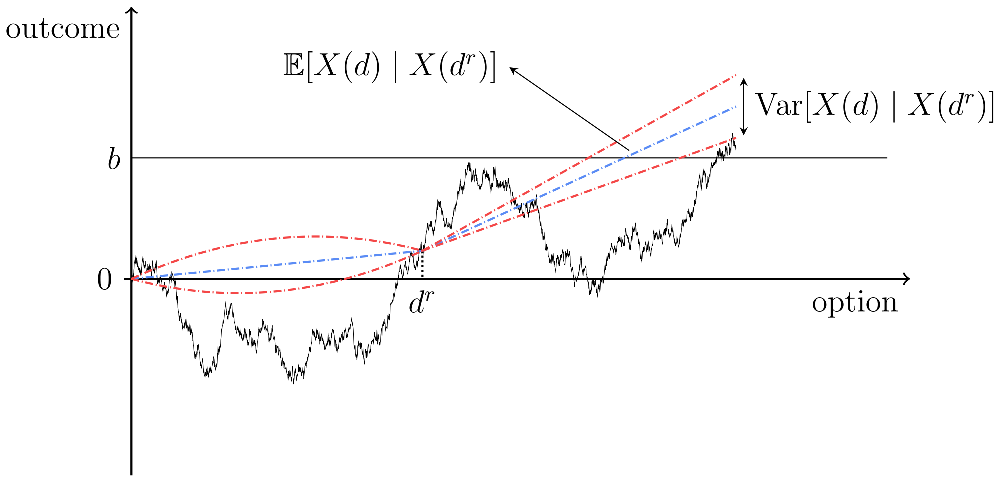

Figure 2: Neutral beliefs with knowledge of option d^r.

known. The drift term, μ, measures the expected rate of change with the variance of beliefs increasing linearly in distance from d^r, capturing the idea that beliefs are more uncertain the further an option is from what is known. The receiver's beliefs are neutral at the beginning of play, with the default option, d_0, providing the anchor point.

For options between two known points, $d \in (d^l, d^r)$, a Gaussian bridge forms. Beliefs are again normally distributed with mean

$$\mathbb{E}\big[X(d) \mid X(d^l), X(d^r)\big] = X(d^l) + \frac{d - d^l}{d^r - d^l}\big(X(d^r) - X(d^l)\big), \tag{3}$$

and variance

$$\mathrm{Var}\big[X(d) \mid X(d^l), X(d^r)\big] = \frac{|d - d^l| \cdot |d - d^r|}{d^r - d^l}\sigma^2. \tag{4}$$

Equations (3) and (4) follow from the projection formulas for jointly normal random variables. The expected outcome is now a simple interpolation of the two neighboring points and independent of the drift. The variance is a concave function across the bridge, reaching a maximum in the middle and approaching zero in the ends. By the Markov property, beliefs depend only on the nearest known points (as they did for $d > d^r$ above). In the limit as n goes to infinity, the beliefs in both cases are depicted in Figure 2.

3.2 Non-deceptive Strategies and Neutral Beliefs

Neutral beliefs, and the absence of indirect informational spillovers, are important because they connect directly to the receiver's inference and the sender's credibility. If a message is effective at dissuading and persuading the receiver, the expert will want to send that message as often as she can. But the more often she does so, the less effective the message is as the credibility of the message itself will be undermined. Successful persuasion, therefore, requires that a message be difficult to imitate.

One way to avoid imitation is to make it impossible. For some strategies, the sender has only a single on-path message that is consistent with her hard information. Full revelation is an example, although there are others. In this class of strategies, incentive compatibility across on-path messages is not so much satisfied as it is rendered moot. The sender either sends the one on-path message available to her or she deviates off the equilibrium path. This means that all deviations are detectable by the receiver and the sender cannot deceive the receiver about her type. We refer to this class of strategies as *non-deceptive*.

Definition 4 *A sender strategy is* non-deceptive *if, in every state of the world, any deviation by the sender can be detected by the receiver. Otherwise, the sender strategy is* deceptive*. An equilibrium is non-deceptive if the sender strategy is non-deceptive.*

The inability to deceive does not necessarily imply a full separation of types, which in our setting would require full revelation.[14] The message available to a type may be a message that many types can send, and upon observing the message the receiver does not know for sure which type he is facing, but he knows that it could only be sent by the types that the strategy dictates send it.[15]

The connection of non-deceptive strategies to neutral beliefs is that they are effectively the same constraint. A non-deceptive strategy generates neutral beliefs, and neutral beliefs imply that the strategy is non-deceptive with probability one. The link connecting these ideas is a lack of conditionality. To not have a choice in on-path messages means that the strategy cannot condition on information that is not revealed. To make this connection tight, we need to allow for zero probability events. We say that a sender strategy is *almost*

[14] In cheap-talk games, no strategy is non-deceptive other than a fully pooling strategy.

[15] An appealing feature of non-deceptive strategies is that they are, in a sense, strategically simple. The sender sends the message that is expected of her or she deviates off path. The inability to deceive also reflects a form of trust and a plain-spoken style of communication that is observed in practice in many relationships.

non-deceptive when, for every message m, the probability that the sender sends m is one, conditionally on the state being compatible with m. We then have the following equivalence.

Lemma 2 *In any equilibrium, the sender strategy is almost non-deceptive if and only if the on-path receiver beliefs are neutral.*

A non-deceptive strategy ties the hands of the sender. It takes away the sender's ability to strategize in her choice of message without being detected by the receiver. This limits, to be sure, what can be communicated in equilibrium, but it guarantees that what is communicated is credible. Moreover, the upside is that it does so in a way that limits informational spillover to only be direct, thereby allowing the sender to control her information to a greater degree.

The class of non-deceptive strategies will play an important role in our analysis. They are not the only strategies that can support equilibria. But they are sufficient to establish that referential advice can deliver leverage to the sender relative to providing a recommendation alone.

4 Conative Advice

4.1 The First-Point Conative Strategy

A natural strategy for the expert is to simply offer a recommendation and nothing more. One way to do this is to reveal a point that achieves a minimum standard, such as an outcome within some fixed distance of the receiver's ideal. This cannot be done with a non-deceptive strategy and neutral beliefs.[16] But it can be done in a way such that the recommendation conditions only on information in one direction. We refer to this as a first-point conative strategy.

Definition 5 *The sender follows a* first-point conative strategy *when, for some $\Delta \geq 0$, the sender reveals the smallest option whose outcome falls in the range $[b - \Delta, b + \Delta]$ and if no such option exists, the sender reveals everything. A* first-point conative equilibrium *is an equilibrium where the sender follows a first-point conative strategy and the receiver implements the recommendation.*

An example of a first-point conative strategy is depicted in Figure 3, where d^c is the recommendation. In the limit, the revealed point will be exactly at the lower boundary, $b - \Delta$. The only exception is when the mapping fails to reach

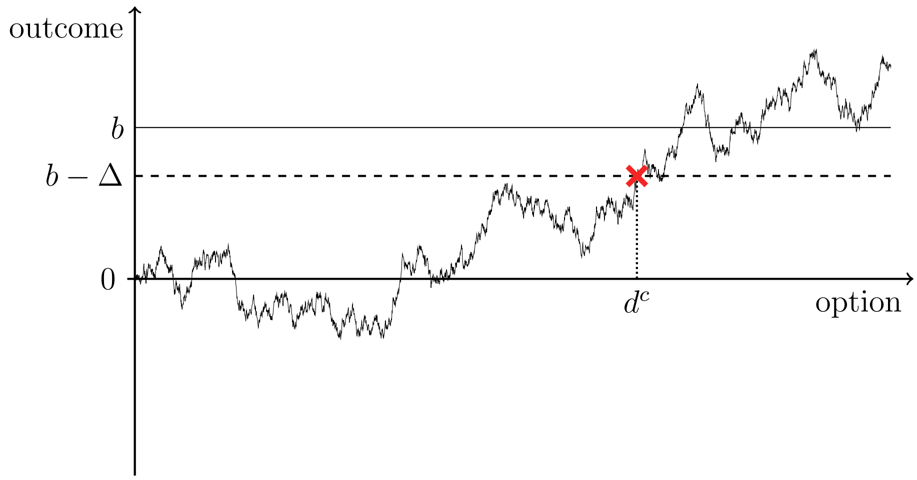

Figure 3: The first-point conative strategy.

this outcome and all information is revealed, although this is a zero probability event.[17]

It is important that in this strategy it is the first point to meet the standard that is revealed. It is important because it ensures on-path incentive compatibility for the sender. The strategy is deceptive and, thus, non-detectable deviations exist. For instance, the sender could reveal the second point that meets the standard and not be detected. Yet, critically, none of these deviations are to the left of the recommendation and profitable for the sender. The first point is important too because the recommendation conditions only on information to the left. Thus, there are no indirect spillovers to the right and beliefs are neutral in that direction.[18] This illustrates how limiting deviations and opportunities to deceive can make incentive compatibility easier to satisfy.

Proposition 1 establishes that the first-point conative strategy can support an equilibrium. It shows, however, that to satisfy incentive compatibility for the receiver, the recommendation must not be too far from the receiver's ideal outcome. Defining:

$$\Delta^{\max} = \frac{\sigma^2}{2\mu} + \frac{\mu}{2\sqrt{n}}, \tag{5}$$

we have the following.

[16]As the strategy must condition on the outcome of the revealed option.

[17]For the discrete space of options, the mapping may jump over the boundary, and it is for this reason that the strategy is defined as a band. As the number of options increases, the probability the mapping does not fall into the band goes to zero.

[18]This is a slight abuse of our terminology in that it applies neutral beliefs directionally.

Proposition 1 *A first-point conative equilibrium exists if and only if $\Delta \leq \Delta^{\max}$.*

The existence of the first-point conative equilibrium confirms that communication need not be complex simply because expertise is complex. Moreover, it shows that an expert can use her expertise whilst protecting it to a maximum degree. Yet is also shows that her ability to persuade the receiver is limited when she does so.

The first-point conative equilibrium that is optimal for the sender is at $\Delta^{\max}$ when the band is the widest. This equilibrium delivers both the smallest option when a point falls within the band and it minimizes the probability that the sender is compelled to fully reveal. This is the receiver's least preferred of these equilibria, although even in this equilibrium the receiver extracts value from the expert. He faces no uncertainty in accepting the recommendation and he obtains an outcome not too far from his ideal.

This raises the question of why. Why must the expert concede so much to the receiver, yet, at the same time, not concede everything such that the receiver obtains his ideal outcome?

The answer is informational spillovers. To persuade the receiver to accept the recommendation, the sender must dissuade him from choosing other, unrevealed options. Unfortunately for the sender, the informational spillover from her recommendation helps the receiver. In fact, a positive complementarity develops such that the better is the recommendation for the receiver—the more persuasive it might be—the more attractive are other options and the less the recommendation dissuades. It is only when the recommendation is not too distant from the receiver's ideal that persuasion dominates and the receiver is willing to accept the recommendation.

To see this, we must understand how the receiver evaluates unrevealed options. To the right of the recommendation the receiver faces a simple risk-return trade-off. This is easiest to see in the limit with the Brownian motion. Experimenting with a larger option increases the expected outcome linearly at rate μ but it also increases the variance at rate σ^2, as given earlier by Equations (1) and (2). How the receiver trades-off risk against return depends on the shape of his utility function. The quadratic-loss functional form delivers a particularly sharp answer to this trade-off as it admits a separable mean-variance representation. For a recommendation d^c, we have for all $d \geq d^c$ that:

$$\mathbb{E}[u_{\mathrm{R}}(X(d) \mid X(d^c))] = -(\underbrace{b - (X(d^c) + \mu(d - d^c))}_{\text{expected distance from } b})^2 - \underbrace{\sigma^2(d - d^c)}_{\text{variance}}. \quad (6)$$

Differentiating delivers the limit threshold $\frac{\sigma^2}{2\mu}$ from Proposition 1. This thresholds represents the boundary on whether the risk of experimenting is worth the return. For outcomes closer than this to b, the receiver accepts the recommendation and does not experiment. He considers such outcomes "good enough" even though he does not obtain his ideal outcome.

Beyond this boundary, outcomes are not "good enough," and the receiver overrides the recommendation and experiments. The return is then worth the risk. In this case the threshold $\frac{\sigma^2}{2\mu}$ represents the expected outcome of his optimal experiment. That is, the receiver is willing to trade-off increased risk for a better expected outcome up until he reaches this threshold, such that he chooses the risky option $d^e > d^c$ with:

$$\begin{aligned}\mathbb{E}[X(d^e) \mid X(d^c)] &= b - \frac{\sigma^2}{2\mu}, \\ \mathrm{Var}[X(d^e) \mid X(d^c)] &= \left(\left(b - \frac{\sigma^2}{2\mu}\right) - X(d^c)\right)\frac{\sigma^2}{\mu}.\end{aligned} \tag{7}$$

This implies that, regardless of the quality of a recommendation, if it is beyond the "good enough" threshold, the receiver's optimal experiment will set the expected outcome equal to that threshold. This does not imply, however, that the receiver's expected utility is unaffected by the recommendation. Observe that the variance of his optimal experiment is decreasing in $X(d^c)$, the recommended outcome. This is because the better is the recommendation, the less boldly the receiver has to experiment, and the less risk he has to endure.

This represents the positive complementarity that plagues the expert. The better is the recommendation for the receiver, the better also are unrevealed outcomes. In revealing an attractive outcome, the sender also reveals that even better outcomes are likely nearby, and this emboldens the receiver to continue to experiment. The informational spillover from the recommendation undermines the recommendation itself. The more the recommendation persuades, the less it dissuades. Proposition 1 establishes that this complementarity only exhausts itself when the recommendation delivers an outcome that the receiver considers "good enough." In equilibrium, therefore, the receiver is able to capture the value of expertise.

The receiver may, of course, choose an option to the left of the recommendation. In this case the informational spillover is both direct and indirect and beliefs are not neutral. In recommending the first option to meet a standard, the sender is also revealing that options to the left don't meet the standard. This indirect spillover, however, shows that these options are even less attractive to the receiver than the recommendation itself, and so are not tempting to him.

To the left of the recommendation, therefore, the indirect spillover helps the sender dissuade the receiver.

To complete the logic of equilibrium in Proposition 1 we must consider deviations by the sender off the equilibrium path. Full revelation cannot be profitable as the receiver would then choose the same or a larger option. This leaves partially revealing deviations. To be profitable the sender cannot reveal the equilibrium recommendation nor any option that does as well or better for the receiver. The receiver will be suspicious upon seeing out-of-equilibrium advice, and it is easy to construct beliefs that cause him to choose the equilibrium recommendation or a larger unrevealed option, thereby supporting equilibrium behavior.

The power of suspicious beliefs supports the existence of multiple first-point conative equilibria. In each equilibrium the receiver expects advice of a certain form and he will be suspicious if it is different. This expectation binds the sender, compelling her to provide advice of the expected form, even if this information is more favorable to the receiver than required in other equilibria. Although this expectation can, in this case, hurt the sender, we will see that in other contexts it can be used to help the sender.

4.2 Conative Advice: The General Case

The first-point conative strategy is just one way to generate conative advice. There are other ways to do so, and many of these ways support equilibrium behavior. Nevertheless, in every conceivable conative strategy, the sender faces similar constraints as in a first-point conative strategy and the sender's ability to persuade the receiver is limited. Theorem 1 characterizes exactly what is possible in equilibrium with conative strategies. It shows, in fact, that the sender can do no better than she can in her most favorable first-point conative equilibrium.

Theorem 1 *For all sequences of equilibria that are conative in the limit, the equilibrium outcomes are in the range $[b - \frac{\sigma^2}{2\mu}, b]$ in the limit. Moreover, for any $x \in [b - \frac{\sigma^2}{2\mu}, b]$ there exists a sequence of equilibria that are conative in the limit in which the outcomes converge to x.*

Establishing this result formally requires us to consider the full family of conative interactions. We show that for a conative interaction to support an equilibrium, the sender's strategy must satisfy a generalization of the first-point conative strategy in Definition 5. As in that definition, the sender must reveal the first option that meets some absolute standard, although that standard need not be a fixed band around b. The standard can vary in the option itself, it can be non-monotonic, it can even have multiple thresholds such that the standard isn't only a single band. Nevertheless, for all these possibilities, we show that any sequence of equilibria that is conative in the limit, the sender can move the outcome no further than $\frac{\sigma^2}{2\mu}$ from the receiver's ideal outcome.

The logic for this result follows from that for the first-point conative strategy if the sender is using a first-point style strategy for all states. As a first-point strategy generates neutral beliefs to the right, a recommendation will only be acceptable to the receiver if it is "good enough" and within $\frac{\sigma^2}{2\mu}$ of his ideal outcome.[19]

The final piece of intuition for Theorem 1 is why the sender's strategy must be of the form of a first-point strategy with neutral beliefs to the right of the recommendation. The only alternative for the sender to do better is to use a deceptive strategy that generates indirect spillovers to the right of the recommendation, and to do so in a way that the non-neutral beliefs are able to dissuade the receiver from experimenting. Unfortunately for the sender, this is not possible in equilibrium. To see why, suppose she recommends option d^c with outcome $X(d^c) < b - \frac{\sigma^2}{2\mu}$ and that this paints a picture of the world such that experimenting to the right is unappealing, so much so that the receiver is willing to accept the recommendation. For this to be possible, some sender types who could send the message don't, and they must choose not to because of information to the right of the recommendation. This is optimal only if these sender types recommend an option to the left of d^c. But if all types who could send this message instead of d^c do so, beliefs at d^c must be neutral. Non-neutral beliefs cannot be sustained in equilibrium here precisely because messages are too easy to imitate. Every sender type who can send a message does so unless they can send an even better message. As this selection conditions only on information to the left of the recommendation, the credibility of a message is undermined if the informational spillover to the right of the recommendation is indirect. In this case, the sender's freedom to deceive undermines her ability to persuade.

[19]Along the sequence of equilibria for finite n, the sender may use a first-point strategy for some states but not for others, creating the complication that the receiver's beliefs need not be neutral, even in the states where the interaction is conative (e.g., if the same recommendation is made in a different state as part of a deceptive strategy). Nevertheless, we show that beliefs necessarily become approximately neutral as n grows large, such that the receiver will implement the recommendation if and only if it is within the range that the receiver considers good enough.

4.3 The Value of Conative Advice

It follows from Theorem 1 that the *sender-optimal conative equilibrium* is the first-point conative equilibrium in which $\Delta = \Delta^{\max}$. In the limit, the outcome of this equilibrium converges almost surely to $b - \frac{\sigma^2}{2\mu}$.

This outcome is better for the sender than outcome b, yet it does not necessarily correspond to a better option than would be chosen in her absence. Without an expert, the receiver experiments on his own, following the logic of Equation (7). In this "no-advice" benchmark, the receiver's choice, d^{na}, has expected outcome $b - \frac{\sigma^2}{2\mu}$, and corresponds to the point where the drift line from the default option intersects this threshold, as depicted in Figure 4.

It is striking that this is exactly the outcome produced by the sender-optimal conative equilibrium. The difference is that this outcome is obtained with certainty in the conative equilibrium whereas in the no-advice benchmark it is obtained only in expectation. Relative to no-advice, therefore, the sender-optimal conative equilibrium will sometimes pull decisions to the left of d^{na}, to option d^r for the mapping in Figure 4, and other times it would push the decision to the right, depending on the realization of the path. The net effect is particularly stark when the receiver has quadratic utility. In this case, the expected decision in the sender-optimal conative equilibrium is exactly equal to the actual decision in the absence of advice. Consequently, for all conative equilibria that converge to a different outcome, the expected outcome is strictly larger than it is without the expert.[20]

Corollary 1 *In any sequence of equilibria that are conative in the limit, the expected equilibrium decision is no smaller than d^{na} in the limit as $n \to \infty$. Moreover, in the first-point sender-optimal equilibrium, the expected equilibrium option converges to d^{na} as $n \to \infty$.*

This says that a sender with linear utility can gain nothing from her expertise with conative advice, and that for many conative equilibria she will be strictly worse off. This conclusion holds with even greater force for a sender with concave utility. A sender with convex utility must temper this loss against the gain in utility from the dispersion of the options chosen.

5 Expert Power: Can the Sender do Better?

In conative equilibria the receiver extracts all of the surplus from expertise. Not only does he benefit from the removal of uncertainty in the outcome, but

[20] With general receiver utility this comparison is not strict, yet the sender's weakness in influencing the decision with conative advice persists. See Appendix A for a discussion.

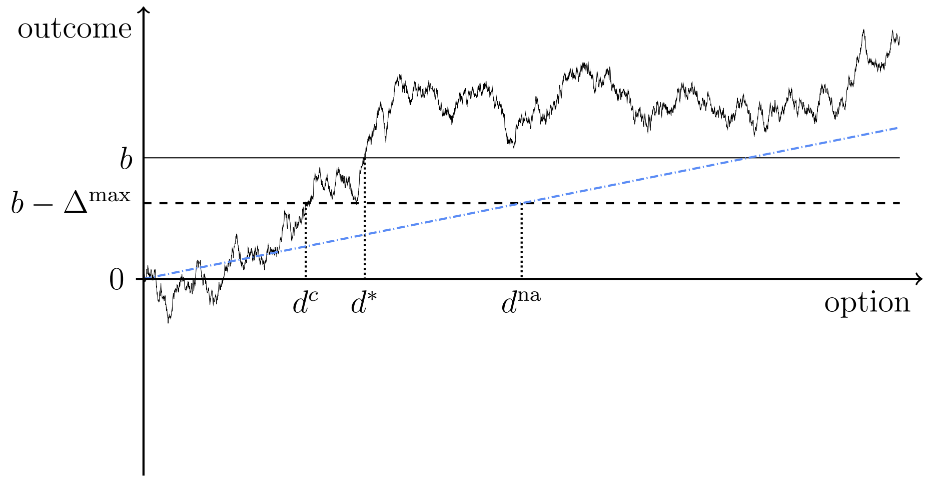

Figure 4: The "no-advice" benchmark, d^{na}.

the outcome itself shifts toward his ideal point in all but the sender-optimal conative equilibrium. Scope for the sender to do better exists. The question we address in this section is whether this is possible and, if so, how much of the surplus the sender can extract.

A natural approach is for the sender to say less. Less than conative is to say nothing, and this forfeits the sender's opportunity to influence the receiver's decision. Saying nothing is also not an equilibrium strategy. "No-advice" is a non-deceptive strategy and the receiver forms neutral beliefs. He faces, therefore, the same decision as if the expert were not present. His optimal choice is to experiment on his own with option d^{na} that has an expected outcome of $b - \frac{\sigma^2}{2\mu}$. A profitable deviation exists for sender types in which the mapping crosses b before d^{na} and not again, such as in Figure 4. By fully revealing the mapping, the sender renders off-path beliefs moot and this compels the receiver to choose the crossing point d^*. This makes the sender better off and, therefore, the strategy of no-advice is not an equilibrium.[21]

The other route for the sender is to say more. It is easy to see that the addition of referential advice by itself need not change equilibrium behavior and deliver leverage to the sender. For instance, each conative equilibrium can be amended to include all information to the left of the recommendation. Because the receiver infers in equilibrium that these options are less attractive than the recommendation, and because information from these options does not spillover to the right of the recommendation, revealing these points does

[21]The use of full revelation as a deviation in this way points to the general difficulty in supporting non-prescriptive equilibria. We take up this issue in Section 5.4.

not change the receiver's optimal decision. The referential equilibrium that this produces is exactly equivalent in outcomes to the conative equilibrium that it extends and, thus, leaves the sender neither better nor worse off.[22]

Indeed, it is possible that referential advice hurts the sender. Another referential equilibrium is for the sender to fully reveal her information. With knowledge of the complete mapping the receiver will obtain his ideal outcome, choosing one of the options where the mapping intersects with b. Even in the best case for the sender in which the option chosen is the first crossing point, this is equivalent to the worst conative equilibrium for the sender and leaves her worse off than in all other conative equilibria.

To help the sender, therefore, referential advice needs to be constructed in just the right way. Enough information to influence the receiver's decision but not too much information such that the information serves the receiver's interests rather than the sender's.

5.1 The Interval Strategy

In the referential extension of conative equilibria, the referential advice is not outcome relevant because it is not decision relevant. The additional outcomes revealed are dominated by the recommendation itself and they do not change beliefs about the remaining unrevealed options. For referential advice to be impactful, it must be that it is decision relevant, and for it to improve upon conative equilibria, it must include information to the right of the recommendation. To that end, consider the following strategy.

Definition 6 *The sender follows the* interval strategy *when the sender reveals the outcomes of options* $\{d_0, d_1, \ldots, d^r\}$*, where* d^r *is the smallest option that satisfies*

$$\max_{d \leq d^r} u_{\mathrm{R}}(X(d)) \geq \max_{d > d^r} \mathbb{E}[u_{\mathrm{R}}(X(d)) \mid X(d^r)], \tag{8}$$

with $d^r = d_n$ *if no such option exists. An* interval equilibrium *is an equilibrium where the sender follows the interval strategy and the receiver implements his most-preferred revealed option.*

The interval strategy represents a stopping rule. In the limit it takes the following form. Starting at the default option, the sender reveals points until the path hits a lower threshold. This threshold is a function of the peak revealed so far, increasing in the height of the peak. Eventually the lower threshold

[22] The referential extension of the conative strategy is non-deceptive, in contrast to the conative strategy itself.

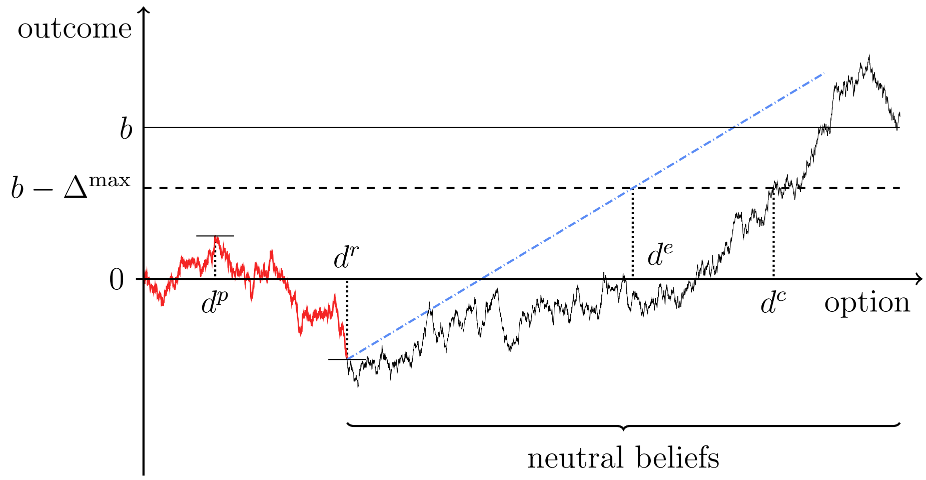

Figure 5: The interval strategy hitting the lower threshold.

equals the peak—at what can be thought of as an upper threshold—and, when it does, the sender necessarily stops revealing. If neither threshold is ever met, the sender fully reveals. The interval strategy is non-deceptive and, thus, beliefs for unrevealed options are neutral.

To understand the thresholds, consider the path depicted in Figure 5 that hits the lower threshold at d^r. Naturally, the recommendation is the option d^p that obtains the peak (the option that is receiver optimal among those revealed) and not the right-most revealed option, d^r. This means that the referential advice *is* decision relevant for the receiver. The points revealed to the right of d^p do not change the receiver's beliefs about d^p itself, but they do change the receiver's beliefs about unrevealed options further to the right. They dissuade only. Nevertheless, by changing the attractiveness of further experimentation, the referential advice changes the relative appeal of the recommendation, and, potentially, the receiver's decision.

By construction, the referential advice is worse for the receiver than the recommendation itself. This matters because of the positive complementarity between information revealed in a recommendation and the receiver's utility from experimenting. In the conative equilibria, this complementarity undoes the power of the sender as the better the outcome that she reveals is, the more emboldened is the receiver to experiment.

In the interval strategy this complementarity is turned around so that it works to the sender's benefit. By revealing bad outcomes, the sender makes experimenting less appealing to the receiver, which, in turn, makes him more likely to accept a recommendation that he otherwise wouldn't. The key is the disconnect between the spillover and the recommendation. With only a recommendation these forces are inseparable. When separated with referential advice, the recommendation and the spillover can be made to work in opposite directions and this enables the sender to influence the receiver's decision to a greater extent.

The lower threshold in the interval strategy is defined as exactly the point at which the receiver is indifferent between accepting the recommendation and experimenting further (the left and the right-hand side of Equation (8), respectively). For the limiting case as n grows large, this becomes:

$$-(b - X(d^p))^2 = -\left(\frac{\sigma^2}{2\mu}\right)^2 - \left(b - \frac{\sigma^2}{2\mu} - X(d^r)\right)\frac{1}{\mu}\sigma^2, \tag{9}$$

where the right-hand side follows from the mean-variance representation in Equation (6) and is the expected utility from option d^e in Figure 5.

This simplifies to:

$$b - X(d^r) = \frac{\mu}{\sigma^2}(b - X(d^p))^2 + \frac{\sigma^2}{4\mu}. \tag{10}$$

As $X(d^p)$ increases, the right-hand side decreases, and, therefore, so too must $X(d^r)$. The threshold increases in the peak and meets the peak when $X(d^p)$ equals $b - \frac{\sigma^2}{2\mu}$. This is exactly the same outcome as defines the sender-optimal conative equilibrium, and marked in the figure by option d^c. If the upper threshold is reached, revelation stops. The recommendation is the right-most revealed option itself and additional dissuasion is not needed.

Proposition 2 establishes that the interval strategy supports an equilibrium. This equilibrium benefits the sender. At worst, the decision is the same as in the sender's most-preferred conative equilibrium. In many cases, however, the path hits the lower threshold earlier and the option chosen is strictly to the left of that chosen in all conative equilibria. As this occurs with strictly positive probability, the interval equilibrium weakly dominates the sender-optimal conative equilibrium.

Proposition 2 *The interval equilibrium exists. It strictly dominates the sender-optimal conative equilibrium with positive probability, and weakly dominates for all states.*

In dissuading the receiver from experimenting, the sender does not convince him that a good option does not exist. By construction, an option whose outcome is arbitrarily close to the receiver's ideal almost always exists if n is large enough. Rather, the receiver is dissuaded from experimenting because he is now more uncertain as to where good options lie. He knows they are to the right of d^r, but he does not know how far. And the further is the outcome of d^r from his ideal at b, the greater is his uncertainty, and the less inclined he is to experiment. Note that he still prefers to experiment than accept the outcome of d^r, but the construction of the interval strategy is such that the right-most option d^r is no longer the reference point, rather it is the peak of the interval at d^p.

5.2 Dominance of the Interval Equilibrium

Referential advice in the interval equilibrium allows the sender to paint a picture of the world that is unfavorable to the receiver, and to do so credibly. It is intuitive that this is possible some of the time, that if the path of outcomes turns down sufficiently, the sender will be able to use the interval strategy to deter the receiver from experimenting. The subtlety of the interval strategy is that this is possible not just some of the time but, in the limit, it is possible all of the time.

In Theorem 2 we establish that as n grows large, the probability approaches one that the lower threshold is reached before the decision in the sender-optimal conative equilibrium. Thus, in the limit the interval equilibrium *strictly* dominates all possible conative equilibria.

In fact, we prove a stronger result. We establish that in a broad class of equilibria, including all prescriptive equilibria that are non-deceptive, that if an equilibrium dominates all conative equilibria, then the interval equilibrium necessarily dominates it.

To formalize this result, we focus on the strategic supply of information to the right of the recommendation. We do this by presuming that strategies are fully-revealing to the left of the recommendation. We say that an interaction is *fully-revealing to the left* if the sender reveals, at least, every option from the default option up to and including the receiver decision. We then have the following.

Theorem 2 *For every sequence of equilibria $\Sigma_1, \Sigma_2, \ldots$ that are conative in the limit, the interval equilibrium strictly dominates Σ_n in the limit as $n \to \infty$. Moreover, if Σ is an equilibrium whose interactions are fully-revealing to the left with probability one and that weakly dominates the sender-optimal conative equilibrium, then the interval equilibrium weakly dominates Σ.*

This implies that not just some sender types can find a peak and a downturn that deters experimentation, but that (almost) every sender type can. These recommendations may get close to the limit threshold of $b - \frac{\sigma^2}{2\mu}$, but one can always be found. This is possible because as the peak approaches $b - \frac{\sigma^2}{2\mu}$, the value of experimenting is itself small and getting smaller. As such, the downward step needed to dissuade the receiver from experimenting also gets smaller, and the noise in the outcome path is sufficient to ensure that such a step almost surely occurs despite the drift of the path being positive.

That all types can dissuade in this way contrasts with existing intuitions. Typically if one set of types separate with a picture of the world that is unfavorable to experimentation, the receiver infers that the remaining types are favorable to experimentation. Instead of condemning these types to a worse outcome, the interval strategy resets, in a sense, at a slightly higher recommendation and allows another set of sender types to separate and paint a slightly less unappealing picture that nevertheless dissuades the receiver. Theorem 2 establishes that this process iterates and that in the limit all sender types can be separated in this way.

The class of strategies that are fully-revealing to the left is broad. It includes the interval strategy, as well as full revelation, and strategies that are deceptive to the right, amongst others. Moreover, many equilibria with strategies outside this class have equivalent equilibria within the class. As noted above, all first-point conative equilibria are equivalent to equilibria in which all information to the left of the recommendation is revealed.

In fact, this equivalence is tight for non-deceptive equilibria. We show in Appendix B (Lemma B.3) that for a non-deceptive strategy to support an equilibrium and be prescriptive, it must be that interactions are fully-revealing to the left with probability one. To be sure, there are non-deceptive strategies that are not equivalent to a strategy that is fully-revealing to the left, yet the lemma shows that such strategies cannot support an equilibrium.[23] Such strategies fail because, with neutral beliefs, full revelation is profitable for sender types that cross b at one of the unrevealed points to the left of the recommendation and never again, in the same way that full revelation undermines no-advice as an equilibrium. Thus, the class of equilibria that are fully-revealing to the left include all non-deceptive equilibria that are prescriptive. [24]

Given the focus until this point of the paper on manipulations of information to the right of the recommendation, the reader may wonder why a restriction on information to the left is necessary at all to establish Theorem 2. Throughout the

[23]For example, the strategy in which the sender always and only reveals the outcome of option d_q, for $q \geq 2$, is neither fully-revealing to the left nor equivalent to a strategy that is.

[24]We discuss non-prescriptive equilibria in the following section.

paper, information to the left of the recommendation has proven innocuous and it is manipulations to the right that have been important. The complication is exactly when information supplied to the left of the recommendation is inter-dependent with information to the right. The fully-revealing to the left restriction imposes a separation between the two sides. If we instead assume strategies such that revelation of information to the left and right of the recommendation are independent, then Theorem 2 is easily extended to *all* prescriptive strategies. We cannot construct an interdependence such that the equilibrium is not equivalent to an equilibrium that is fully-revealing to the left, let alone one that dominates the interval equilibrium, but we cannot, alas, rule it out.[25]

Simulations. We can obtain further insight into the interval strategy numerically. Figure 6 depicts the average option chosen for 10,000 simulations as the complexity of the underlying issue varies. As can be seen, the sender's leverage can be considerable and is maximized at intermediate levels of complexity. Even without advice, the sender is better off the more complex is the issue as uncertainty makes the receiver more tentative in his experiment. For low complexity (low σ^2), the receiver's choice approaches 2 and the expected outcome converges on b, whereas for high complexity (high σ^2) the experiment approaches 0 (following from the key threshold $b - \frac{\sigma^2}{2\mu}$). At either extreme the sender has little leverage. For high complexity, she has little leverage because there is simply little to leverage. The receiver is already making a very favorable decision. At the other extreme, the sender has little leverage because the outcome follows a more narrow path and there aren't the radical peaks and troughs that dissuade the receiver to a significant degree. Even in this case persuasion and dissuasion is always possible, yet the simulation shows that the power of dissuasion is small. The middle range of moderate complexity is the sweet spot for the sender and where her leverage is greatest.

The simulations also provide insight into why the receiver accepts the sender's advice, even when the sender holds so much leverage. In accepting the

[25]This seemingly represents a technical rather than substantive difficulty. The purpose of information to the right of the recommendation is to dissuade, and Theorem 2 establishes that, on its own, this cannot be used to dominate the interval equilibrium. Similarly, information to the left cannot do this and, in fact, is often unnecessary in equilibrium. The Markov property of the outcome function implies that there is no direct informational spillover from the left of a recommendation to the right. Thus, the indirect spillover—the deception—is caused by the strategy alone and not the realization of outcomes. In effect, such strategies allow the sender to randomize within a pure strategy the information she reveals to the right of the recommendation, and this randomization makes the analysis demanding.

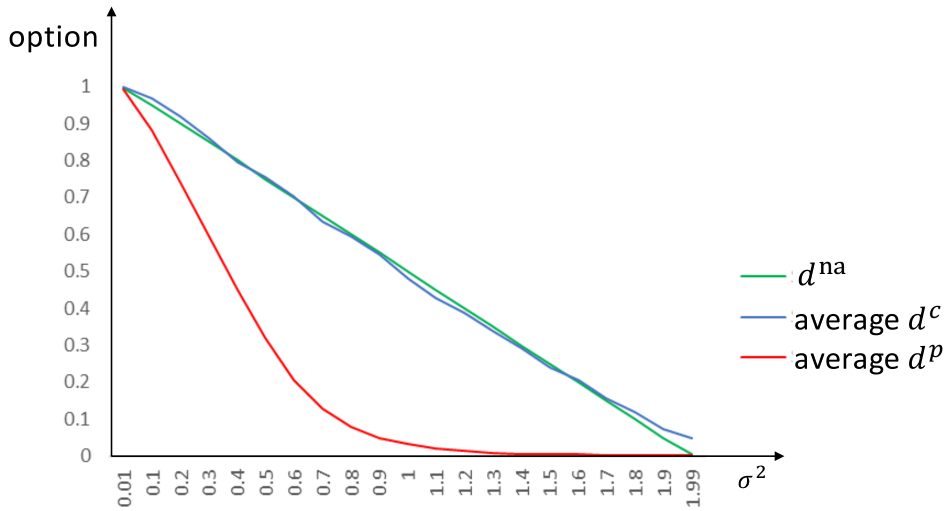

Figure 6: The receiver's decision as a function of σ^2 for no-advice (green), the sender-optimal conative equilibrium (blue) and the interval equilibrium (red), for parameter values $\mu = b = 1$, $n = 1000$, and sample size $10,000$ for every value of σ^2.

recommendation, the receiver is accepting an outcome potentially far from his ideal, and in many cases he would obtain a better outcome by experimenting on his own. Using the same parameter values as in Figure 6, and fixing $\sigma^2 = 1$, simulations show that this, in fact, occurs a majority of the time. The receiver is made worse off 60% of the time from accepting the sender's advice than she would be from going it alone. She nevertheless accepts the advice because while the upside of going it alone is bounded, the downside is not. The interval equilibrium, no matter how favorable to the sender, guarantees the receiver an outcome no further than b from his ideal. Going it alone, in contrast, could leave her with an arbitrarily bad outcome. The value of advice, therefore, is as insurance, and the sender's ability to offer that insurance is what allows her to move the decision in her favor. This explains why we rely on experts yet so often feel ripped-off in doing so.

5.3 How Much Referential Advice is Needed? Credibility and the Interval Strategy

The key pieces of information in the interval equilibrium are the peak and the end point. The peak persuades while the end-point dissuades. This raises the question as to whether the additional information revealed—the rest of the interval—is necessary. This question is particularly pertinent given that Theorem 2 leans on equilibria fully-revealing to the left that presume all of this information is provided.

The additional information beyond the peak and the end-point plays an important role in establishing credibility and, therefore, in making referential advice effective. While all of the information is not necessary to support equilibrium, some of it is required, and, in particular, a better equilibrium for the sender cannot be constructed by conveying less information.

To see this, consider the strategy that follows the interval strategy but only reveals from the recommendation to the right. This differs from the interval equilibrium in omitting information to the left of the recommendation. It can readily be verified that this strategy too supports an equilibrium and that it is equivalent to the interval strategy. This equivalence is akin to that between conative and the conative extension presented above. With the omission of some information, the amended interval strategy is deceptive, and thus the sender can now deviate on path. Critically, however, these deviations only move the decision to the right and are unprofitable. By construction, there is no suitable peak and end pair of points to the left as otherwise the interval strategy would have played it.

The referential advice to the right of the recommendation is not so easily dismissed. Were the sender to reveal only the peak and the end-point, even more on-path deviations would be opened up and, critically, some of these would move the recommendation to the left and be profitable. This undermines the equilibrium. For instance, by revealing a lower outcome to the right of where the interval strategy would stop, the sender can potentially support a recommendation with an even less appealing outcome for the receiver to the left of that in the interval equilibrium.

The information between the peak and the end-point is important because it provides credibility to the sender. The receiver needn't ask *why* did the sender reveal the peak that she did? By seeing all of the points between the peak and the end, the receiver can be sure there is no better peak for him, and thus accept the recommendation with confidence. Without this referential advice, the receiver cannot be sure he is seeing what he is supposed to see.

The common thread across these equilibria is that they implement a relative rather than an absolute standard (as in conative equilibria). Convincing the receiver that the recommendation meets a relative standard requires credibility and this, in turn, requires the sender to provide additional, referential advice. In these equilibria advice fulfills three roles. It persuades, it dissuades, and it provides credibility. In our model persuasion and dissuasion each lean on only a single piece of information. The interval equilibrium shows that the third requirement—that of credibility—can require more information, and possibly much more.

5.4 Other Referential Equilibria

Referential Advice with Less Information. Referential equilibria exist that are neither fully-revealing to the left nor equivalent to an equilibrium that is. Advice in some of these equilibria can be very limited. Indeed, in some states, these equilibria can require the sender to reveal only a single piece of referential advice. Here is an example.

Example 3 *Given some $T^{**} < 0$, the sender reveals the smallest option d such that $X(d) \leq T^{**}$, and if no such option exists, reveals all the options.*

If the lower threshold is reached, revealing that point dissuades the receiver from experimenting on his own and thereby persuades him to accept the default option as a recommendation. The referential advice, when the threshold is met, is a single piece of information. If the threshold is not, the sender fully reveals her information.

This equilibrium improves on the interval equilibrium in some states. In the limit she is able to persuade the receiver to choose the default option with positive probability, whereas in the interval equilibrium this is a zero probability event. The cost of doing so, however, is that she fully reveals if the threshold is not met, and in these cases she does strictly worse than not only the interval equilibrium but also the sender-optimal conative equilibrium.

Full revelation is a rather excessive response to the threshold not being met, and this raises the question of whether the sender can do better in these states. To see what can be accomplished in the limit, and to see the tension in such equilibria, consider the following example.

Example 4 *Suppose that n is a perfect square such that option $d_{q\sqrt{n}} = q$, a fixed distance from d_0 for all such n.*[26] *For some $T^{*} < 0$, the sender's strategy is as follows:*

[26] More generally, we can consider the case of any n if we replace $d_{q\sqrt{n}}$ by $d_{\lfloor q\sqrt{n}\rfloor}$, where $\lfloor q\sqrt{n}\rfloor$ refers to the integer part of $q\sqrt{n}$.

(i) she reveals option $d_{q\sqrt{n}}$ if $X(d_{q\sqrt{n}}) \leq T^*$, otherwise

(ii) she fully reveals all options if $X(d) \geq b - \Delta^{\max}$ for some $d \in [0, d_{q\sqrt{n}}]$, or

(iii) she follows the sender-optimal first-point conative strategy with $\Delta^{\max}$.

Referential advice in this equilibrium can still involve a single piece of information, and the sender wields this minimal advice to persuade the receiver to choose the default option. Moreover, rather than always completely revealing her information when this threshold is not met, the sender transitions to the sender-optimal conative strategy in part (iii).

The conative strategy cannot be used all the time in this situation, however. Attempting to persuade with the default option carries a cost when it doesn't come off, and this is evident in part (ii). When the sender fails to reveal that the outcome of $d_{q\sqrt{n}}$ is sufficiently bad, information spills over indirectly into the smaller options. Knowing that $X(d_{q\sqrt{n}})$ is above T^*, the receiver's beliefs for all smaller options are also higher. Thus, he is unwilling to accept an outcome $\Delta^{\max}$ from his ideal as without the risk of a very low outcome, he would be emboldened to experiment further should that be the recommendation.

The strategies in Examples 3 and 4 can be combined and sliced in many different ways. The resulting equilibria can outperform the interval equilibrium in some states by inducing a smaller action up to and including the default option. Nevertheless, no matter how this is done, Theorem 2 tells us that it has to involve a cost. In other states the equilibrium must lead to larger actions and, in fact, that in some of these states the option chosen is worse than even the sender-optimal conative equilibrium.

Non-prescriptive Equilibria. A non-prescriptive equilibrium requires that the receiver choose an option that is not revealed. As such, such an equilibrium is necessarily referential. In our environment it is difficult to construct equilibria that are non-prescriptive. We have been unable to do so, although we have also been unable to rule them out.

We saw earlier that no-advice cannot support equilibrium behavior as, with neutral beliefs, full revelation is profitable for some types. This argument extends to all non-deceptive strategies. Full revelation undermines any non-deceptive strategy that is non-prescriptive. The only exception is if the option the receiver chooses is immediately adjacent to a revealed point as then there is no space for the mapping to cross b between a revealed option and what the receiver will choose. We refer to such equilibria as *near-prescriptive*.

Definition 7 *An equilibrium is* near-prescriptive *when, for every on-path message m, either the receiver chooses an option disclosed in m, or the receiver chooses an option just above or just below an option disclosed in m.*

Such an equilibrium does not capture the spirit of non-prescriptiveness, however, for as n grows large, the option space grows dense and these equilibria effectively become prescriptive. Nevertheless, for non-deceptive strategies this is the only type of non-prescriptive equilibrium that is possible.

Lemma 3 *If an equilibrium is non-deceptive, then it is near-prescriptive.*

For a truly non-prescriptive equilibrium to exist, therefore, the strategy must be deceptive. To rule out the full-revelation deviation, it must also be that any message with a gap cannot profitably be sent by a type in which the path crosses b in that gap and nowhere else to the right. Intuition suggests that this makes supporting an equilibrium difficult, although demonstrating this formally with non-neutral beliefs (and without closed form representations) is challenging.[27]

Receiver-Optimal Equilibria. The amount of referential advice required in equilibrium is put into sharp relief for receiver-optimal equilibria. Full revelation is receiver optimal, of course, but it is not uniquely so. One may wonder how little information can the sender provide in such an equilibrium and whether that advice can be conative.

With a finite set of options, advice is necessarily referential in any receiver-optimal equilibrium. In fact, we show in Appendix B (Proposition B.1) that in any receiver-optimal equilibrium, the sender must reveal the best option for the receiver and *all* options to the right of it. She can reveal more, up to full revelation, but she can reveal no less than this.

The reason referential advice is required is because, with a finite set of options, the *best* option for the receiver is a relative standard. The receiver does not know what outcome is her best among those available, and so he

[27]For instance, consider the following strategy that is non-prescriptive for some states and prescriptive for others: Denote by d^* the smallest option such that $X(d^*) \geq b$; then for some d^{np} large, (i) If $d^* \leq d^{\mathrm{np}}$, reveal all options up to and including d^*, (ii) otherwise reveal nothing. In case (i) the receiver chooses option d^*, and in case (ii) he chooses d^{np}. Case (ii) is non-prescriptive, and the receiver knows that the outcome is strictly below b. Nevertheless, if d^{np} is sufficiently large, the density may be packed sufficiently tightly to b that he prefers d^{np} than to experiment further to the right. Clearly, however, even were this strategy able to support an equilibrium, it does not dominate the sender-optimal conative equilibrium, let alone the interval equilibrium.

needs to be convinced that the recommendation made is indeed the best. He is only convinced when he knows that the sender has no profitable deviations available, and for that to hold, the receiver must reveal all information to the right. The strategy is deceptive when the sender reveals only these points although, as with other examples above, this strategy has a non-deceptive equivalent equilibrium in full revelation.

This example provides insight into the distinction between a relative and an absolute standard in advice. While the receiver's optimal option is relative in a finite sample, it is absolute in a continuum. With probability one an option exists that produces outcome b, and if the receiver is advised to take that option by the sender, he requires no additional advice to convince him that it is the best option for him. In the limit, therefore, the relative standard becomes absolute, and advice in a receiver-optimal equilibrium can be conative.

In Appendix B (Proposition B.2) we show that a sequence of conative equilibria implement the receiver-optimal outcome in the limit and that equilibria in this sequence are better for the sender than any *exactly* receiver-optimal equilibria. By letting the size of the band, Δ, in a first-point conative equilibrium converge to zero, the sender must reveal a better and better option to the receiver, such that in the limit the option revealed produces outcome b. We show that this can be done slowly enough as n increases, such that the probability an outcome falls within this band goes to one, and the interaction is conative almost surely in the limit.

6 Concluding Discussion

Expertise is everywhere and its importance verges on being self-evident. Yet grasping why and how it matters, and teasing its effects out empirically, has proven more elusive than one may have expected. The objective of this paper has been to shed more light on the role of expertise, how it manifests in advice, and when it matters. At an abstract level, we hope that the use of referential information resonates with intuition and experience, and opens up new questions. At a practical level, we aim for our results to open up new channels of understanding and new interpretations of old data. Before concluding, we offer briefly here several areas where this opportunity is most promising.

We have analyzed an environment with unrestricted communication. In practice, a design choice for decisionmakers is how much information they allow experts to provide to them. CEOs famously prefer "executive summaries" over lengthy reports, Congress requires that only the text of a bill be reported to the floor by a committee without supporting reports, and common law dictates that only judicial decisions themselves form binding precedent and not any supporting arguments contained in a judge's written opinion. The wisdom of these restrictions depends on what equilibrium is expected to be played were communication unrestricted. If it is the interval equilibrium, then the CEO is better off with only a recommendation, whereas if the decisionmaker can expect (or mandate) full disclosure of private information, restricting the channel of communication helps the expert and not the decisionmaker.[28]

Beliefs play an important role in our analysis. They are rich, broad, and, at times, subtle. This suggests that equilibrium refinements may shed additional insight into which equilibria are the more reasonable. Standard refinements can be applied here, although their ability to refine prediction is limited; for instance, the never-weak-best-response (NWBR) refinement eliminates none of the equilibria of interest. A more promising, and novel, route that our framework opens up is to distinguish between deviations that over communicate versus those that under communicate. A sender over-communicates when she sends an equilibrium message *plus* additional information. Under-communication is then deviations that are not a super-set of any equilibrium message. It seems plausible that over-communication conveys sincerity and openness, whereas under-communication indicates deception and breeds mistrust. If so, over-communicated information is then taken at face value and beliefs are neutral.[29] This refinement eliminates all non-prescriptive equilibria that are Pareto inefficient, yet the interval equilibrium and all first-point conative equilibria survive.[30] Exploring these ideas in theory and particularly in the lab promise to yield more insight into the psychology and strategy of communication.[31]

[28]This suggests that the impact of expertise may be identified through what and how much an expert communicates and the range of outcomes that it produces. Conative equilibria exist only in a relatively narrow band, whereas referential equilibria range from the worst possible for the expert up to the interval equilibrium and perhaps beyond. Evidence of expertise should also exist beyond the choice itself and in the beliefs a decisionmaker holds over all options.

[29]One may interpret over-communication through the lens of "speeches" that motivated early refinements. The sender is showing the equilibrium outcome and then, in effect, saying "here is additional information that will make us both better off."

[30]The Pareto requirement eliminates the possible equilibrium described in Footnote 27. It also implies that the sender can over-communicate all information to the left of the recommendation without upsetting the equilibrium, meaning that the surviving non-prescriptive equilibria have near-prescriptive analogues.

[31]There is scope also for refinements based on traditional arguments. For example, the classic idea that beliefs are inherently uncertain, such that a receiver would accept a recommendation that is close enough to his ideal even if he is uncertain it is the absolute best, would eliminate full revelation as an equilibrium, or any exactly receiver optimal equilibria.

Although the focus has been on equilibria when the number of options is large, the results are characterized for finite n. Indeed, for referential advice to help the expert in equilibrium there can be as few as three options (as with only two options there is no space for a recommendation and referential advice that dissuades). Our results show simply that the bigger the set is, the more likely the opportunity will arise for referential advice to be able to influence the decisionmaker, and for referential to dominate conative equilibria.[32] A similar conclusion emerges if the outcome space is bounded. The insurance value of advice is a relative effect, and even in bounded domains the decisionmaker would be tempted to accept a "good enough" outcome to avoid potentially worse outcomes.

One stark assumption in the model is that the expert knows the state of the world precisely. This aids tractability but is not necessary. Our results are unchanged if we add an independent noise term to the expert's knowledge of each outcome. A variant more in the spirit of our model is to instead suppose that the precision of the expert's knowledge also varies in the option itself, increasing in the distance an option is from the known default option. In Appendix A we formalize this idea via a noise term that itself is modeled as a Brownian motion and anchored at the default option. The striking implication to come from this is that the conative equilibria no longer exist whereas the referential equilibria do, providing more support for our focus on referential advice. Conative equilibria break down as a recommendation that is far from the default implies the expert is very unsure of the recommendation herself, and for a sufficiently distant recommendation the receiver will prefer an option near or at the default. The interval equilibrium, in revealing all outcomes from the default to the right, does not succumb to this problem.

The model we develop in this paper exposes clearly how strategic communication is deeply intertwined with experimentation. The receiver's decision to accept advice is made in comparison to what he can achieve by experimenting on his own, and the benefit of experimentation is itself affected by the advice that is conveyed. While our model differs from the communication literature in relaxing the perfect correlation of states, it is notable that it simultaneously differs from the experimentation literature in relaxing the independence of states. By parameterizing the correlation, our model demonstrates how the literatures can be connected. We explore a static model, as is standard in the communication literature, although there is no logical barrier to it being

[32] Referential advice would likely be more effective in settings in which the expert were able to convince the decisionmaker that a good option doesn't exist. Our setting makes the task more difficult by ruling out that possible channel of influence.

dynamic as is the experimentation literature. We hope that the connection between these areas of economic decisionmaking can be explored more deeply in further research.

Appendices

The appendices are organized as follows.

Appendix A offers four extensions of the main model. Section A.1 considers receiver utility beyond the quadratic case. Section A.2 explores the case in which the sender has preference for larger options. Section A.3 discusses the case in which the sender is imperfectly informed about the outcomes. Section A.4 extends the message space of the main model by allowing the sender to send, as message, a set of states that includes the true state.

Appendix B includes results omitted from the main text. Section B.1 provides sufficient conditions on the existence of off-path suspicious beliefs needed to sustain an equilibrium. Section B.2 shows that, on the equilibrium path of a prescriptive equilibrium, the sender is always informative. Section B.3 shows that, in every non-deceptive prescriptive equilibrium, interactions are fully-revealing to the left. Section B.4 formalizes the facts mentioned in Section 5.4 of the main text regarding receiver-optimal equilibria.

Finally, Appendix C contains the proofs of the results stated in the main text.

A Extensions

A.1 General Receiver Utility

To simplify the presentation, in the main text we have developed the model for quadratic receiver utility. Our results do not depend on this restriction, with the sole exception of Corollary 1. In this section we develop general conditions on receiver utility under which our results continue to hold. The intuition in this general setting remains unchanged, even though the details can be technical and complex. All the proofs for our results (in Appendix C) are provided under general receiver utility and, to that end, we establish in this section several intermediary results relevant to this environment.

Without quadratic utility, how the receiver trades off risk and return depends on the degree of underlying uncertainty and the expected outcome itself. This means, for example, that absent expert advice, the receiver's favored option will not always deliver outcome $b - \frac{\sigma^2}{2\mu}$, rather it will depend on the default outcome. Despite this, the principles that underlie all of our equilibria, and that lead to dominance of the interval equilibrium, carry through. Relaxing quadratic utility, our results continue to hold if the following assumptions on receiver utility hold. We define

$$R(x) = -\frac{u''_{\mathrm{R}}(x)}{u'_{\mathrm{R}}(x)} \quad \text{and} \quad P(x) = -\frac{u'''_{\mathrm{R}}(x)}{u''_{\mathrm{R}}(x)}$$

as the coefficients of absolute risk aversion and absolute prudence, respectively.[33]

ASSUMPTION 1. u_{R} is smooth to the fourth order and exponentially dominated.[34]

ASSUMPTION 2. $u''_{\mathrm{R}} < 0$ and $u_{\mathrm{R}}(x)$ is maximized at $x = b > 0$, the ideal outcome of the receiver.

ASSUMPTION 3. On the range $(-\infty, b)$, $R' > 0$.

ASSUMPTION 4. If $P \neq 0$, $P' < 0$.

ASSUMPTION 5. $R(0) < 2\mu/\sigma^2$.

Assumption (1) is a technical condition to ensure expectations are well defined. Assumption (2) ensures that u_{R} is strictly concave and an ideal outcome for the receiver exists. Assumptions (3) and (4) ensure that the optimization problem the receiver faces is concave (see Lemma A.1 below), using the concepts of absolute risk aversion and absolute prudence from the literature on decision making under uncertainty (see Callander and Matouschek, 2019, for a discussion of these conditions). Finally, Assumption (5) is the analog of the requirement $b > \frac{\sigma^2}{2\mu}$ that ensures the sender's problem is non-trivial.

If receiver utility meets Assumptions (1)–(5), the logic for conative equilibria is the same, although the condition describing the first-point conative equilibria is more involved. Let $\underline{x}$ be the unique outcome $x < b$ at which the coefficient of absolute risk aversion is equal to $2\mu/\sigma^2$. Similarly, let $\bar{x}$ be the unique outcome $x > b$ such that $u_{\mathrm{R}}(x) = u_{\mathrm{R}}(\underline{x})$. Therefore, the receiver utility of any outcome outside the range $[\underline{x}, \bar{x}]$ is less than the receiver utility of all outcomes inside that range. By Assumption (5), we have $\underline{x} \in (0, b)$. Lemma A.2 below then implies that the sender cannot get her first best in equilibrium if n is large enough (following the argument used in the main text to show the necessity of equilibrium advice). Observe that, for the case of quadratic receiver utility in

[33] Assumptions (1) and (2) imply that the coefficient of absolute risk aversion is well defined and positive on the range $(-\infty, b)$, and that the coefficient of absolute prudence is well defined everywhere.

[34] Formally, u_{R} is four times continuously differentiable, and for all $\alpha > 0$, $u_{\mathrm{R}}^{(k)}(x)e^{-\alpha|x|} \to 0$ as $|x| \to \infty$, where $u_{\mathrm{R}}^{(k)}$ denotes the derivative of order k.

the main text, $\underline{x} = b - \frac{\sigma^2}{2\mu}$ and $\bar{x} = b + \frac{\sigma^2}{2\mu}$. The existence and uniqueness of $\underline{x}$ follow from Assumptions (3) and (5), together with the fact that the coefficient of absolute aversion $-u''_{\mathrm{R}}(x)/u'_{\mathrm{R}}(x)$ becomes unbounded as x approaches b because, by Assumption (2), u'' is bounded above negatively while u' vanishes. The existence and uniqueness of $\bar{x}$ follow from Assumption (1).

We now extend the definition of first-point conative sender strategy to the case of general receiver utility.

Definition 8 *In the case of general receiver utility, we say that the sender follows a* first-point conative strategy *when there exists* $\underline{\Delta}, \overline{\Delta} > 0$*, with* $u_{\mathrm{R}}(b - \underline{\Delta}) = u_{\mathrm{R}}(b + \overline{\Delta})$*, and such that the sender reveals the smallest option whose outcome falls in the range* $[b - \underline{\Delta}, b + \overline{\Delta}]$ *and if no such option exists, the sender reveals everything.*

The sender-optimal conative strategy is then the first-point conative strategy with the largest range $[b - \underline{\Delta}, b + \overline{\Delta}]$ that supports a prescriptive equilibrium, analogously to the case of quadratic receiver utility, and as n grows to infinity this largest range becomes $[\underline{x}, \bar{x}]$. Regardless of the default outcome there exists an outcome that the receiver is willing to accept risk-free rather than to engage in more experimentation.

Similarly, the interval equilibrium carries the same logic, although the threshold that stops revelation is more complicated, depending on how the receiver trades off risk from his optimal experiment should he ignore the advice. This means that the exact equilibrium behavior will vary from that with quadratic utility, and potentially vary a lot for specific mappings.

Despite this variation, the comparison of the interval equilibrium to the sender-optimal conative equilibria carries through exactly as before. The thresholds for each equilibrium may vary, but as the logic of the interval equilibrium depends on the level of the sender-optimal conative equilibrium, it remains the case that the necessary condition will be satisfied before the conative threshold is reached.

The one result that does not extend to general receiver preferences is Corollary 1 that reveals a precise equivalence in the expected outcome of the sender-preferred conative equilibrium and when the expert is absent (and provides no advice). The logic of that result is general although the exact equality is special. It depends on the independence of the receiver's optimal experiment from the default outcome that only holds with quadratic utility. The difference is, however, small in the sense that the leverage of the expert is small in the receiver-optimal conative equilibrium with the relative comparison depending on the curvature of utility and not any fundamental property of strategic communication.

We end this section with two technical lemmas needed in the proofs of our main results to account for general receiver utility. In each lemma, Z is an independent random variable that follows the standard normal distribution.

Lemma A.1 *For all $x_0 \in \mathbf{R}$, the mapping $\Delta \mapsto \mathbb{E}\Big[u_{\mathrm{R}}(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big]$ defined for $\Delta \geq 0$ is strictly concave.*

Proof. Let $x_0 \in \mathbf{R}$ and $f(\Delta) = \mathbb{E}\Big[u_{\mathrm{R}}(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big]$. First, when absolute prudence is strictly decreasing, by continuous differentiation,

$$\frac{(u_{\mathrm{R}}''')^2 - u_{\mathrm{R}}''''u_{\mathrm{R}}''}{(u_{\mathrm{R}}'')^2} < 0,$$

except possibly on a set of measure zero. Since $u_{\mathrm{R}}'' < 0$, $u_{\mathrm{R}}'''' \leq 0$, and $u_{\mathrm{R}}''' < \sqrt{u_{\mathrm{R}}''u_{\mathrm{R}}''''}$, except possibly on a set of measure zero. Then, observing that

$$\begin{aligned}
\mathbb{E}\Big[u_{\mathrm{R}}'(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)Z\Big] &= \sigma\sqrt{\Delta}\ \mathbb{E}\Big[u_{\mathrm{R}}(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big], \text{ and} \\
\mathbb{E}\Big[u_{\mathrm{R}}''(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)Z\Big] &= \sigma\sqrt{\Delta}\ \mathbb{E}\Big[u_{\mathrm{R}}'(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big],
\end{aligned}$$

we get

$$f'(\Delta) = \mu\ \mathbb{E}\Big[u_{\mathrm{R}}'(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big] + \frac{\sigma^2}{2}\ \mathbb{E}\Big[u_{\mathrm{R}}''(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big],$$

and

$$\begin{aligned}
\frac{1}{\mu^2}f''(\Delta) = {} & \mathbb{E}\Big[u_{\mathrm{R}}''(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big] \\
& + 2\left(\frac{\sigma^2}{2\mu}\right)\ \mathbb{E}\Big[u_{\mathrm{R}}'''(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big] \\
& + \left(\frac{\sigma^2}{2\mu}\right)^2\ \mathbb{E}\Big[u_{\mathrm{R}}''''(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big].
\end{aligned}$$

Finally, using that $u_{\mathrm{R}}'''' \leq 0$ and $u_{\mathrm{R}}''' < \sqrt{u_{\mathrm{R}}''u_{\mathrm{R}}''''}$ (except possibly on a set of measure zero),

$$\frac{1}{\mu^2} f''(\Delta) \\ < - \mathbb{E}\left[\left(u_{\mathrm{R}}''(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z) - \left(\frac{\sigma^2}{2\mu}\right) u_{\mathrm{R}}''''(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\right)^2\right],$$

so that $f''(\Delta) < 0$ and f is strictly concave. ■

Lemma A.2 *Let f be the mapping $\Delta \mapsto \mathbb{E}\Big[u_{\mathrm{R}}(x_0 + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big]$ defined for $\Delta \geq 0$.*

- *If $x_0 \geq \underline{x}$, $f(\Delta)$ is maximized at $\Delta = 0$.*
- *If $x_0 < \underline{x}$, there exists $\overline{\Delta} > 0$ such that for all $\Delta \in (0, \overline{\Delta})$, $f(\Delta) > f(0)$.*

Proof. As in the proof of Lemma A.1, we have

$$f'(0) = \mu u_{\mathrm{R}}'(x_0) + \frac{\sigma^2}{2} u_{\mathrm{R}}''(x_0),$$

and f is strictly concave by Lemma A.1. By Assumptions (3) and (5), $f'(0) \leq 0$ if $x_0 \geq \underline{x}$, so $f(0)$ is a maximal value by concavity of f. Similarly, $f'(0) > 0$ if $x_0 < \underline{x}$, so f is strictly increasing in a neighborhood of 0. ■

A.2 Sender Prefers Larger Options

The assumption that the sender prefers smaller options is not a normalization. It implies that the option that is best for the sender is the one that, in the absence of any advice, is safest for the receiver. We now explore the opposite case in which the sender wants to convince the receiver to choose the option about which, in the absence of advice, he is most uncertain about.

So suppose that the sender's preferences are strictly increasing in d, using $u_{\mathrm{S}}(d) = +d$ as the sender's utility function. The receiver's utility function remains quadratic as in Section 2. Since the sender plays no role in the no-advice benchmark of Section 4.3, this change does not affect the option the receiver would choose in the absence of any advice, which is still given by d^{na}, the option that produces outcome $b - \frac{\sigma^2}{2\mu}$ in expectation.

Now suppose the receiver knows that $X(d^*) = x^*$, and let us briefly revisit the case of optimal receiver decision assuming neutral beliefs. Note that, the receiver's preferences being unchanged, the case $x^* < b$ is already treated as part of Section 4 and Theorem 1. The case $x^* > b$—relevant when the sender prefers larger decisions—is more subtle. If $x^* > b$, then, letting

$$\beta(d) = \min\left\{\frac{1}{2}(\sqrt{b^2 + 2d\sigma^2} - b), b\right\},$$

we have the following:

- If $d^* < \frac{4b^2}{\sigma^2}$, so that $\beta(d) = \frac{1}{2}(\sqrt{b^2 + 2d\sigma^2} - b)$, then for $x^* < b + \beta(d^*)$, the receiver's optimal decision is d^*. As x^* increases, the optimal decision decreases. In the limit as x^* grows to infinity, the optimal decision converges to 0.

- If $d^* > \frac{4b^2}{\sigma^2}$, so that $\beta(d) = b$, then for $x^* < b + \beta(d^*) = 2b$, the receiver's optimal decision is d^*. For $x^* \in [2b, \frac{\sigma^2}{2b}d^*]$, the optimal decision is 0. Then, as x^* increases above $\frac{\sigma^2}{2b}d^*$, the optimal decision gradually increases towards d^*, and then decreases again towards 0. In the limit as $x^* \to \infty$, the optimal decision converges to zero.

The first result concerns the existence of an equilibrium with conative interactions.

Proposition A.1 *If $x_C : \mathcal{D} \mapsto \mathbf{R}$ is nondecreasing with $x_C(d) \in [b, b + \beta(d)]$ for all options d, then there exists a sequence of equilibria that are conative in the limit where $x_C(d^*)$ is the outcome of equilibrium decision d^*.*

Proof. Let $\Delta(d) = x_C(d) - b$, fix n, and consider the following sender strategy: The sender reveals the largest option d whose outcome belongs to $[b - \Delta(d), b + \Delta(d)]$. If no such option exists, the sender reveals all options. Then, suppose that, upon receiving an on-path message, the receiver operates as follows: If the entire state is revealed, the receiver simply chooses the utility maximizing option. If two or more options are optimal, as in the main text, the receiver chooses the sender-preferred option—here, the largest one. Otherwise, the receiver chooses the only revealed option.

The sender's strategy is deceptive but optimal when restricted to on-path messages, by the same argument as in the proof of Theorem 1.

Let us show that the receiver's strategy is also optimal among on-path messages. The case of the sender revealing the full state is immediate. Let us focus on the case in which the sender reveals only one option d^* whose outcome x^* belongs to $[b - \Delta(d), b + \Delta(d)]$. The receiver is never better off choosing an option strictly larger than d^*, because x_C is nondecreasing and so any such option yields an outcome even further away from b than is x^*.

If $d < d^*$, the receiver believes that $X(d)$ is distributed normally with mean $\mathbb{E}[X(d) \mid X(d^*)=x^*] = dx^*/d^*$ and variance $\text{Var}[X(d) \mid X(d^*)=x^*] = \sigma^2(d^* - d)d/d^*$. We consider two cases.

- If $x^* \leq b$ then choosing option $d < d^*$ yields an outcome which, on average, is further away from b than is x^*, and in addition increases the variance of the outcome. The receiver is better off choosing d^*.

- If $x^* > b$ then choosing option $d < d^*$ yields the receiver's expected utility

$$
\begin{aligned}
&- (b - \mathbb{E}[X(d) \mid X(d^*)=x^*])^2 - \text{Var}[X(d) \mid X(d^*)=x^*] \\
&\qquad\qquad\qquad\qquad = -\left(b - \frac{dx^*}{d^*}\right)^2 - \sigma^2 \frac{(d^* - d)d}{d^*}
\end{aligned}
$$

  whereas the receiver's expected utility when taking option d^* is $-(b-x^*)^2$. Hence, the receiver is better off choosing d^* when the difference of the two terms, which simplifies to

$$
\frac{d^* - d}{(d^*)^2}\big((d + d^*)x^* - dd^*\sigma^2 - 2bd^*x^*\big), \tag{A.1}
$$

  is non-positive. We have $d^* - d \geq 0$, we remark that the term $(d+d^*)x^* - dd^*\sigma^2 - 2bd^*x^*$ is linear in d, so if it evaluates non-positively at the two extremes $d = 0$ and $d = d^*$, it also evaluates non-positively at all options between the two extremes. If $d = 0$, then

$$
(d + d^*)x^* - dd^*\sigma^2 - 2bd^*x^* = (x^* - 2b)x^*d^* \leq 0
$$

  because $x_C(d^*) \leq b + \beta(d^*) \leq 2b$. If $d = d^*$ then

$$
(d + d^*)x^* - dd^*\sigma^2 - 2bd^*x^* = d^*(2(x^* - b)x^* - d^*\sigma^2)
$$

  which is quadratic in x^* and is non-positive if and only if $x^* \in [\frac{1}{2}(\sqrt{b^2 + 2d\sigma^2} - b), \frac{1}{2}(\sqrt{b^2 + 2d\sigma^2} + b)]$, and, so, is non-positive because for all d, $x_C(d)$ is in the interval $[b, b + \beta(d)] \subset [\frac{1}{2}(\sqrt{b^2 + 2d\sigma^2} - b), \frac{1}{2}(\sqrt{b^2 + 2d\sigma^2} + b)]$. Hence Equation (A.1) is non-positive and the receiver's maximal expected utility is reached when choosing d^*.

Thus, the receiver's strategy is optimal among on-path messages.

Finally, observe that if the outcome of any option d falls in $[b-\Delta(d), b+\Delta(d)]$, the sender is never strictly better off deviating by revealing the full state: Doing so would yield a decision that is at least as large as the decision expected in equilibrium.

Lemma B.1 (or rather, its analog, which continues to hold for the sort of sender's preferences considered here) applies: The strategies defined above can be completed with appropriate off-path beliefs and decisions to form an equilibrium.

Sending n to infinity, the outcome path becomes distributed as a Brownian motion with drift $\mu > 0$. This Brownian motion (almost surely) hits the frontier defined by the function $x_C(\cdot)$, which implies, by the same argument as in the proof of Theorem 1, that as $n \to \infty$ the equilibria just defined become conative. In addition, the sender then reveals the pair $(d^*, x_C(d^*))$ where d^* is the largest option for which $X(d^*) = x_C(d^*)$. ■

In addition, the following result implies that the sender's preferred conative equilibrium, in the limit, implements outcome $b + \beta(d^*)$ for equilibrium decision d^*. Its proof is analogous to the proof of the first part of Theorem 1 and is omitted.

Lemma A.3 *In any sequence of equilibria that are conative in the limit, the probability that the equilibrium outcome of equilibrium decision d^* is above $b + \beta(d^*)$ vanishes as n grows large.*

The logic of conative advice remains unchanged. Previously the sender revealed the smallest option with outcome sufficiently close to the receiver's ideal. With increasing preferences, she now reveals the largest option that is sufficiently close. This outcome is then above b rather than below. A difference is that the threshold that defines "close enough" is now non-constant. The reason is that the relevant domain over which the receiver may experiment rather than accept advice is now a Brownian bridge. The utility of experimentation depends on the slope of that bridge, and that in turn depends on how large is the option that the sender recommends. The larger is the option, the flatter is the bridge and the less tempted to experiment is the receiver, which allows the sender to pull the outcome further above b. At the other extreme, the smaller is the option, the steeper is the bridge, such that as the recommended option approaches d_0, the outcome it produces must itself approach b. For quadratic receiver utility, this threshold is given by $\frac{1}{2}\big(b + \sqrt{b^2 + 2\sigma^2 d}\big)$, as marked in Figure 7a. The construction of this strategy now yields neutral beliefs to the left of the revealed option rather than to the right.

It is evident in this case that, in expectation, conative equilibria do shift the decision in the sender's favor relative to what the receiver would choose in the absence of advice. Without advice, the receiver's risk aversion creates

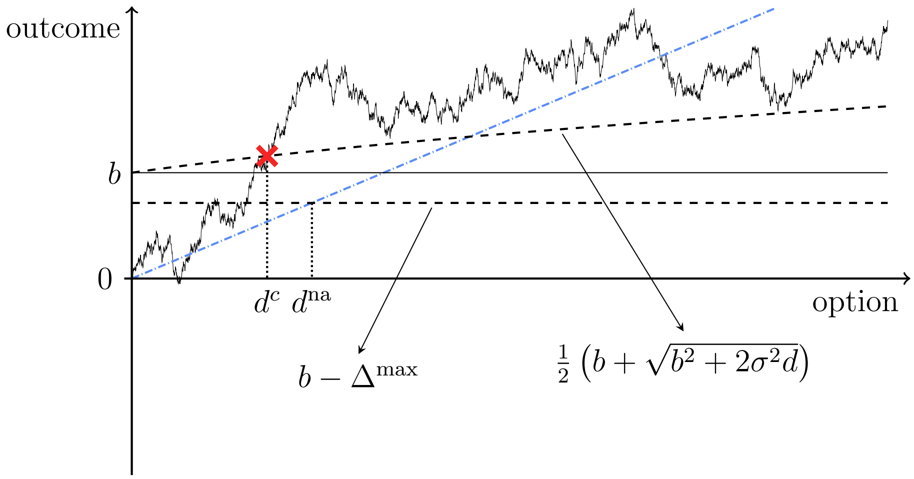

(a)

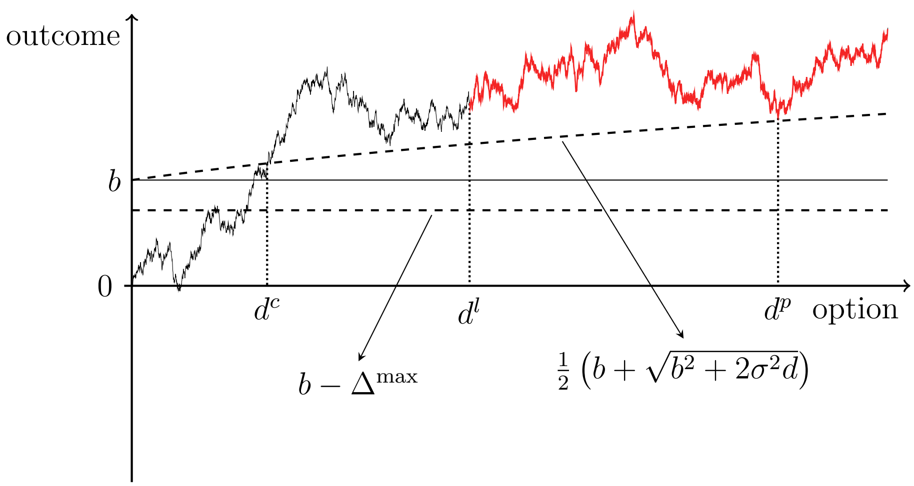

(b)

Figure 7: Strictly increasing sender preferences.

a timidity that works in the opposite direction to the sender's preferences, whereas it previously worked in her favor.

It still remains the case, however, that the sender can sway the decision more in her favor by providing referential advice than she can by making a recommendation alone, as shown by our second result below.

Proposition A.2 *There exists a sequence of referential equilibra $\Sigma_1, \Sigma_2, \ldots$ such that, for any sequence of equilibria $\Sigma'_1, \Sigma'_2, \ldots$ that are conative in the limit, Σ_n provides an expected utility for the sender that is greater than Σ'_n in the limit as $n \to \infty$.*

Proof. Fix n. We build an equilibrium Σ_n whose structure is similar to the interval equilibrium of the main model. Let d^l be the largest option that satisfies either one of these two properties: Either $b \leq X(d^l) \leq b + \beta(d^l)$, or, $X(d^l) > b + \beta(d^l)$ and

$$\max\{-(X(d) - b)^2 : d \geq d^l\} \geq -b^2 \quad \text{if} \quad d^l \geq \frac{4b^2}{\sigma^2} \quad \text{and} \quad \frac{X(d^l)}{d^l} \leq \frac{\sigma^2}{2b},$$

$$\max\{-(X(d) - b)^2 : d \geq d^l\} \geq \frac{d^l\sigma^2(d^l\sigma^2 - 4b(X(d^l) - b))}{4(X(d^l)^2 - d^l\sigma^2)} \quad \text{if} \quad d^l < \frac{4b^2}{\sigma^2} \quad \text{or} \quad \frac{X(d^l)}{d^l} > \frac{\sigma^2}{2b}.$$

Let $d^l = 0$ if no such option exists.

Consider the following sender strategy: The sender reveals d^l and all the options to the right of d^l. Then, upon receiving an on-path message, the receiver chooses the option that maximizes his utility among the set of options that are revealed by the sender, and as before, if two or more options maximize the receiver's utility, the receiver chooses the largest one.

Two facts are worth noting. First, $-b^2$ is the maximum possible expected utility the receiver can get if he chooses an option d to the left of d^l, for the case when $d^l \geq \frac{4b^2}{\sigma^2}$ and when $\frac{X(d^l)}{d^l} \leq \frac{\sigma^2}{2b}$. Second, it can be shown that

$$\frac{d^l\sigma^2(d^l\sigma^2 - 4b(X(d^l) - b))}{4(X(d^l)^2 - d^l\sigma^2)}$$

is an upper bound on the maximum possible utility that the receiver can obtain by choosing an option d less than d^l if either $d^l < \frac{4b^2}{\sigma^2}$ or $\frac{X(d^l)}{d^l} > \frac{\sigma^2}{2b}$ (this upper bound becomes tight as n grows large). These two facts follow from the above description regarding the receiver's optimal option under neutral beliefs, and from solving the relevant maximization problems. The calculations are tedious but straightforward and are omitted.

These two facts combined imply that the sender's and the receiver's strategies just described are optimal when restricted to on-path messages. In addition, the sender is never strictly better off deviating by revealing all options. In consequence, as in the proof of Proposition A.1, the analog of Lemma B.1 for the case of sender preferences considered here guarantees that such a sender strategy supports a prescriptive equilibrium.

In the limit as $n \to \infty$, we get a compact characterization of the equilibrium: d^l becomes the largest option that satisfies either

$$X(d^l) = b + \beta(d^l),$$

or, if $d^l < \frac{4b^2}{\sigma^2}$,

$$\max\{-(X(d) - b)^2 : d \geq d^l\} \geq \frac{d^l\sigma^2(d^l\sigma^2 - 4b(X(d^l) - b))}{4(X(d^l)^2 - d^l\sigma^2)}.$$

In the limit, such an option almost surely exists.

Let $\Sigma'_1, \Sigma'_2, \ldots$ be a sequence of equilibria that are conative in the limit. The above-mentioned thresholds that determine the option d^l imply that the equilibrium decision in the sender's preferred conative equilibrium is never greater than the equilibrium decision of this referential equilibrium, and that with positive probability, it is strictly smaller. ∎

The referential equilibria in Proposition A.2 are analogous to the interval equilibria of Section 5.1. Figure 7b illustrates the interval strategy used in the limiting case $n \approx \infty$. Rather than reveal from the left, the sender now reveals from the right, with all revealed outcomes above b. The recommendation is again the closest revealed outcome to b, although it is now the low point in the mapping, and the referential advice is information to the left of that recommendation with the key dissuasive point being the peak, now marked by d^l. In the same way that previously information to the left of the recommendation could be omitted without upsetting the equilibrium, information to the right of the recommendation can be omitted here, reducing the amount of information needed to be communicated in equilibrium.

Because the domain where the receiver may experiment is now a bridge, the construction of equilibrium is somewhat more delicate with the thresholds dependent on the location of the revealed option. Nevertheless an equilibrium can be constructed. For many paths the interval equilibrium yields a strictly better choice by the receiver than do all conative equilibria, but it is no longer true that this holds with probability one. With positive probability the requisite combination of high and low points can't be found, and the interval equilibrium implements the same choice by the receiver as does the sender-optimal conative equilibrium. Thus, even here, it is still the case that the interval equilibrium weakly dominates all conative equilibria as the option set grows large.

A.3 Sender is Imperfectly Informed

We now consider another direction in which the model can be extended, allowing for the sender to be imperfectly informed about the outcomes.

In our context, it is natural to suppose that, just like the receiver, the sender is better informed about options that are closer to the default option. To capture this notion, we assume that the sender observes the realization of a signal function that is correlated with the outcome function. For each option, she then decides whether to reveal her signal. As in the main model, we assume revelations have to be truthful. The utilities for the sender and the receiver are as in the main model (Section 2). The outcome of option d continues to be denoted $X(d)$ and is distributed as in the main model. However, instead of observing directly $X(d)$ for every option d, the sender now observes $Y(d)$, with $Y(d_0) = X(d_0) = 0$ and

$$Y(d_i) = Y(d_{i-1}) + X(d_i) - X(d_{i-1}) + \frac{\varepsilon}{\sqrt[4]{n}}\xi_i,$$

for $i = 1, \ldots, n$, where ξ_i is independently drawn from the standard normal distribution, and ε captures the amount of noise in the signals the sender gets to observe. In the case $\varepsilon = 0$, the sender is perfectly informed about the state, as in our main model. As ε grows larger, the sender becomes gradually less informed until, in the limit in which $\varepsilon = \infty$, she knows as little as the receiver does. Finally, for any positive ε, the sender is more uncertain about the outcomes of options that are further away from the default option, just like the receiver. As n grows large, X becomes distributed as a Brownian motion with drift μ and scale σ, and Y becomes distributed as a Brownian motion with drift μ and scale $\sqrt{\sigma^2 + \varepsilon^2}$.

For every option d, let $Z(d)$ denote the best estimate of $X(d)$—a minimizer of the mean-squared error—given the sender's information:

$$Z(d) = \mathbb{E}[X(d) \mid Y(d'), \forall d'].$$

Let $\gamma = \varepsilon^2/(\sigma^2 + \varepsilon^2)$. The projection formulas for jointly normal random variables imply

$$Z(d) = (1-\gamma)Y(d) + \gamma\mu d, \qquad \text{and} \qquad \mathrm{Var}[X(d) \mid Y(d'), \forall d'] = \gamma\sigma^2 d.$$

In particular, as n grows to infinity, $Z(\cdot)$ is distributed as a Brownian motion with drift μ and scale $(1-\gamma)\sigma$. So, compared to *original* outcome function, the drift of the *estimated* outcome function is the same, but the scale is reduced by the factor $1-\gamma$, which captures the informativeness of the sender's signals. As $\gamma \to 0$, signals become perfectly informative and estimations become confounded with true outcomes, while as $\gamma \to 1$, signals become perfectly uninformative and estimations become equal to the unconditional expected outcomes.

It is worth noting that, for each option d, the value of signal $Y(d)$ is a sufficient statistic to compute the distribution of $X(d)$ conditional on all of the sender's information:

$$\mathbb{E}[X(d) \mid Y(d)] = \mathbb{E}[X(d) \mid Y(d'), \forall d'],$$
$$\mathrm{Var}[X(d) \mid Y(d)] = \mathrm{Var}[X(d) \mid Y(d'), \forall d'].$$

In addition, if $d \geq d'$,

$$\mathbb{E}[X(d) \mid Y(d')] = \mathbb{E}[X(d) \mid Y(d''), \forall d'' \leq d'] = Z(d') + (d-d')\mu, \qquad \text{(A.2)}$$

and

$$\mathrm{Var}[X(d) \mid Y(d')] = \mathrm{Var}[X(d) \mid Y(d''), \forall d'' \leq d'] = \gamma d'\sigma^2 + (d-d')\sigma^2. \quad \text{(A.3)}$$

The expected utility of the receiver who takes option d given knowledge of $Y(d)$ or, equivalently, $Z(d)$, is

$$-(Z(d)-b)^2 - \gamma\sigma^2 d. \qquad \text{(A.4)}$$

Compared to the receiver's expected utility conditional on $X(d)$, $-(X(d)-b)^2$, notice the presence of the second term $-\gamma\sigma^2 d$. This term captures the disutility the receiver gets for choosing larger options, due to the compounded noise in the sender's signals.

Overall, the case of an imperfectly informed sender can be analyzed in much the same way as our main model by noticing that, rather than revealing outcomes depending on when $X(\cdot)$ hits different thresholds, the sender reveals signals (or equivalently, outcome estimates) depending on when $Z(\cdot)$ hits suitably adjusted thresholds.

The key difference between this extension and our main model is that there are no longer any conative equilibria, specifically, the probability of non-conative interactions remains bounded away from zero as n grows large. To see why, suppose the sender's strategy were to reveal the signal of the smallest option at which the estimate function $Z(\cdot)$ hits threshold $b - \frac{\sigma^2}{2\mu}$ (the threshold associated with the sender-optimal conative equilibrium with perfect information). If the estimate function hits the threshold early enough, it is a best response for the receiver to choose the revealed option. But if the estimate function hits the threshold late—for a large option—the receiver is now better off picking an option to the left of the revealed one, perhaps even the default option. In such cases, the receiver understands that the sender is very unsure about the outcome of the option she is revealing and recommending. Rather than follow such a risky recommendation, he turns it down for something closer to the safe default decision. This fact holds more generally and is captured in Proposition A.3 below.

Proposition A.3 *If the sender is imperfectly informed about the state, i.e., $\varepsilon > 0$, then there is no sequence of equilibria that are conative in the limit.*

Proof. Suppose by contradiction that there exists a sequence of equilibria that are conative in the limit, which we write $\Sigma_1, \Sigma_2, \ldots$.

Let $\underline{z} = b - \frac{\sigma^2}{2\mu}$. If the receiver observes the value $Y(d)$ or $Z(d)$ of option d, and if $Z(d) \geq \underline{z}$, then under neutral beliefs, the receiver's expected utility is no greater when choosing option $d' > d$ than when choosing option d. In contrast, if n is sufficiently large and $Z(d) < \underline{z}$, then the receiver gets more expected utility by choosing some option $d' > d$. These facts follow from the same arguments as in Section 4 and imply, by the same arguments as in the proof of Theorem 1, that as n grows large, the probability that the equilibrium outcome of Σ_n is below $\underline{z}$ vanishes.

Since, as n grows large, the estimated outcome path is a Brownian motion starting at 0 with drift μ and scale $(1-\gamma)\sigma$, for every option d_T, there is a positive probability that the outcomes of all the options to the left of d_T are below $\underline{z}$. And, because the receiver incurs a disutility linear in the option chosen, as shown in Equation (A.4), if d_T is large enough, the receiver would rather decide the default option, whose outcome is known. Hence, the probability of non-conative interactions remain bounded away from zero as n grows large. ■

However, there exist prescriptive equilibria in which the sender either reveals the signal of a single option, or reveals all the signals. These equilibria are the analog of the first-point conative equilibria in the main model, but with a threshold or band that is non-constant. For example, let

$$\Delta(d) = \sqrt{\frac{\sigma^4}{4\mu^2} - \gamma\sigma^2 d}, \quad \text{and} \quad d_M = \frac{\sigma^2}{4\gamma\mu^2},$$

and consider the following sender strategy: The sender reveals the signal of the smallest option $d \leq d_M$ whose value falls within the range $[b - \Delta(d), b + \Delta(d)]$. If no such option exists, then the sender reveals all the signals.

Proposition A.4 *The "one-or-all" sender strategy described above is a prescriptive equilibrium strategy.*

Proof. The optimality of the sender's strategy for on-path messages is immediate, by the same arguments as in the second part of Theorem 1. And similarly, the sender is never strictly better off revealing all the signals. Let us show that, for on-path messages, the receiver is best off choosing an option whose signal is disclosed by the sender.

First, note that for all $d \leq d_M$, the range $[b - \Delta(d), b + \Delta(d)]$ is included in $[b - \frac{\sigma^2}{2\mu}, b + \frac{\sigma^2}{2\mu}]$. If the sender communicates the signal of only one option d^*, applying the receiver's beliefs given by Equations (A.2) and (A.3) and using the same logic as in the second part of Theorem 1, the receiver is never strictly better off choosing an option to the right of d^*.

Second, note that, if the sender discloses the signal of a single option d^* such that the signal value is on the boundary of the range $[b - \Delta(d^*), b + \Delta(d^*)]$, the receiver's expected utility, when the receiver chooses d^*, is independent of d^* and is equal to

$$-\left(b - b \pm \sqrt{\frac{\sigma^4}{4\mu^2} - \gamma\sigma^2 d^*}\right) - \gamma\sigma^2 d^* = -\frac{\sigma^4}{4\mu^2}.$$

If the receiver decides $d < d^*$ for which $Z(d) \notin [b - \Delta(d), b + \Delta(d)]$, then the receiver's expected utility, given $Y(d)$ or equivalently given $Z(d)$, is less than $-\frac{\sigma^4}{4\mu^2}$. Hence, the receiver is never strictly better off deviating to the left when the sender reveals the signal of only one option, because according to the sender's strategy, the sender's estimated outcomes of all the options d to the left of the revealed option fall outside the range $[b - \Delta(d), b + \Delta(d)]$.

Thus, the receiver decides d^*. Equilibrium existence then follows from the construction of suspicious off-path beliefs and off-path decisions as done in Lemma B.1, and this equilibrium is, by definition, prescriptive. ■

In this one-or-all equilibrium, as n grows large, and as γ vanishes, d_M becomes infinite and the probability that the estimated-outcome path remains outside of the band defined by $[b - \Delta(\cdot), b + \Delta(\cdot)]$ vanishes. Then, the sender discloses the signal of only one option, and $\Delta(d) \approx \Delta^{\max}$, the threshold of the sender-optimal conative equilibrium under perfect information. So, for large option sets and a sender close to being perfectly informed, the equilibrium behavior becomes arbitrarily close to that of the sender-optimal conative equilibrium of Section 4.

The fact that the receiver is uncertain about the outcomes of large revealed options does not cause any issues for referential advice. Consider the following sender strategy: The sender reveals the signals of options $\{d_0, \ldots, d^r\}$, where d^r is the smallest option that satisfies

$$\max_{d \leq d^r} \mathbb{E}[u_{\mathrm{R}}(X(d)) \mid Y(d)] \geq \max_{d > d^r} \mathbb{E}[u_{\mathrm{R}}(X(d)) \mid Y(d^r)], \tag{A.5}$$

and $d^r = d_n$ if no such option exists. The following proposition asserts that this sender strategy supports a prescriptive equilibrium. By construction, equilibrium interactions are referential.

Proposition A.5 *The sender strategy just described is a prescriptive equilibrium strategy.*

Proof. Note that the left-hand side of Equation (A.5) is equal to

$$\max_{d \leq d^r} \mathbb{E}[u_{\mathrm{R}}(X(d)) \mid Y(d'), \forall d' \leq d^r],$$

the maximum expected utility the receiver can achieve by choosing one of the options disclosed by the sender, conditionally on the hard information included in the sender's message. Similarly, note that the right-hand side of Equation (A.5) is equal to

$$\max_{d > d^r} \mathbb{E}[u_{\mathrm{R}}(X(d)) \mid Y(d'), \forall d' \leq d^r],$$

the maximum expected utility the receiver can achieve by choosing one of the options not disclosed by the sender, conditionally on the hard information included in the sender's message and assuming neutral beliefs. Thus, by construction, if the receiver holds neutral beliefs he is never strictly better off choosing an option not included in the sender's message—interactions must be prescriptive. Observe that the sender strategy is non-deceptive, so that the receiver forms neutral beliefs upon receiving an on-path message, and so that the sender's strategy is trivially optimal among on-path messages, as for each state there exists only one possible on-path message.

Finally, the sender is obviously never strictly best-off deviating to reveal all the signals. Setting suspicious off-path beliefs and off-path decisions as in Lemma B.1 ensures equilibrium existence. ■

Note that as ε vanishes, the above sender strategy converges to the interval strategy of Section 5.1. Thus, as ε vanishes and n grows large, the equilibrium that results strictly dominates the all-or-one equilibrium of Proposition A.4.

Simple but tedious calculations show that, for every $\varepsilon \geq 0$, as n grows large, the sender chooses, as d^r, the smallest option d whose estimated outcome hits the threshold

$$b - \frac{\sigma^2}{4\mu} - \frac{\mu}{\sigma^2}(Z(\widehat{d}) - b)^2,$$

where $\widehat{d}$ is the smallest option less than or equal to d that maximizes the receiver's expected utility conditionally on the sender's information, thus generalizing the threshold obtained for perfectly informed senders (see Equation (9) in Section 5.1).

A.4 General Messages

In this section, we allow the sender to send, as message, any set of states that includes the true state. Let $\overline{\mathcal{M}}$ be this extended set of possible messages. To distinguish between $\mathcal{M}$ as defined in our main model and $\overline{\mathcal{M}}$, we refer to $\mathcal{M}$ as the *regular message space*, and to $\overline{\mathcal{M}}$ as the *extended message space*. The goal is to show that equilibrium outcomes obtained under the regular message space are robust to generalization.

With the extended message space, a strategy for the sender is a mapping M from Θ to $\overline{\mathcal{M}}$, a strategy for the receiver is a mapping D from $\overline{\mathcal{M}}$ to $\mathcal{D}$, and a belief function, that captures the receiver's belief on possible states, is a mapping from $\overline{\mathcal{M}}$ to $\Delta(\Theta)$.

Proposition A.6 *Any prescriptive equilibrium under the regular message space can be extended to an equilibrium under the extended message space.*

Proof. Given an equilibrium (M, D, B) under the regular message space, consider the strategy profile $(\overline{M}, \overline{D}, \overline{B})$ under the extended message space, defined as follows:

- For each state θ, let $\overline{M}(\theta) = \{\theta' \in \Theta \mid \forall d \in \mathcal{D},\, m(d){=}X(d;\theta')\}$ with $m \equiv M(\theta)$.

- We define $\overline{D}$ and $\overline{B}$ over the set of on-path messages $\overline{M}(\Theta)$ as follows: for each state θ, $\overline{D}(\overline{M}(\theta)) = D(M(\theta))$, and $\overline{B}(\overline{M}(\theta)) = B(M(\theta))$. (Note that, for any pair of states (θ, θ'), we have that $M(\theta) = M(\theta')$ if and only if $\overline{M}(\theta) = \overline{M}(\theta')$.)

- Finally, we define $\overline{D}$ and $\overline{B}$ over the set of off-path messages analogously to the proof of Lemma B.1 in Section B.1. For any state θ, let $d^{\dagger}(\theta)$ be the option that maximizes the receiver's utility in state θ, and if two or more utility maximizing options exist, let $d^{\dagger}(\theta)$ be the largest one. Let $m \notin \overline{M}(\Theta)$ and let $d^* = \max_{\theta \in m} d^{\dagger}(\theta)$. Let θ^* be such that $d^{\dagger}(\theta^*) = d^*$. Let $\overline{B}(m)$ be the belief that assigns probability one to θ^*, and let $\overline{D}(m) = d^*$.

It is easily verified that that the triple $(\overline{B}, \overline{D}, \overline{M})$ just defined satisfies the conditions required of a perfect Bayesian equilibrium. By construction, $\overline{B}$ is an adequate belief function for the receiver, which follows Bayes rule for on-path messages whenever possible. Also, the receiver always chooses an optimal option given his beliefs. In every state, by definition, the sender chooses a message that is, at least, optimal among the on-path messages—that is, the sender is never strictly better off deviating with on-path messages. Besides, whenever the state is fully revealed, the receiver chooses an option greater than or equal to the option chosen in the equilibrium (M, D, B). Thus, again by definition of $(\overline{B}, \overline{D}, \overline{M})$, the sender is never strictly better off deviating by revealing the state fully. Finally, in every state θ, if the sender deviates from sending message $M(\theta)$ to sending off-path message $m \notin \overline{M}(\Theta)$, the option selected by the receiver is larger than or equal to the decision that would occur had the sender disclosed the entire state θ. Hence, the sender is never strictly better off announcing an off-path message. Therefore, $(\overline{B}, \overline{D}, \overline{M})$ is an equilibrium. ■

B Auxiliary Results

B.1 On the Existence of Suspicious Beliefs

When constructing an equilibrium, it is convenient to focus on the on-path behavior of the receiver, thus providing an incomplete strategy profile. Under some general conditions, the incomplete profile can be completed with off-path suspicious beliefs and decisions so as to be an equilibrium. The purpose of this section is to formalize this fact.

Consider mappings $M : \Theta \to \mathcal{M}$, $D : M(\Theta) \to \mathcal{D}$ and $B : M(\Theta) \to \Delta(\Omega)$. We interpret these mappings as describing a sender strategy and a receiver strategy and belief function restricted to messages that are expected on the equilibrium path. Consider the following conditions:

1. Receiver beliefs, captured by B, follow Bayes rule whenever possible, and upon observing an on-path message, the receiver maximizes utility: For all $m \in M(\Theta)$,

$$D(m) \in \arg\max_{d \in \mathcal{D}} \mathbb{E}^{B(m)}[u_{\mathrm{R}}(X(d))].$$

2. The sender maximizes utility among the possible equilibrium messages: For all $m \in M(\Theta)$, and all states θ compatible with m, $D(M(\theta)) \leq D(m)$. In addition, the sender is never strictly better off revealing the full state if the receiver, upon observing the full state, were to choose the utility-maximizing option that—if not unique—is worst for the sender.

3. On-path interactions are prescriptive: For all $m \in M(\Theta)$, the sender reveals $D(m)$.

These conditions stipulate the rationality and optimality of sender and receiver behaviors with respect to the set of messages that are expected to be observed on the equilibrium path, and require in addition that the sender is not strictly better off revealing her type.

The tuple (M, D, B) forms an incomplete strategy profile. Lemma B.1, below, asserts that under the above conditions, off-path suspicious beliefs exist and so the incomplete profile can be completed to form an equilibrium.

Lemma B.1 *Assume that Conditions (1)–(3) above are satisfied. The incomplete strategy profile (M, D, B) can be extended to a complete strategy profile in which off-path beliefs and decisions are suspicious. The resulting strategy profile is an equilibrium.*

Proof. First, we extend B and D to the entire message space $\mathcal{M}$ as follows. For any state θ, let $d^\dagger(\theta)$ be the option that maximizes the receiver's utility in state θ, and if two or more utility maximizing options exist, let $d^\dagger(\theta)$ be the largest one. For any $m \notin M(\Theta)$, we define $D(m)$ as

$$D(m) = \max_{\theta \in \Gamma(m)} d^\dagger(\theta).$$

Because there are finitely many decisions, there exists θ^* such that $d^\dagger(\theta^*) = D(m)$. Let $B(m)$ be the belief that assigns probability one to θ^*.

Second, we observe that the off-path beliefs and decisions just defined are suspicious. Indeed, by assumption, for every $m \in M(\Theta)$, and all $\theta \in \Gamma(m)$,

$$D(m) \leq \max\left(\arg\max_{d\in\mathcal{D}} u_{\mathrm{R}}(X(d;\theta))\right).$$

Thus, for all $m \notin M(\Theta)$, all $m' \in M(\Theta)$ such that $\Gamma(m) \cup \Gamma(m') \neq \emptyset$, and all $\theta \in \Gamma(m) \cup \Gamma(m')$,

$$D(m') \leq \max\left(\arg\max_{d\in\mathcal{D}} u_{\mathrm{R}}(X(d;\theta))\right) \leq \max\left(\arg\max_{d\in\mathcal{D}} \mathbb{E}^{B(m)}[u_{\mathrm{R}}(X(d))]\right).$$

Hence, off-path beliefs are suspicious. By definition, off-path decisions are suspicious as well.

Overall, together with the assumptions made on the incomplete strategy profile, since beliefs and decisions are suspicious, the sender best responds given the receiver's strategy. On path, receiver beliefs satisfy Bayes rule and off path, they are consistent with the hard information revealed. Finally, the receiver always makes optimal decisions given his beliefs. Hence, the completed strategy profile (M, D, B) is an equilibrium. ■

Observe that for a given sender strategy, the receiver strategy for on-path messages is almost uniquely determined in an equilibrium with prescriptive interactions—the only flexibility is about how ties are broken, but since ties occur with probability zero, such flexibility is not relevant for our results, as discussed in Section 2. When a sender strategy supports a prescriptive equilibrium, this equilibrium is unique up to off-path beliefs and decisions and tie-break rules. Therefore, it can be convenient to focus on the sender strategy, with the understanding that it is associated with an essentially unique equilibrium. The corollary to Lemma B.1 that follows gives the conditions for the sender strategy to be a prescriptive equilibrium strategy.

Let $M : \Theta \to \mathcal{M}$ be a sender strategy and let $L(m)$ be the smallest minimizer of $d \mapsto \mathbb{E}[u_{\mathrm{R}}(X(d;\theta)) \mid M(\theta){=}m]$. Consider the following conditions:

4. For all $m \in M(\Theta)$, and all states θ compatible with m, $L(M(\theta)) \leq L(m)$. In addition, the sender is never strictly better off revealing the full state if the receiver, upon observing the full state, were to choose the utility-maximizing option that—if not unique—is worst for the sender.

5. For all states θ, the message $M(\theta)$ reveals $L(M(\theta))$.

Corollary B.1 *If a sender strategy M satisfies Conditions (4) and (5), then there exists a receiver strategy D and a belief function B such that (M, D, B)*

is a prescriptive equilibrium, and in this equilibrium, the receiver chooses the sender-preferred option in case of ties.

B.2 On No-advice Equilibrium Messages

Here we show that, in all prescriptive equilibria, communication is always informative when the dimension of the state space is large enough: All sender types reveal the outcome of a non-default option.

Lemma B.2 *If n is large enough, in all prescriptive equilibria, in every state of the world, the sender reveals at least one non-default option.*

Proof. Let (M, D, B) be a prescriptive equilibrium. Suppose, by contradiction, that for some state, the sender reveals the default option (already known to the receiver) and no other options. In this case, since the equilibrium is prescriptive, the receiver chooses the default option upon receiving this trivial message. Hence, in every state, the sender gets her ideal option, d_0.

Let $\mathcal{F}$ be the information—formally a σ-algebra—generated by the sender's message strategy M. By the law of iterated expectation,

$$\mathbb{E}[u_{\mathrm{R}}(X(d_1))] = \mathbb{E}[\,\mathbb{E}[u_{\mathrm{R}}(X(d_1)) \mid \mathcal{F}]].$$

If n is large enough, we have that

$$\mathbb{E}[u_{\mathrm{R}}(X(d_1))] > u_{\mathrm{R}}(0),$$

and so there exists an on-path message m for which

$$\mathbb{E}[u_{\mathrm{R}}(X(d_1;\theta)) \mid M(\theta)=m] > u_{\mathrm{R}}(0).$$

Thus, for n large enough, the receiver is strictly better off choosing a non-default option for at least some on-path messages, which creates a contradiction. ■

B.3 On Non-deceptive Equilibria

In this section we formally state and prove the assertion made in Section 5.2 that in every non-deceptive prescriptive equilibrium, interactions are fully-revealing to the left, and thus Theorem 2 regarding the domination of the interval equilibrium applies.

Lemma B.3 *If a non-deceptive equilibrium is prescriptive, interactions are fully-revealing to the left with probability one.*

Proof. Lemma B.3 does not depend on the particular form of receiver utility.

The proof proceeds by contradiction. Let (M, D, B) be a prescriptive equilibrium where M is non-deceptive, and with positive probability equilibrium interactions are not fully-revealing to the left.

The set of states where the receiver gets his ideal outcome for at least one option has probability zero. Therefore, there exists a state $\theta^\dagger$ such that, in this state, the interaction is not fully-revealing to the left, and $X(d; \theta^\dagger) \neq b$ for all $d \in \mathcal{D}$. Let $m^\dagger = M(\theta^\dagger)$, and let $d^\dagger$ be an option to the left of $D(m^\dagger)$ that is not revealed in $m^\dagger$.

Because M is non-deceptive, in every state compatible with $m^\dagger$, the sender communicates $m^\dagger$. There exists a state θ compatible with $m^\dagger$ such that $X(d^\dagger; \theta) = b$ and for every $d \neq d^\dagger$, $X(d^\dagger; \theta) \neq b$. In this state, in equilibrium, the receiver's decision is $D(m^\dagger)$, yet the sender can guarantee herself the strictly preferred decision $d^\dagger < D(m^\dagger)$ by revealing all the options. Hence we have a contradiction. ■

B.4 On Receiver-Optimal Equilibria

In this section we provide two results about receiver-optimal equilibria discussed in Section 5.4 of the main text.

Proposition B.1 *In all receiver-optimal equilibria, with probability one, the sender reveals all the options to the right of the equilibrium decision.*

Proof. This proposition does not depend on the receiver utility function.

Fix n and consider an equilibrium (M, D, B) that is receiver-optimal: For all states θ, $D(M(\theta))$ maximizes $d \mapsto u_{\mathrm{R}}(d; \theta)$. Let Ω be the set of states θ such that, for all options d, $X(d; \theta) \neq b$. Observe that Ω has probability one. Suppose that, for some $\theta \in \Omega$, $m = M(\theta)$ is such that for some option $d > D(m)$, m does not include the outcome of d. Let θ' be a state with $u_{\mathrm{R}}(X(d; \theta')) > u_{\mathrm{R}}(X(d; \theta))$. Then, in state θ', the sender is strictly better off revealing m, but the receiver is strictly better of choosing d instead of $D(m)$. Hence we have a contradiction. ■

Proposition B.2 *There exists a sequence of first-point conative equilibria such that interactions are conative in the limit and outcomes converge to the receiver's ideal.*

Proof. We prove Proposition B.2 for the case of quadratic receiver utility, its extension to the case of general receiver utility of Section A.1 is straightforward.

For any n, let $\Delta_n = 1/\sqrt[8]{n}$, and denote by Ω_n the set of states θ such that for all d, $X(d;\theta) \notin [b-\Delta_n, b+\Delta_n]$. Thus, Ω_n is the set of states such that the first-point conative strategy with parameter Δ_n reveals the entire state, as opposed to revealing a single option. We show that $\Pr[\theta \in \Omega_n] \to 0$ as $n \to \infty$.

Observe that if $\theta \in \Omega_n$, then at least one of the following is true:

(1) $X(d_n;\theta) < b - \Delta_n$, or

(2) for some $d_i \le d_{n-1}$, $X(d_{i+1};\theta) - X(d_i;\theta) > 2\Delta_n$.

Let Z be an independent random variable that follows the standard normal distribution. First, observe that $X(d_n)$ is normally distributed with mean $\mu\sqrt{n}$ and variance $\sigma^2\sqrt{n}$. Thus,

$$\begin{aligned}\Pr[X(d_n) < b - \Delta_n] &< \Pr[X(d_n) < b] \\ &= \Pr\left[Z < \frac{b-\mu\sqrt{n}}{\sigma\sqrt[4]{n}}\right].\end{aligned}$$

Second, fixing $d_i \le d_{n-1}$,

$$\Pr[X(d_{i+1}) - X(d_i) > 2\Delta_n] = \Pr\left[Z > \frac{2/\sqrt[8]{n} - \mu/\sqrt{n}}{\sigma/\sqrt[4]{n}}\right].$$

For n large enough,

$$\frac{2/\sqrt[8]{n} - \mu/\sqrt{n}}{\sigma/\sqrt[4]{n}} > \frac{\sqrt[8]{n}}{\sigma},$$

and

$$\begin{aligned}\Pr[X(d_{i+1}) - X(d_i) > 2\Delta_n] &< \Pr\left[Z > \sqrt[8]{n}/\sigma\right] \\ &\le e^{-\sqrt[4]{n}/\sigma^2},\end{aligned}$$

where we use the Chernoff bound $\Pr[Z > t] \le e^{-t^2/2}$ if $t \ge 0$. Overall we get, for n large enough,

$$\Pr[\theta \in \Omega_n] \le \Pr\left[Z < \frac{b-\mu\sqrt{n}}{\sigma\sqrt[4]{n}}\right] + ne^{-\sqrt[4]{n}/\sigma^2},$$

and observing that $(b-\mu\sqrt{n})/(\sigma\sqrt[4]{n}) \to -\infty$, we have $\Pr[\theta \in \Omega_n] \to 0$ as $n \to \infty$.

Thus, as n grows large, the first-point conative strategy with parameter Δ_n defined above reveals all the options with vanishing probability. If this strategy supports a prescriptive equilibrium, then it means that interactions become conative with probability one in the limit. It is easily verified that, for any $\Delta > 0$ small enough, the first-point conative strategy with parameter Δ is a prescriptive equilibrium strategy—the case of quadratic utility is studied in Section 4, and the general case follows from the second part of Theorem 1. $\blacksquare$

C Proofs

C.1 Proof of Lemma 2

Lemma 2 does not depend on the particular form of the receiver utility function.

Consider an equilibrium (M, D, B). First, suppose M is almost non-deceptive. Let Ω be any (measurable) set of states. Let m be an on-path message. Note that $M^{-1}(m) \subseteq \Gamma(m)$. And thus, applying Bayes rule,

$$\begin{aligned}\Pr[\theta \in \Omega \mid \theta \in M^{-1}(m)] &= \frac{\Pr[\Omega \cap M^{-1}(m) \mid \Gamma(m)]}{\Pr[M^{-1}(m) \mid \Gamma(m)]} \\ &= \Pr[\Omega \cap M^{-1}(m) \mid \Gamma(m)] \\ &= \Pr[\Omega \mid \Gamma(m)],\end{aligned}$$

where we used compact notation on the right-hand side and where both equalities owe to the fact that $\Pr[M^{-1}(m) \mid \Gamma(m)] = 1$, because M is assumed to be almost non-deceptive. Hence, on-path receiver beliefs are neutral.

Second, suppose that on-path receiver beliefs are neutral. Let m be an on-path message and let $\Omega = \Gamma(m) \backslash M^{-1}(m)$. Then, $\Pr[\Omega \mid \Gamma(m)] = \Pr[\Omega \mid M^{-1}(m)] = 0$, and hence, $\Pr[M^{-1}(m) \mid \Gamma(m)] = 1$. So, the sender strategy is almost non-deceptive.

C.2 Proof of Theorem 1

The statement of Theorem 1 holds for quadratic receiver utility. Here, we prove an extended version of Theorem 1 for general receiver utility under the assumptions of Section A.1, of which quadratic utility is a special case. We define $\underline{x}$ and $\bar{x}$ as in Section A.1.

Theorem C.1 *For all sequences of equilibria that are conative in the limit, the equilibrium outcomes are in the range* $[\underline{x}, b]$ *in the limit. Moreover, for any*

$x \in [\underline{x}, b]$ *there exists a sequence of equilibria that are conative in the limit in which the outcomes converge to* x.

We begin by proving the second part of Theorem C.1.

Let $\underline{x}' \in [\underline{x}, b)$ and let $\bar{x}'$ be the unique outcome greater than b that satisfies $u_{\mathrm{R}}(\bar{x}') = u_{\mathrm{R}}(\underline{x}')$ (existence and uniqueness follow from Assumption (1) in Section A.1). Notice that by concavity of the receiver utility, the outcomes in $[\underline{x}', \bar{x}']$ are better for the receiver than the outcomes outside $[\underline{x}', \bar{x}']$.

Fix n, and consider the (extended) first-point conative strategy in which the sender reveals the smallest option whose outcome belongs to $[\underline{x}', \bar{x}']$ and if no such option exists, the sender reveals all the options. This sender strategy accounts for asymmetric receiver utilities, as explained in Section A.1. Then, suppose that, upon receiving an on-path message, the receiver operates as follows: If the full state is revealed, the receiver chooses the utility-maximizing option. If two or more options are optimal, the receiver chooses the smallest one. Otherwise, the receiver chooses the only revealed option.

The sender's strategy is deceptive but optimal when restricted to on-path messages. If the state includes any outcome in the range $[\underline{x}', \bar{x}']$, then all on-path messages reveal an option whose outcome is in that range. The sender, who prefers smaller to larger options, is best off revealing the smallest option. If instead the state does not include any outcome in the range $[\underline{x}', \bar{x}']$, then there exists only one possible on-path message, so the sender's message is trivially optimal among on-path messages.

The receiver's strategy is also optimal among on-path messages. The case of the sender revealing the full state is immediate. If instead the sender reveals only one option d^* whose outcome x^* belongs to the range $[\underline{x}', \bar{x}']$, then all options to the left of d^* have an outcome outside this range, making the receiver worse off, while beliefs to the right of d^* are neutral, and as $x^* \geq \underline{x}$, by Lemma A.2, the receiver is not strictly better off choosing an option to the right of d^*.

Finally, observe that if some outcome falls in $[\underline{x}', \bar{x}']$, the sender is never strictly better off deviating by revealing the full state: Doing so would yield a decision that is never less than the decision expected in equilibrium.

Therefore, Lemma B.1 applies: The strategies defined above can be completed with appropriate off-path beliefs and decisions to form an equilibrium.

Sending n to infinity, the outcome path becomes distributed as a Brownian motion with drift $\mu > 0$. For any given $\varepsilon > 0$, this Brownian motion almost surely crosses $[\underline{x}', \underline{x}' + \varepsilon]$, so that, as the option set grows large, the probability that the sender reveals a single option whose outcomes is in $[\underline{x}', \underline{x}' + \varepsilon]$ converges to one. Hence, in the sequence of equilibria just defined, equilibrium outcomes converge to $\underline{x}'$ in probability, and the probability of conative interactions converges to one. Finally, the sequence of equilibria of Proposition B.2 guarantees the existence of a sequence of equilibria conative in the limit and whose outcomes converge to b.

This proves the second part of Theorem C.1.

We now prove the first part. Fix any sequence of equilibria such that the probability of conative equilibrium interaction converges to one as n grows large. We show that the probability that equilibrium outcomes fall outside the range $[\underline{x}, b]$ vanishes.

First observe that, as n grows large, the probability that the equilibrium outcome is greater than b vanishes. Indeed, by now familiar arguments, if it was not the case, with non-zero probability, for any sufficiently large n, revealing the full state would yield a smaller receiver decision and thus would be a profitable deviation for the sender.

Second, we show that if $x < \underline{x}$, then as n grows large, the probability that the equilibrium outcome is less than or equal to x vanishes. Considering the n-th equilibrium in the sequence, let Ω_n be the set of states for which equilibrium interaction is conative. Let $\Omega_n(d)$ be the subset of Ω_n for which the sender reveals option d, and $\mathcal{O}_n(d)$ be the set of possible outcomes $X(d;\theta)$ for $\theta \in \Omega_n(d)$. Note that $\mathcal{O}_n(d)$ may be empty for some options d. Thus, if $\theta \in \Omega_n$, in equilibrium the sender reveals the smallest option d whose outcome is in $\mathcal{O}_n(d)$.

Suppose (d^*, x^*) is such that $x^* \in \mathcal{O}_n(d^*)$ and $d^* < d_n$. Let $d > d^*$. The receiver's belief on $X(d)$ after observing that $X(d^*) = x^*$ (and nothing more) is characterized by the cumulative distribution

$$y \mapsto \Pr[X(d;\theta) \leq y \mid X(d^*;\theta){=}x^*, \theta \in \Omega_n(d^*)].$$

Note that

$$\begin{aligned}
&\Pr[X(d;\theta) \leq y \mid X(d^*;\theta){=}x^*, \theta \in \Omega_n(d^*)] \\
&\qquad = \Pr[X(d;\theta) \leq y \mid X(d^*;\theta){=}x^*, \theta \in \Omega_n, X(d';\theta) \notin \mathcal{O}_n(d')\, \forall d' < d^*] \\
&\qquad = \Pr[X(d;\theta) \leq y \mid X(d^*;\theta){=}x^*, \theta \in \Omega_n],
\end{aligned}$$

where the first equality owes to the sender's best response, and the second equality owes to the Markov property of the outcome function.

Because the probability of Ω_n converges to one as $n \to \infty$, the value of $\Pr[X(d^*;\theta) \leq y \mid X(d^*;\theta) = x^*, \theta \in \Omega_n]$ becomes arbitrarily close to the value of $\Pr[X(d^*;\theta) \leq y \mid X(d^*;\theta) = x^*]$.[35] Consequently, the expected utility of the

[35] In general, if, for every n, A_n and B_n are two non-null events and $\Pr[B_n] \to 1$, then

receiver when choosing $d > d^*$, conditionally on observing that $X(d^*) = x^*$ under a conative interaction, becomes arbitrarily close to the expected utility of the receiver when choosing d, but only conditionally on $X(d^*) = x^*$.

That is, as $n \to \infty$, the receiver's beliefs become neutral to the right of the option revealed by the sender. Lemma A.2 then implies that the receiver is strictly better off choosing an option to the right of the option revealed if the outcome of that option is less than $\underline{x}$ and n is large enough.

In addition, as n grows large, with a probability that converges to one, the sender strictly prefers to reveal the full state than to reveal as only data point the outcome of the largest option d_n, because by revealing all the options, the receiver makes a decision which is strictly less than d_n for a set of states whose probability converges to one. Therefore, as n grows large, with probability converging to one, the sender reveals a single option which is not the largest available option.

Putting the last two facts together, it follows that the probability that the equilibrium outcome is less than $\bar{x}$ vanishes as n grows large.

C.3 Proof of Corollary 1

We prove Corollary 1 for the special case of quadratic receiver utility (the result does not always hold in the general case of Section A.1).

As explained in Section 4, if the sender does not provide any information, the receiver chooses the option $d = d^{\mathrm{na}}$ that maximizes his expected utility

$$-(\mathbb{E}[X(d)] - b)^2 - \mathrm{Var}[X(d)] = -(\mu d - b)^2 - \sigma^2 d.$$

Then, as $n \to \infty$, $d^{\mathrm{na}} \to \underline{x}/\mu$, with $\underline{x} = b - \frac{\sigma^2}{2\mu}$.

In the sender-optimal conative equilibrium, the equilibrium outcome converges to $\underline{x}$ as n grows large, and the equilibrium decision is equal to the first hitting time (in the language of stochastic calculus) of the barrier $\underline{x}$ for the outcome path, then distributed as a Brownian motion with drift μ and scale σ. As it turns out, this average hitting time is also equal to $\underline{x}/\mu$ (see, for example, Dixit, 1993, p. 56). In addition, for any sequence of equilibria that are conative in the limit, the equilibrium outcomes eventually belong to the range $[\underline{x}, b]$ by Theorem 1, consequently the equilibrium decisions become lower-bounded by the first hitting time of the barrier $\underline{x}$, and so on average, are no less than $\underline{x}/\mu$.

$|\Pr[A_n|B_n] - \Pr[A_n]| \to 0$.

C.4 Proof of Proposition 2

In this proof, the receiver utility function can be assumed to be quadratic or to satisfy the general assumptions of Section A.1.

The existence of the interval equilibrium is straightforward. Fix n. Suppose the sender follows the interval strategy, and that upon observing an on-path message, the receiver chooses the utility maximizing option among the options revealed by the sender (and, in case of ties, chooses the sender-preferred option). Notice that the interval strategy is a non-deceptive sender strategy, and therefore, the strategy is trivially optimal when restricted to on-path messages. In addition, the receiver forms neutral beliefs after observing an on-path message. Therefore, by the Markov property of the outcome function, the receiver utility for an option d to the right of the right-most revealed option d^r is equal to $\mathbb{E}[u_{\mathrm{R}}(X(d)) \mid X(d^r)]$, which, by definition of the interval strategy, is never strictly higher than the utility of the best option among the options revealed. So the receiver best responds to on-path messages. Of course, the sender is never strictly better off deviating by revealing all the options. Equilibrium existence then follows from Lemma B.1.

By definition of the sender strategy, this interval equilibrium weakly dominates the sender-optimal conative equilibrium for all states. If no-advice is not an equilibrium, then with positive probability the right-most revealed option of the interval equilibrium has an outcome below zero, in which case it strictly dominates the sender-optimal conative equilibrium.

C.5 Proof of Theorem 2

We divide Theorem 2 into two propositions, and prove each separately.

Proposition C.1 *For every sequence of equilibria $\Sigma_1, \Sigma_2, \ldots$ that are conative in the limit, the interval equilibrium strictly dominates Σ_n in the limit as $n \to \infty$.*

Proposition C.2 *If Σ is an equilibrium whose interactions are fully-revealing to the left with probability one and that weakly dominates the sender-optimal conative equilibrium, then the interval equilibrium weakly dominates Σ.*

C.5.1 Proof of Proposition C.1

Throughout this proof, we consider the case of general receiver utility under the assumptions of Section A.1, of which quadratic utility is a special case. We define $\underline{x}$ and $\bar{x}$ as in Section A.1.

Recall that in the limit case $n \to \infty$, outcomes become distributed as a Brownian motion with drift μ and scale σ. In the case of quadratic receiver utility, if the receiver is given the outcome x of decision d, with $x \leq b - \frac{\sigma^2}{2\mu}$, chooses decision $d + \Delta$, with $\Delta \geq 0$, and if the receiver's beliefs are neutral to the right of d, then the receiver's expected utility is equal to

$$-(x + \mu\Delta - b)^2 - \Delta\sigma^2$$

whose maximum is reached for $\Delta \geq 0$ and is equal to

$$-\frac{\sigma^2}{\mu}(b - x) + \frac{\sigma^4}{4\mu^2}.$$

If $x = b - \frac{\sigma^2}{2\mu}$, then the maximum is reached for $\Delta = 0$, and if $x > b - \frac{\sigma^2}{2\mu}$, the maximum is reached for $\Delta > 0$. Besides, the utility of any option d whose outcome is known to be x is $-(x - b)^2$. Therefore, in this limiting case, the value of d^r is the smallest to satisfy

$$b - X(d^r) = \frac{\mu}{\sigma^2} \min\{(b - X(d))^2 : d \in [0, d^r]\} + \frac{\sigma^2}{4\mu}. \tag{C.1}$$

In turn, the receiver decision d^p minimizes $d \mapsto (b - X(d))^2$ for $d \leq d^r$. Of course, if d^r is set to the first hitting time of $b - \frac{\sigma^2}{2\mu}$, then d^r satisfies Equation (C.1). As $\mu > 0$, this first hitting time is almost surely finite, so d^r is finite with probability one.

In the remainder of this proof, we denote by Z an independent random variable that follows the standard normal distribution, and we define

$$f(x, \Delta) = \mathbb{E}\Big[u_{\mathrm{R}}(x + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big],$$

for $x \in \mathbf{R}$ and $\Delta \geq 0$. Note that, given two options $d_j > d_i$, $f(x, d_j - d_i)$ represents the expected utility of a receiver who chooses option $d_k = j$, observes that $X(d_i) = x$, and who forms neutral beliefs to the right of d_i. Let $\Delta^{\mathrm{opt}}(x)$ be a maximizer of f. By Lemma A.1, this maximizer exists and is unique.

The proof utilizes Lemmas C.1 and C.2 below.

Lemma C.1 *We have* $\Delta^{\mathrm{opt}}(x) = O(|x - \underline{x}|)$.

Proof. For $C > 0$, let

$$g(\Delta) = \mathbb{E}\Big[u'_{\mathrm{R}}(\underline{x} - C\Delta + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big] + \frac{\sigma^2}{2\mu}\,\mathbb{E}\Big[u''_{\mathrm{R}}(\underline{x} - C\Delta + \mu\Delta + \sigma\sqrt{\Delta}Z)\Big].$$

By the same arguments as in the proof of Lemma A.1,

$$g'(\Delta) = (\mu - C)\ \mathbb{E}[u_{\mathrm{R}}''(\cdot)] + \frac{\sigma^2}{2}\ \mathbb{E}[u_{\mathrm{R}}'''(\cdot)] + (\mu - C)\frac{\sigma^2}{2\mu}\ \mathbb{E}[u_{\mathrm{R}}''''(\cdot)] + \frac{\sigma^4}{4\mu}\ \mathbb{E}[u_{\mathrm{R}}'''''(\cdot)].$$

To simplify notation, the expression in the parenthesis $(\cdot)$ refers to $\underline{x} - C\Delta + \mu\Delta + \sigma\sqrt{\Delta}Z$. In particular,

$$g'(0) = (\mu - C)u_{\mathrm{R}}''(\underline{x}) + \frac{\sigma^2}{2}u_{\mathrm{R}}'''(\underline{x}) + (\mu - C)\frac{\sigma^2}{2\mu}u''''(\underline{x}) + \frac{\sigma^4}{4\mu}u'''''(\underline{x}).$$

Thus, $g'(0)$ is linear in C. If $C = 0$ and u_{R} is quadratic, then $g'(0) = \mu u_{\mathrm{R}}''(\underline{x}) < 0$ by Assumption (2). If $C = 0$ and u_{R} is not quadratic,

$$\begin{aligned}
\frac{g'(0)}{\mu} &= u''(\underline{x}) + 2\left(\frac{\sigma^2}{2\mu}\right)u'''(\underline{x}) + \left(\frac{\sigma^2}{2\mu}\right)^2 u''''(\underline{x}) \\
&< u''(\underline{x}) + 2\left(\frac{\sigma^2}{2\mu}\right)\sqrt{|u''(\underline{x})u''''(\underline{x})|} + \left(\frac{\sigma^2}{2\mu}\right)^2 u''''(\underline{x}) \\
&= -\left(\sqrt{|u''(\underline{x})|} - \left(\frac{\sigma^2}{2\mu}\right)\sqrt{|u''''(\underline{x})|}\right)^2,
\end{aligned}$$

where we used the inequality $u'''(\underline{x}) < \sqrt{|u''(\underline{x})u''''(\underline{x})|}$, consequence of Assumption (4). Thus, in both the quadratic and non-quadratic case, if $C = 0$, $g'(0)$ is negative, and by linearity of $g'(0)$ in C, $g'(0)$ remains negative if C is positive but small enough. In the remainder of the proof of this lemma, we assume C meets this condition. By definition of $\underline{x}$, $g(0) = 0$, and as $g'(0) < 0$, if Δ is small enough, $g(\Delta) < 0$. Notice that

$$\frac{\partial f(x, \Delta)}{\partial \Delta} = \mu g(\Delta),$$

following the same argument as in the proof of Lemma A.1. As f is strictly concave in Δ by Lemma A.1, $\Delta^{\mathrm{opt}}(x)$ is the only value of Δ satisfying $g(\Delta) = 0$; if $\Delta < \Delta^{\mathrm{opt}}(x)$ then $g(\Delta) > 0$; if $\Delta > \Delta^{\mathrm{opt}}(x)$, then $g(\Delta) < 0$. Since $g(|x - \underline{x}|/C) < 0$ if $|x - \underline{x}|$ is small enough, $\Delta^{\mathrm{opt}}(x) < |x - \underline{x}|/C$ is $|x - \underline{x}|$ is small enough, which concludes the proof. ■

Lemma C.2 *There exists $K > 0$ and $\bar{\delta} \in (0, \underline{x})$ such that for all $\delta \in (0, \bar{\delta})$, and all $\Delta \geq 0$,*

$$u_{\mathrm{R}}(\underline{x} - \delta) > f(\underline{x} - \delta - K\delta^2, \Delta).$$

Proof. By Lemma C.1, there exists $C > 0$ such that if $|x - \underline{x}|$ is small enough, then $\Delta^{\text{opt}}(x) \leq C|x - \underline{x}|$. Let

$$K = \frac{2\mu}{\sigma^2} + 2\sigma^2 C(1 + C\mu)\left|\frac{u_{\mathrm{R}}'''(\underline{x})}{u_{\mathrm{R}}'(\underline{x})}\right|.$$

Applying Taylor's theorem with the Peano form of the remainder, we get

$$u_{\mathrm{R}}(\underline{x} - \delta) = u_{\mathrm{R}}(\underline{x}) - \delta u_{\mathrm{R}}'(\underline{x}) + \frac{\delta^2}{2} u_{\mathrm{R}}''(\underline{x}) + o(\delta^2) = g_1(\delta) + o(\delta^2),$$

with

$$g_1(\delta) = u_{\mathrm{R}}(\underline{x}) - \delta\Big(1 + \frac{\mu}{\sigma^2}\delta\Big) u_{\mathrm{R}}'(\underline{x}),$$

where we use the equality $-u_{\mathrm{R}}''(\underline{x})/u_{\mathrm{R}}'(\underline{x}) = 2\mu/\sigma^2$.

Applying Taylor's theorem a second time but with the Lagrange form of the remainder, we get

$$u_{\mathrm{R}}(x) = u_{\mathrm{R}}(\underline{x}) + (x - \underline{x})u_{\mathrm{R}}'(\underline{x}) + \frac{(x - \underline{x})^2}{2} u_{\mathrm{R}}''(\underline{x}) + \frac{(x - \underline{x})^3}{6} u_{\mathrm{R}}'''(\underline{x}) + Q(x),$$

with

$$Q(x) = \frac{(x - \underline{x})^4}{24} u_{\mathrm{R}}''''(\xi(x))$$

and ξ some real function. Since $u_{\mathrm{R}}'''' \leq 0$, which follows from Assumption (4) as shown in the proof of Lemma A.1, we have $Q \leq 0$. Therefore,

$$f\big(\underline{x} - \delta - K\delta^2, \Delta\big) \leq f_1\big(\underline{x} - \delta - K\delta^2, \Delta\big) + f_2\big(\underline{x} - \delta - K\delta^2, \Delta\big)$$

with

$$\begin{aligned} f_1(x, \Delta) = u_{\mathrm{R}}(\underline{x}) + u_{\mathrm{R}}'(\underline{x})\ \mathbb{E}\Big[x - \underline{x} + \mu\Delta + \sigma\sqrt{\Delta}Z\Big] \\ - \frac{\mu}{\sigma^2} u_{\mathrm{R}}'(\underline{x})\ \mathbb{E}\Big[(x - \underline{x} + \mu\Delta + \sigma\sqrt{\Delta}Z)^2\Big], \end{aligned}$$

where we use again $-u_{\mathrm{R}}''(\underline{x})/u_{\mathrm{R}}'(\underline{x}) = 2\mu/\sigma^2$, and

$$f_2(x, \Delta) = \frac{1}{6} u_{\mathrm{R}}'''(\underline{x})\ \mathbb{E}\Big[(x - \underline{x} + \mu\Delta + \sigma\sqrt{\Delta}Z)^3\Big].$$

We have, after simplification,

$$
\begin{aligned}
f_1(\underline{x} - \delta - K\delta^2, \Delta) &= u_\mathrm{R}(\underline{x}) - \delta(1 + K\delta)u_\mathrm{R}'(\underline{x}) - \frac{\mu}{\sigma^2}\big(\delta + K\delta^2 - \mu\Delta\big)^2 u_\mathrm{R}'(\underline{x}) \\
&\leq u_\mathrm{R}(\underline{x}) - \delta(1 + K\delta)u_\mathrm{R}'(\underline{x}),
\end{aligned}
$$

since $u_\mathrm{R}'(\underline{x}) > 0$ as u_R is strictly concave and maximized at $b > \underline{x}$. Therefore,

$$
g_1(\delta) - f_1\big(\underline{x} - \delta - K\delta^2, \Delta\big) \geq \Big(K - \frac{\mu}{\sigma^2}\Big)u_\mathrm{R}'(\underline{x})\delta^2.
$$

We also have, after simplification,

$$
f_2(\underline{x} - \delta - K\delta^2, \Delta) = -\frac{1}{6}u_\mathrm{R}'''(\underline{x})\big(\delta + K\delta^2 - \mu\Delta\big)\big(3\Delta\sigma^2 + (\delta + K\delta^2 - \mu\Delta)^2\big).
$$

Therefore,

$$
\big|f_2\big(\underline{x} - \delta - K\delta^2, \Delta\big)\big| \leq \frac{1}{6}|u_\mathrm{R}'''(\underline{x})|\big(\delta + K\delta^2 + \mu\Delta\big)\big(3\Delta\sigma^2 + (\delta + K\delta^2 + \mu\Delta)^2\big),
$$

so if $\Delta \leq 2C\delta$ and δ is small enough,

$$
\big|f_2\big(\underline{x} - \delta - K\delta^2, \Delta\big)\big| \leq 2C\sigma^2(1 + C\mu)|u_\mathrm{R}'''(\underline{x})|\delta^2.
$$

Observe that if δ is small enough,

$$
\Delta^{\mathrm{opt}}(\underline{x} - \delta - K\delta^2) \leq C\big(\delta + K\delta^2\big) \leq 2C\delta.
$$

Overall, if δ is small enough, for all Δ,

$$
\begin{aligned}
u_\mathrm{R}(\underline{x} - \delta) - f\big(\underline{x} - \delta - K\delta^2, \Delta\big) &\geq u_\mathrm{R}(\underline{x} - \delta) \\
&\qquad - f\big(\underline{x} - \delta - K\delta^2, \Delta^{\mathrm{opt}}(\underline{x} - \delta - K\delta^2)\big) \\
&\geq 2C\sigma^2(1 + C\mu)|u_\mathrm{R}'''(\underline{x})|\delta^2 \\
&\qquad + \Big(K - \frac{\mu}{\sigma^2}\Big)u_\mathrm{R}'(\underline{x})\delta^2 + o(\delta^2),
\end{aligned}
$$

which by choice of the constant K is positive if δ is small enough, concluding the proof of Lemma C.2. ■

With a slight abuse of notation, we consider the limiting case $n = \infty$, for which the outcome function follows a Brownian motion with drift μ and scale σ. We define τ_a, τ_b as follows:

- τ_a is the smallest option d such that $X(d) = \underline{x}$.

- τ_b is the smallest option d such that

$$\max\{u_{\mathrm{R}}(X(d')) : d' \in [0,d]\} \geq \sup\{\,\mathbb{E}[u_{\mathrm{R}}(X(d') \mid X(d)] : d' \in (d,\infty)\}.$$

Note that $\tau_a \geq \tau_b$, and that $\tau_a < \infty$ with probability one, because $\mu > 0$. In the interval strategy, d^r is τ_b, while τ_a is the option revealed in the sender-optimal conative strategy. We prove that, with probability one, $\tau_a > \tau_b$.

The proof proceeds by contradiction. Let us suppose that $\tau_a < \tau_b$ with probability $\varepsilon > 0$.

Let $K > 0$ and $\bar{\delta} > 0$ be defined as in Lemma C.2. Let $\delta \in (0, \bar{\delta})$. Let τ_δ be the stopping time defined by the smallest option d such that $X(d)$ hits value $\underline{x} - \delta$. We then consider two stopping times based on τ_δ. First, a stopping time τ_L defined as the smallest option $d > \tau_\delta$ such that

$$X(d) = \underline{x} - \delta - K\delta^2.$$

Second, a stopping time τ_U defined as the smallest option $d > \tau_\delta$ such that

$$X(d) = \underline{x}.$$

(Of course, $\tau_U = \tau_a$, but the context in which τ_U is used is different, so we find it convenient to define it independently.)

Since $\mu > 0$, τ_δ is finite with probability one.

Thus, the probability that $\tau_U < \tau_L$ given that $\tau_\delta < \infty$ is the same as the probability that $\tau_U < \tau_L$. By standard results (see, for example, Dixit, 1993, pp. 51–54), for any Brownian motion starting at zero, with drift μ and scale σ, and whose value at zero is $\underline{x} - \delta$, the probability of reaching the upper barrier $V^H \equiv \underline{x}$ before reaching the lower barrier $V^L \equiv \underline{x} - \delta - K\delta^2$ is

$$\frac{e^{-2V^L\mu/\sigma^2} - e^{-2(\underline{x}-\delta)\mu/\sigma^2}}{e^{-2V^L\mu/\sigma^2} - e^{-2V^H\mu/\sigma^2}} = \frac{e^{-2(\underline{x}-\delta-K\delta^2)\mu/\sigma^2} - e^{-2(\underline{x}-\delta)\mu/\sigma^2}}{e^{-2(\underline{x}-\delta-K\delta^2)\mu/\sigma^2} - e^{-2\underline{x}\mu/\sigma^2}} = K\delta + o(\delta).$$

Hence, the probability that $\tau_U < \tau_L$ converges to zero as $\delta \to 0$. We have $\tau_U = \tau_a$, and if $\tau_b > \tau_\delta$, then we also have $\tau_L \geq \tau_b$. As $\tau_a < \tau_b$ with probability ε and $\tau_\delta < \tau_a$ by definition, it follows that $\tau_U < \tau_L$ with probability ε. This contradicts the fact that the probability that $\tau_U < \tau_L$ converges to zero as $\delta \to 0$.

Therefore, with probability one, $\tau_a > \tau_b$ in this limiting case. Theorem 2 then follows from the fact that as $n \to \infty$, outcomes become distributed as a Brownian motion with drift μ and scale σ.

C.5.2 Proof of Proposition C.2

Throughout this proof, we consider the case of general receiver utility under the assumptions of Section A.1, of which quadratic utility is a special case. We define $\underline{x}$ and $\bar{x}$ as in Section A.1. In addition, for every n, let $\underline{x}^n$ be defined as the value of $x < b$ that satisfies, for every $i = 1, \ldots, n$,

$$\mathbb{E}[u_{\mathrm{R}}(X(d_i)) \mid X(d_{i-1}){=}x] = u_{\mathrm{R}}(x).$$

Existence and uniqueness of $\underline{x}^n$ are guaranteed by Lemmas A.1 and A.2. If $x < \underline{x}^n$, then $\mathbb{E}[u_{\mathrm{R}}(X(d_i)) \mid X(d_{i-1}){=}x] > u_{\mathrm{R}}(x)$. Note that as $n \to \infty$, $\underline{x}^n \to \underline{x}$ and for the case of quadratic receiver utility, $\underline{x}^n = b - \Delta^{\max}$, where $\Delta^{\max}$ is defined in Equation (5) of Section 4.

We begin by introducing several elements of language that will be used throughout the proof.

Sub-message/Super-message A message m' is a *sub-message* of a message m when, for every option d, m always reveals d if m' reveals d. A message m' is a *super-message* of a message m if m is a sub-message of m'.

Simple/Complex messages A message m is *simple* when it reveals exactly the outcomes of consecutive options $d_i, d_{i+1}, \ldots, d_j$ for some $i < j$. A message m is *complex* if there are options $d_i < d_j < d_k$ such that m reveals both d_i and d_k but does not reveal d_j.

Message length The *length* of a message is the number of options it reveals.

Principal/Sequel messages Every complex message m is composed of a principal and a sequel. The *principal* of m is the longest sub-message of m that is simple and reveals d_0 (if m does not reveal d_0, then the principal is the empty message). The *sequel* of m is the sub-message of m that excludes the principal but includes the remaining information: It reveals all the options that m reveals except for the options already revealed in the principal of m.

Lemma C.3 *Let (M, D, B) be an equilibrium. If m is a simple on-path message that reveals at least the default option and such that, on a probability one set of states θ conditionally on $\theta \in \Gamma(m)$, $M(\theta)$ is a super-message of m, then, on a probability one set of states θ conditionally on $\theta \in \Gamma(m)$, the interval equilibrium weakly dominates (M, D, B).*

Proof. Let (M, D, B) be an equilibrium and let m_0 be a simple on-path message that satisfies the conditions of Lemma C.3. Suppose m_0 discloses exactly the options $d_0, \ldots, d_I$. Let $\mathcal{S}$ be the set of all the super-messages m of m_0 with $m \in M(\Theta)$ and let

$$\Omega = \bigcup_{m \in \mathcal{S}} M^{-1}(m).$$

Let $m \in \mathcal{S}$. In every state $\theta^\dagger \in \Omega \subseteq \Gamma(m_0)$, the sender has the possibility to send m_0 and induce a decision in $\{d_0, \ldots, d_I\}$, thus $D(m) \in \{d_0, \ldots, d_I\}$, which implies

$$\max_{0 \le i \le I} u_{\mathrm{R}}(X(d_i; \theta^\dagger)) \ge \max_{I < j \le n} \mathbb{E}[u_{\mathrm{R}}(X(d_j; \theta)) \mid M(\theta){=}m].$$

Let $\mathcal{F}$ be the information (σ-algebra) generated by M. By the law of iterated expectations, we have, for every $d \in \mathcal{D}$,

$$\mathbb{E}[u_{\mathrm{R}}(X(d; \theta)) \mid \theta \in \Omega] = \mathbb{E}[\, \mathbb{E}[u_{\mathrm{R}}(X(d; \theta)) \mid \mathcal{F}, \theta \in \Omega]].$$

Hence,

$$\max_{0 \le i \le I} u_{\mathrm{R}}(X(d_i; \theta^\dagger)) \ge \max_{I < j \le n} \mathbb{E}[u_{\mathrm{R}}(X(d_j; \theta)) \mid \theta \in \Omega].$$

Finally, as $\Pr[\theta \in \Omega \mid \theta \in \Gamma(m_0)] = 1$, we also have, for every $d \in \mathcal{D}$,

$$\mathbb{E}[u_{\mathrm{R}}(X(d; \theta)) \mid \theta \in \Omega] = \mathbb{E}[u_{\mathrm{R}}(X(d; \theta)) \mid \theta \in \Gamma(m_0)],$$

so

$$\max_{0 \le i \le I} u_{\mathrm{R}}(X(d_i; \theta^\dagger)) \ge \max_{I < j \le n} \mathbb{E}[u_{\mathrm{R}}(X(d_j; \theta)) \mid \theta \in \Gamma(m_0)],$$

which implies that for every state in Ω, the interval equilibrium yields an equilibrium decision that is no greater than the equilibrium decision of (M, D, B).
■

We now proceed to the main proof. Let us fix n and consider any equilibrium $\Sigma = (M, D, B)$ that satisfies the conditions of Proposition C.2. Let (M^I, D^I, B^I) denote the interval equilibrium. If $\underline{x}^n < 0$ then the interval equilibrium implements the sender's first best (no-advice is an equilibrium) and Proposition C.2 is immediately satisfied. In the rest of this proof we assume $\underline{x}^n > 0$.

We prove by induction on N the following statement:

> *For every $N \le n-1$, if, with probability one, M reveals at least the options $d_0, \ldots, d_N$, then the interval equilibrium weakly dominates Σ.*

Observe that Proposition C.2 is included in the case $N = 0$. If $N = n - 1$, then, with probability one, the sender communicates a simple message, and by Lemma C.3, the interval equilibrium weakly dominates Σ.

By contradiction, we prove that if the induction statement holds for N then it holds for $N-1$. Suppose that the induction statement holds for N, that, with probability one, M reveals at least the options $d_0, \ldots, d_{N-1}$ and Σ weakly dominates the sender-optimal conative equilibrium, and that Σ strictly dominates the interval equilibrium with positive probability. Then, there exists a message $m_0 \in \mathcal{M}$ (not necessarily an on-path message) revealing exactly the options $d_0, \ldots, d_{N-1}$ such that, conditionally on the state being compatible with m_0,

(1) with probability one, interactions are fully-revealing to the left, M reveals at least $d_0, \ldots, d_{N-1}$ and Σ weakly dominates the sender-optimal conative equilibrium; and,

(2) with positive probability, interactions are fully-revealing to the left, M reveals $d_0, \ldots, d_{N-1}$ but does not reveal d_N, and Σ strictly dominates the interval equilibrium.

In Steps 1–4 below, we show that there exists some $K \in \{N+1, \ldots, n\}$ and a message m_1 that reveals exactly $d_0, \ldots, d_{K-1}$ such that, in the equilibrium Σ, m_1 is an on-path message, $D(m_1) = d_{K-1}$, but the receiver is strictly better off choosing the unrevealed option d_K, thus contradicting the assumption that Σ is an equilibrium.

Step 1: Let $\mathcal{Q}$ be the set of complex messages in $M(\Theta)$ whose principal is m_0. In this first step, we establish Equations (C.2)–(C.4) below:

$$m_0 \notin M(\Theta), \tag{C.2}$$
$$m_0(d_i) < \underline{x}^n \quad \text{for all} \quad i \leq N-1, \tag{C.3}$$
$$0 < \Pr[M(\theta) \in \mathcal{Q} \mid \theta \in \Gamma(m_0)], \tag{C.4}$$

in particular, $\mathcal{Q}$ is not empty.

Equation (C.2) is implied by Lemma C.3.

Let Ω be the set of states compatible with m_0 such that for all $\theta \in \Omega$, $M(\theta)$ reveals $d_0, \ldots, d_{N-1}$ but does not reveal d_N, and, in state θ, Σ strictly dominates the interval equilibrium.

Equation (C.3) owes to the fact that, if $\theta \in \Omega$, then $D^I(M^I(\theta)) \geq N$ but $D(M(\theta)) \leq N-1$, so that $m_0(d_i) = X(d_i; \theta) < \underline{x}^n$ if $i \leq N-1$.

Finally, by assumption,

$$\Pr[\theta \in \Omega \mid \theta \in \Gamma(m_0)] > 0,$$

and by Lemma C.3, if $\theta \in \Omega$ then $M(\theta)$ is a complex message whose principal is m_0. We then have

$$\Pr[M(\theta) \in \mathcal{Q} \mid \theta \in \Gamma(m_0)] \geq \Pr[\theta \in \Omega \mid \theta \in \Gamma(m_0)],$$

which implies Equation (C.4).

Step 2: Let $\mathcal{Q}_k$ be the subset of $\mathcal{Q}$ such that each message of Q_k reveals d_k but does not reveal any option $d_N, \ldots, d_{k-1}$. Observe that $\{\mathcal{Q}_{N+1}, \ldots, \mathcal{Q}_n\}$ is a partition of $\mathcal{Q}$. Thus, by Equation (C.4), there exists k such that

$$\Pr[M(\theta) \in \mathcal{Q}_k \mid \theta \in \Gamma(m_0)] > 0.$$

Let K be the largest such index.

For every $\delta > 0$, let $\mathcal{M}_\delta$ be the set of simple messages m_1 that extend m_0 in the following way:

(i) if $k \leq N-1$, $m_1(d_k) = m_0(d_k)$;

(ii) if $N-1 < k < K$, $m_1(d_k) = \underline{x} - \dfrac{K-(k+1)}{K-N}\delta + \varepsilon_k$, for some $\varepsilon_k \in [0, \delta/N)$;

(iii) if $k \geq K$, $m_1(d_k) = \emptyset$;

(iv) $m_1(d_i) < m_1(d_j) < \underline{x}^n$ for every $i \in \{0, \ldots, N-1\}$ and every $j \in \{N, \ldots, K-2\}$;

(v) finally, $\underline{x}^n \leq m_1(d_{K-1}) < b$.

Observe that if $m_1 \in M(\Theta) \cap \mathcal{M}_\delta$, then $D(m_1) = d_{K-1}$. The lemma below shows that $M(\Theta) \cap \mathcal{M}_\delta$ is not empty.

Lemma C.4 *For every $\delta > 0$, $M(\Theta) \cap \mathcal{M}_\delta \neq \emptyset$.*

Proof. Consider the set of all the states that are compatible with a message of $\mathcal{M}_\delta$ and such that the outcomes of options to the right of d_K yield a strictly higher receiver utility than the outcomes of the other options. Conditionally on the state being compatible with m_0, this set of states has positive probability, so by our condition (1) above, in at least one of these states $\theta^\dagger$ the sender reveals at least $d_0, \dots, d_{N-1}$ in Σ, interactions are fully-revealing to the left, and the receiver decision is not greater than that of the interval equilibrium in the same state.

Let $m^\dagger = M(\theta^\dagger)$. Since m_0 is not an on-path message (see Equation (C.2)), $m^\dagger$ reveals some option(s) to the right of d_{N-1}. Naturally $m^\dagger$ cannot reveal any option to the right of d_{K-1} for the interval equilibrium to be weakly dominated. Then, because interactions are fully-revealing to the left, and by monotonicity of the outcomes over options $d_N, \dots, d_{K-1}$, $m^\dagger$ must be a simple message that reveals exactly the options $d_0, \dots, d_k$ for $k \leq K-1$.

Suppose $k < K-1$. Observe that for all $i = d_N, \dots, d_{K-2}$, as $X(d_{i-1}; \theta^\dagger) < \underline{x}^n$, we have,

$$\mathbb{E}\big[u_{\mathrm{R}}(X(d_i)) \mid X(d_{i-1}; \theta^\dagger)\big] > u_{\mathrm{R}}(X(d_{i-1}; \theta^\dagger).$$

Thus, if receiver beliefs are neutral upon receiving $m^\dagger$, the receiver would be strictly better off deviating from the prescribed equilibrium decision. If beliefs are not neutral upon receiving $m^\dagger$, then for every state compatible with $m^\dagger$, the sender sends a super-message of $m^\dagger$, as implied by the prescriptive interactions and noting that the receiver decision is the largest option revealed in $m^\dagger$. By the same argument as in the proof of Lemma C.3, for at least one such message, the receiver is strictly better off deviating—as otherwise the law of iterated expectations would be violated.

Therefore, $k = K-1$ and $M(\Theta) \cap \mathcal{M}_\delta \neq \emptyset$. ■

Step 3: In this step, we establish the inequality

$$\Pr[M(\theta) \in \mathcal{Q} \mid \theta \in \Gamma(m_1)] = \Pr[M(\theta) \in \mathcal{Q}_K \mid \theta \in \Gamma(m_1)] > 0. \tag{C.5}$$

If $k < K$, then $M^{-1}(Q^k) \cap \Gamma(m_1) = \emptyset$. Thus,

$$\Pr\big[M(\theta) \in Q^k \mid \theta \in \Gamma(m_1)\big] = 0.$$

If $k > K$, then by definition of K,

$$\Pr[M(\theta) \in Q_k \mid \theta \in \Gamma(m_0)] = 0,$$
$$\Pr[M(\theta) \in Q_K \mid \theta \in \Gamma(m_0)] > 0.$$

Observing that the (Gaussian) conditional distributions of $(\theta_K, \dots, \theta_n)$ given $\Gamma(m_0)$ and $\Gamma(m_1)$, respectively, are equivalent, we also have

$$\Pr[M(\theta) \in Q_k \mid \theta \in \Gamma(m_1)] = 0, \Pr[M(\theta) \in Q_K \mid \theta \in \Gamma(m_1)] > 0.$$

Equation (C.5) then follows from the fact that $\{\mathcal{Q}_{N+1}, \ldots, \mathcal{Q}_n\}$ is a partition of $\mathcal{Q}$.

Step 4: This step establishes that there exists $C > 0$ such that for every $\delta > 0$ and every $m_1 \in \mathcal{M}_\delta$,

$$\mathbb{E}[u_{\mathrm{R}}(X(d_K;\theta)) \mid M(\theta) \notin \mathcal{Q}, \theta \in \Gamma(m_1)] > u_{\mathrm{R}}(\underline{x}^n) + C.$$

We have

$$\begin{aligned}
&\mathbb{E}[u_{\mathrm{R}}(X(d_K)) \mid X(d_{K-1}){=}\underline{x}^n + \varepsilon_{K-1}] \\
&= \mathbb{E}[u_{\mathrm{R}}(X(d_K;\theta)) \mid \theta \in \Gamma(m_1)] \\
&= \Pr[M(\theta) \notin \mathcal{Q} \mid \theta \in \Gamma(m_1)]\ \mathbb{E}[u_{\mathrm{R}}(X(d_K;\theta)) \mid M(\theta) \notin \mathcal{Q}, \theta \in \Gamma(m_1)] \\
&\quad + \sum_{k=1}^{K} \Pr[M(\theta) \in \mathcal{Q}_k \mid \theta \in \Gamma(m_1)]\ \mathbb{E}[u_{\mathrm{R}}(X(d_k;\theta)) \mid M(\theta) \in \mathcal{Q}_k, \theta \in \Gamma(m_1)].
\end{aligned}$$

By Step 3, we have

$$\begin{aligned}
&\sum_{k=1}^{K} \Pr[M(\theta) \in \mathcal{Q}_k \mid \theta \in \Gamma(m_1)]\ \mathbb{E}[u_{\mathrm{R}}(X(d_k;\theta)) \mid M(\theta) \in \mathcal{Q}_k, \theta \in \Gamma(m_1)] \\
&\quad = \Pr[M(\theta) \in \mathcal{Q}_K \mid \theta \in \Gamma(m_1)]\ \mathbb{E}[u_{\mathrm{R}}(X(d_k;\theta)) \mid M(\theta) \in \mathcal{Q}_K, \theta \in \Gamma(m_1)].
\end{aligned}$$

By continuity, there exists $\alpha > 0$ such that, for all $\delta > 0$ small enough and all $m_1 \in \mathcal{M}_\delta$,

$$\Pr[M(\theta) \in \mathcal{Q} \mid \theta \in \Gamma(m_1)] > \alpha \Pr[M(\theta) \in \mathcal{Q}_K \mid \theta \in \Gamma(m_0)].$$

In addition, the smoothness conditions imposed on u_{R} imply that

$$x \mapsto \mathbb{E}[u_{\mathrm{R}}(X(d_K)) \mid X(d_{K-1}){=}x]$$

is continuously differentiable, so that there exists $A > 0$ such that if ε is positive but small enough,

$$\mathbb{E}[u_{\mathrm{R}}(X(d_K)) \mid X(d_{K-1}){=}\underline{x}^n + \varepsilon] \geq u_{\mathrm{R}}(\underline{x}^n) - A\varepsilon.$$

Thus, for all $\delta > 0$ small enough and all $m_1 \in \mathcal{M}_\delta$,

$$\mathbb{E}[u_{\mathrm{R}}(X(d_K)) \mid X(d_{K-1}){=}m_1(d_{K-1})] \geq u_{\mathrm{R}}(\underline{x}^n) - A\varepsilon_K.$$

If $m \in \mathcal{Q}_K$, then $D(m) \leq d_{N-1}$, and so

$$u_{\mathrm{R}}(m(d_K)) \leq \max_{0 \leq i \leq N-1} u_{\mathrm{R}}(m(d_i)) < u(\underline{x}^n),$$

so there exists $B > 0$ such that

$$\mathbb{E}[u_{\mathrm{R}}(X(d_K;\theta)) \mid M(\theta) \in \mathcal{Q}_K] < u_{\mathrm{R}}(\underline{x}^n) - B.$$

Step 5: If $\theta \in \Gamma(m_1)$ then either (1) $M(\theta) \in \mathcal{Q}$ in which case $D(M(\theta)) \leq d_{N-1}$ and Σ strictly dominates the interval equilibrium in state θ, or (2) $M(\theta) \notin \mathcal{Q}$ in which case $D(M(\theta)) = D^I(M^I(\theta)) = d_{K-1}$.

Thus,

$$\mathbb{E}[u_{\mathrm{R}}(X(d_K;\theta)) \mid M(\theta) \notin \mathcal{Q}, \theta \in \Gamma(m_1)] = \mathbb{E}[u_{\mathrm{R}}(X(d_K;\theta)) \mid M(\theta){=}m_1].$$

By Step 4,

$$\mathbb{E}[u_{\mathrm{R}}(X(d_K;\theta)) \mid M(\theta){=}m_1] > u_{\mathrm{R}}(\underline{x}) + C.$$

The smoothness conditions imposed on u_{R} mean that if δ is small enough, for every $m_1 \in \mathcal{M}_\delta$,

$$u_{\mathrm{R}}(D(m_1)) = u_{\mathrm{R}}(\underline{x} + \varepsilon_{K-1}) \leq u_{\mathrm{R}}(\underline{x}) + D\varepsilon_{K-1} \leq u_{\mathrm{R}}(\underline{x}) + D\delta,$$

and thus, if δ is small enough,

$$\mathbb{E}[u_{\mathrm{R}}(X(d_K;\theta)) \mid M(\theta){=}m_1] > u_{\mathrm{R}}(D(m_1))$$

and there exists a profitable deviation for the receiver. Therefore, Σ cannot be an equilibrium.

C.6 Proof of Lemma 3

Lemma 3 does not depend on the particular form of receiver utility.

The proof proceeds by contradiction. Let (M, D, B) be an equilibrium where M is non-deceptive but that is not near-prescriptive.

There exists a state $\theta^\dagger$, such that the equilibrium decision in state $\theta^\dagger$ is an undisclosed option d_k, and such that option d_{k-1} is also undisclosed. Let $m^\dagger = M(\theta^\dagger)$.

Since M is non-deceptive, beliefs are neutral, so the variance of $X(d_k)$ given the receiver's information upon observing $m^\dagger$ is positive and thus for every option d, it must be the case that

$$\mathbb{E}\big[u_{\mathrm{R}}(X(d);\theta) \bigm| M(\theta){=}m^{\dagger}\big] < u_{\mathrm{R}}(b).$$

Let

$$\alpha = \max_{d\in\mathcal{D}} \mathbb{E}\big[u_{\mathrm{R}}(X(d);\theta) \bigm| M(\theta){=}m^{\dagger}\big].$$

Since the sender strategy is non-deceptive, in every state compatible with $m^{\dagger}$, the sender sends $m^{\dagger}$. Thus, there exists a state θ that satisfies

$$u_{\mathrm{R}}(X(d;\theta)) \leq \alpha < u_{\mathrm{R}}(b)$$

for every $d \in \mathcal{D}$, $d \neq d_{k-1}$, and $X(d_{k-1};\theta) = b$. And thus, in this state θ, the sender is strictly better off revealing all the options. Hence we have a contradiction.

References

Ben-Porath, E., E. Dekel, and B. L. Lipman (2019). Mechanisms with evidence: Commitment and robustness. *Econometrica 87*(2), 529–566.

Callander, S. (2008). A theory of policy expertise. *Quarterly Journal of Political Science 3*(2), 123–140.

Callander, S. (2011). Searching and learning by trial and error. *The American Economic Review 101*(6), 2277–2308.

Callander, S. and N. Matouschek (2019). The risk of failure: Trial and error learning and long-run performance. *American Economic Journal: Microeconomics 11*(1), 44–78.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica 50*(6), 1431–1451.

Dixit, A. (1993). *The Art of Smooth Pasting*. Routledge.

Dye, R. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research 23*(1), 123–145.

Dziuda, W. (2011). Strategic argumentation. *Journal of Economic Theory 146*(4), 1362–1397.

Garfagnini, U. and B. Strulovici (2016). Social experimentation with interdependent and expanding technologies. *The Review of Economic Studies 83*(4), 1579–1613.

Gibbons, R., N. Matouschek, and J. Roberts (2013). Decisions in organizations. In R. Gibbons and J. Roberts (Eds.), *Handbook of Organizational Economics*, pp. 373–431. Princeton University Press.

Gilligan, T. and K. 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.

Glazer, J. and A. Rubinstein (2004). On optimal rules of persuasion. *Econometrica 72*(6), 1715–1736.

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics 24*(3), 461–483.

Hart, S., I. Kremer, and M. Perry (2017). Evidence games: Truth and commitment. *The American Economic Review 107*(3), 690–713.

Jakobson, R. (1960). Linguistics and poetics. In T. Sebeok (Ed.), *Style in Language*, pp. 350–377. MIT Press.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *The American Economic Review 101*(6), 2590–2615.

Matthews, S. A. and A. Postlewaite (1985). Quality testing and disclosure. *The RAND Journal of Economics 16*(3), 328–340.

Meyer, M. (2017). Communication with self-interested experts. Lecture notes, Nuffield College, Oxford University.

Milgrom, P. and J. Roberts (1986). Relying on the information of interested parties. *The RAND Journal of Economics 17*(1), 18–32.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics 12*(2), 380–391.

Rappoport, D. (2020). Evidence and skepticism in verifiable disclosure games. SSRN Working Paper No. 2978288.

Rayo, L. and I. Segal (2010). Optimal information disclosure. *Journal of Political Economy 118*(5), 949–987.

Seidmann, D. J. and E. Winter (1997). Strategic information transmission with verifiable messages. *Econometrica 65*(1), 163–169.

Shin, H. (2003). Disclosures and asset returns. *Econometrica 71*(1), 105–133.

Sobel, J. (2013). Giving and receiving advice. In D. Acemoglu, M. Arellano, and E. Dekel (Eds.), *Advances in Economics and Econometrics: Tenth World Congress*, pp. 373–431. Princeton University Press.

Weber, M. (1922). *Economy and Society.* University of California Press.

Weber, M., H. Gerth, and C. Mills (1958). *From Max Weber: Essays in Sociology.* Oxford University Press.