The Power of Referential Advice*

Steven Callander[†] Nicolas S. Lambert[‡] Niko Matouschek[§]

August 2020

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 changes communication as it creates an expectation for additional information that ties the hands of the expert. This can hurt the expert as she may be compelled to reveal more information than she would like, up to and including full revelation. It can also help the expert as, by tying her hands, her messages become more credible. 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 Bardi, 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. Referential advice matters to communication because it creates an expectation for more information and this expectation ties the hands of the expert. Done poorly, the expert is compelled to provide more information than she would like, up to and including full revelation of her expertise. Done well, however, the expert is able to shape the decision maker's beliefs in a way that systematically sways decisions in her favor. We identify an equilibrium in which the expert is always strictly better off when she provides referential advice than she is 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. As a modeling tool, this formulation has yielded considerable insight, yet it conflicts with our intuitions about what it means to be an expert. Weber (1922) famously emphasized the large gap in knowledge between an expert and a decisionmaker and argued that this gap is the source of an expert's power, observing that "the political master finds himself in the position of the 'dilettante' who stands opposite the 'expert'."[1]

We take inspiration from Weber and introduce a novel, richer conception of expertise, one that allows space for referential advice. 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.

The key novelty of this approach is that it allows the expert to communicate precisely yet imperfectly. The spillover of information from expert advice is incomplete. This means that the doctor not only has more information than the patient, but when she prescribes medicine for a migraine, that prescription does not reveal all of her expertise about, say, heart disease or any other illness.[2]

In the canonical model, in contrast, the spillover of information from a recommendation is complete. In revealing the ideal treatment for one malady, a doctor reveals the ideal treatment for all maladies. 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 and we show that the expert is able to do so 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.

The existence of recommendation-only *conative* equilibria imply that com-

[1]The logic for this advantage is described by Bendor et al. (1985, p.1042): "A bureau's influence rests... as Weber noted, [in] its control of information, its ability to manipulate... information about the consequences of different alternatives."

[2]Subsequent literature has generalized the informational structure of the classical models and explored the limits of unraveling, although in directions different from ours. We discuss this work below.

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

plex 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. Our main result is that this is not true.

Constraining conative equilibria is the fundamental problem of informational spillover. The expert can minimize the amount of information she reveals, but she can't prevent that information from spilling over into other options. Should the decisionmaker reject the expert's recommendation, he can still use the information contained in that recommendation for his own benefit.[4] 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 all conative equilibria is equal to or larger than what the decisionmaker would choose in the absence of expertise. This means that a risk-neutral expert wanting smaller choices is indifferent between her most preferred conative equilibrium and not having expertise at all.[5]

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.

The influence of referential information is not that the expert reveals bad outcomes when the realized states of the world are unfavorable to the decisionmaker and stays quiet otherwise (which would fall to standard adverse selection arguments). Referential advice works instead through a different channel. In our model, an ideal option for the decisionmaker almost surely exists and, therefore, no amount of bad information will convince him otherwise. The decisionmaker's core problem, however, is that he does not know which

[4]For example, the persuasiveness of an overly-cautious doctor is limited because when she prescribes 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 is enough.

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

option is his ideal. By strategically providing referential information, and by exploiting the correlation across states, the expert is able to manipulate and spread out the decisionmaker's uncertainty such that no single decision is particularly attractive. Combining this ability with a carefully chosen recommendation, the expert is able to present a picture of the world that persuades the decisionmaker to accept a recommendation that he otherwise wouldn't were it presented alone and without referential information. The power of this result is that the logic holds not just some of the time but all of the time: The expert is always able to induce a more favorable decision with referential advice than with a recommendation alone.

At its core, the difference between referential and conative equilibria is credibility. Conative advice cannot carry the weight of both persuasion and dissuasion because in minimizing communication, the expert, ironically, has too much freedom to communicate, freedom that she can use to misrepresent her true type and deceive the decisionmaker. In many strategic environments that freedom is powerful. In a game of strategic communication that freedom is destructive. In conative equilibria the freedom to deceive undermines the expert's ability to persuade. Referential advice is credible because it ties the hands of the expert. In referential equilibria the expert potentially provides a lot of information. Some of it is used to persuade and dissuade directly, the remainder is important in that it makes imitation difficult. Precisely because imitation is difficult, the expert is unable to deceive the decisionmaker, and this is what makes her message credible. Credibility is what makes the persuasion and dissuasion possible and, when constructed in just the right way, induces the decisionmaker to accept a recommendation that he otherwise wouldn't.

Our model of expertise introduces a novel application of stochastic processes. Formally, we analyze a decision problem with a finite set of options but we allow that set to become arbitrarily large such that in the limit it approximates the real line. The realization of the correlated state variables then produces a mapping from options to outcomes that is a discrete stochastic process. We construct the correlation across states such that in the limit this path approximates a Brownian motion and we interpret the variance of the Brownian motion as parameterizing the *complexity* of the decision problem and, as the expert knows the outcome of each option, as the complexity of expertise. This construction offers advantages in richness, tractability, and realism, and we leverage these to provide insights into strategic communication.[6] The canonical models of expertise as a single piece of information (Milgrom 1981; Crawford and Sobel

[6]The Brownian motion representation of uncertainty has also recently found application in models of search (Callander 2011; Garfagnini and Strulovici 2016).

1982) correspond to the special case of our model in which the correlation across states is perfect and the mapping is perfectly linear.

Modeling expertise in a richer, more realistic way, significantly expands the space of types, messages, and beliefs. This richness leads to a multiplicity of equilibria. It also means that a complete characterization of referential equilibria is beyond our reach. We view this as a reasonable trade-off of analytic power for modeling richness. That said, we do fully characterize what the expert can achieve through a recommendation alone. We then show that a particular referential equilibrium strictly dominates *all* conative equilibria, such that the expert is always made better off by providing referential advice, and that this equilibrium possesses appealing properties relative to other referential equilibria.

Related Literature

Experts play a role in almost every aspect of economic, social, and political life, and intuition strongly suggests that they benefit from their expertise. Indeed, Weber et al. (1958) concluded that in politics the power of experts was preeminent: "Under normal conditions, the power position of a fully developed bureaucracy is always overtowering." Documenting this advantage empirically, however, can be challenging. Nevertheless, over time, broad and compelling evidence has accumulated that experts not only influence decisions but that they shape them to their personal advantage: Division managers manipulate headquarters into funding too many projects (Milgrom and Roberts 1988); realtors manipulate homeowners into selling too quickly and cheaply (Levitt and Syverson 2008); and OBGYNs manipulate patients into having too many C-sections (Gruber and Owings 1996), among other evidence. The contribution of our model is to provide a novel theoretical foundation for this expert advantage even when the decisionmaker knows the expert does not have his interests at heart.

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; 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.

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.[7]

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.

The alternative approach to communication with hard information is cheap talk communication (Crawford and Sobel 1982). The analysis with cheap-talk in our environment is trivial and immediate: No informative equilibria exist. If different messages induce different decisions by the receiver then, because the expert cares only about the action taken, it follows that at least one message must be strictly suboptimal. Callander (2008) shows that informative communication is possible if the expert cares instead about the outcome. Analyzing the limit case of our model in which the mapping is a Brownian motion, he shows that this creates a common interest—both expert and receiver wish to avoid extreme outcomes—and that this common interest supports

[7]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.

informative advice. Nevertheless, in the equilibrium he identifies, referential advice plays no role and communication is purely conative: In equilibrium, the sender recommends an action that maps into her own ideal point and, as long as the sender's preferences are not too different from his, the receiver implements that option. The sender has no incentive to deviate as she obtains her ideal outcome, and the receiver implements the recommendation as he prefers the sender's ideal outcome rather than face the risk of choosing on his own. This balance is not relevant for the preferences we analyze here and 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 an absence of commitment. In our model, neither the receiver nor the sender 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.

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; Gibbons et al. 2013). We follow the convention of using the female pronoun for the sender and the male pronoun for the receiver.

Technology: There is a set of options $\mathcal{D} = \{d_0, d_1, \ldots, d_n\}$, 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. Thus, for each option d, $X(d)$ is random. Our main interest is on large option sets, as n approaches infinity, in which case the set of feasible options becomes the non-negative half line $[0, \infty)$, and the outcome function $X(\cdot)$ becomes the realization of a Brownian motion with drift μ and scale σ.

We denote the state of the world by $\theta = (\theta_1, \ldots, \theta_n)$ and the set of possible states of the world by $\Theta \equiv \mathbf{R}^n$. 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 Section 6.1 for details.

The sender's utility is strictly monotonic and declining in the choice of option; an example is the linear form $u_{\mathrm{S}}(d) = -d$. 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 Section 6.2.

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 the 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 $\sigma^2/(2\mu)$, which also reflects the complexity of the underlying decision problem. As $\sigma^2/(2\mu) \to 0$ the receiver can infer much about the environment from his knowledge of the default option, whereas as $\sigma^2/(2\mu) \to \infty$ such inference is increasingly useless.

To ensure that communication is non-trivial, we impose the following restriction on the receiver's preference relative to the complexity of the environment:[8]

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

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.[9]

Timing: First, nature draws the outcome function—the state of the world—and the sender learns the realization. Second, the sender sends her message. 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: The solution concept is perfect Bayesian equilibrium. A strategy for the sender is a mapping M from the set of all possible states Θ to the set of all possible messages $\mathcal{M}$. A strategy for the receiver is a mapping D from the set of all possible messages $\mathcal{M}$ to the set of all possible options $\mathcal{D}$. The receiver's beliefs are described by a belief mapping B that assigns belief $B(m)$—a probability distribution over states, $B(m) \in \Delta(\Theta)$—to every possible message $m \in \mathcal{M}$. Strategies M and D and belief mapping B form a perfect Bayesian equilibrium if *(i.)* the sender's strategy M maximizes her utility given D, *(ii.)* the receiver's strategy D maximizes his expected utility given B, and *(iii.)* given M, on path, the receiver's beliefs satisfy Bayes' rule whenever

[8]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 3.2).

[9]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 B.2.

possible and, off path, they are consistent with any hard information that has been revealed. We refer to a tuple (M, D, B) as a *strategy profile*. 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. An interaction contains a *recommendation* if an option that is revealed is chosen by the receiver. If it contains only a recommendation then the interaction is *conative* (by convention, we include in this definition the recommendation and choice of the default option). Advice other than a recommendation 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*.[10] 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)* $D'(M'(\theta))$.

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

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

as n goes to infinity.[11] 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.[12] Moreover, for any given equilibrium interaction, there are many off-path beliefs that can support it as an equilibrium. We refer to them 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.

Suspicious beliefs ensures the deviation may be unprofitable. To be 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.

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

[11] So, 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.

[12] 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 $M(\Theta)$.

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.[13] In Appendix B.1, 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 as a single equilibrium the class of equilibria that differ only in off-path suspicious beliefs and decisions.[14]

In some cases, particularly when the expert reveals less information than expected, the suspicious beliefs necessary must be skeptical and induce the receiver to choose the right-most option, d_n. The existence of a worst option for all sender types that can send a message is important for the existence of equilibrium, as has been known generally since the work of Seidmann and Winter (1997). Because our type space is intertwined with the space of options, we bound both with finite n. This choice is technical rather than substantive. Our primary interest is in equilibria as n grows large and the option space approaches the positive real line, which brings us close to the standard models in the literature.

3 Preliminaries

3.1 On-Path Beliefs and Strategic Communication

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.

[13]In Section 5.5 we offer an equilibrium refinement to select equilibria that do not require any suspicion when the sender "over-communicates."

[14]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.

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

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

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

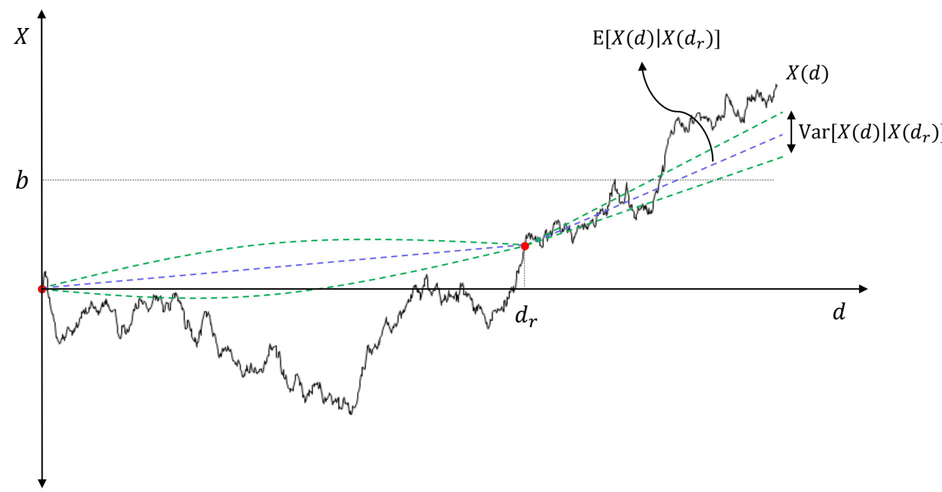

Figure 1: Neutral beliefs with knowledge of option d_r and option d_0.

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

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

and variance

$$\text{Var}[X(d) \mid X(d_l), X(d_r)] = \frac{|d - d_l| \cdot |d - d_r|}{d_r - d_l}\sigma^2. \tag{4}$$

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 1.

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. The credibility of a message derives from which sender types send it. However, the more effective a message is at dissuasion, the more sender types will want to send it, undermining the message itself. For dissuasion to be credible, therefore, it must 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.[15] 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.[16]

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*

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

[16]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.[17]

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.

3.2 The Necessity of Equilibrium Advice

Before moving on to the equilibria that do exist, we begin with one that doesn't exist. In every equilibrium the sender must communicate information as revealing nothing is not an equilibrium. This result is relatively straightforward, although the ideas that underlie it will prove generally useful throughout our analysis, and we develop it explicitly here.

The strategy of "no-advice" is non-deceptive—the revelation of any information is off-path—and, by Lemma 2, the receiver's beliefs are neutral. The receiver then faces a choice. He can accept the default option d_0 with a known but unappealing outcome, or he can venture off on his own, experimenting with a risky option that offers the prospect of a more appealing outcome. This generates a risk-return trade-off that depends on the complexity of the underlying environment. Larger options improve the expected outcome at rate μ but the variance increases linearly at rate σ^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. Recalling the default outcome is set to 0, we have:

$$\mathbb{E}[u_{\mathrm{R}}(X(d))] = -(b - \mu d)^2 - \sigma^2 d. \tag{5}$$

[17] All omitted proofs can be found in the appendices.

In the limit as $n \to \infty$, a simple optimization then gives the optimal choice to be d^{na} where:

$$\mathbb{E}[X(d^{\text{na}})] = b - \frac{\sigma^2}{2\mu} \qquad \text{and} \qquad \text{Var}[X(d^{\text{na}})] = \left(b - \frac{\sigma^2}{2\mu}\right)\frac{\sigma^2}{\mu}. \tag{6}$$

The receiver does not choose the option with expected outcome of b. Instead, he stops short at $b - \sigma^2/(2\mu)$. Beyond that point the improved return is not worth the additional risk. Figure 2 depicts this choice.

The maintained assumption $b > \sigma^2/(2\mu)$ ensures that for n large enough this choice is preferred over the security of the default option. For smaller n, integer problems may leave d_0 as the receiver's optimal when b is just beyond this threshold. To ensure more generally that the sender experiments in the absence of advice requires:

$$b > \frac{\sigma^2}{2\mu} + \frac{\mu}{2\sqrt{n}}, \tag{7}$$

with the final term vanishing as n grows large.

That the distance $\sigma^2/(2\mu)$ does not depend on the value of b, or, more accurately, the distance between the receiver's ideal and the default outcome, is a particular property of quadratic utility. What is not special is that the size of this gap determines the option chosen, how much risk it involves and, thus, the receiver's utility from his optimal experiment. The bigger the gap between the receiver's ideal and the default outcome, the bolder must the receiver experiment to find a desirable option, and the riskier that option is, lowering his utility. A better default outcome is good for the receiver, therefore, because even though it is abandoned, it means that he is more confident in his choice.

With no on-path deviations, and beliefs for off-path deviations constrained only by the hard information, intuition suggests beliefs can be found to render any deviation unprofitable (indeed, off-path beliefs need only be suspicious enough to induce a choice at or to the right of d^{na}—meaning an option $d \geq d^{\text{na}}$). That all is true, but it matters only to the extent that beliefs need to be formed at all by the receiver. A deviation of special importance in our setting is for the sender to reveal all the options, thus rendering off-path beliefs moot. For some sender types this deviation is profitable and, thus, no-advice is not an equilibrium strategy.

Lemma 3 *No-advice is not an equilibrium if* $b > \frac{\sigma^2}{2\mu} + \frac{\mu}{2\sqrt{n}}$*: In all equilibria, there is a sender type who discloses the outcome of at least one non-default*

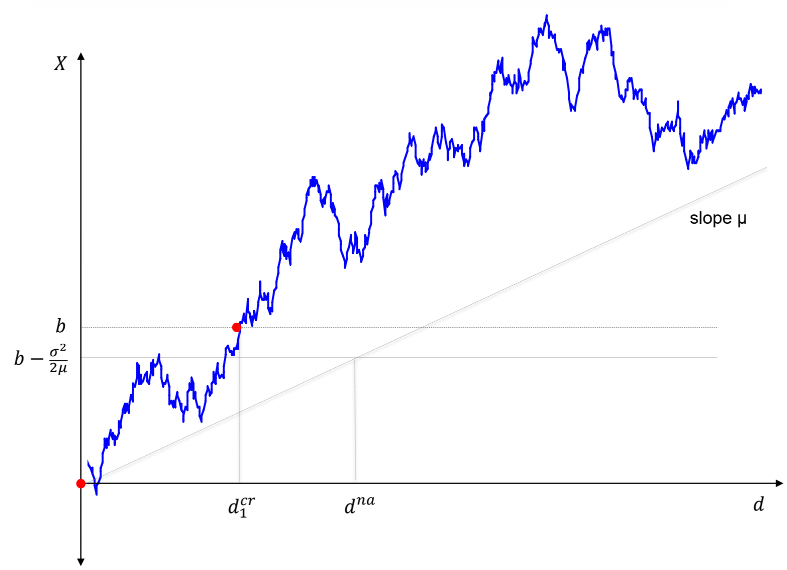

Figure 2: Receiver's optimal choice without advice d^{na}.

option.

An example of a profitable deviation is depicted in Figure 2. This outcome path crosses b only once at an option to the left of d^{na}. The receiver cannot ignore hard information and must choose this option. Thus, the deviation is profitable.[18] This depends, of course, on d^{na} not equalling d_0, which, for finite n, necessitates the restriction on b in the lemma.

The failure of no-advice as an equilibrium reflects the classic unravelling intuition of hard information games. With communication that is non-prescriptive, the receiver chooses an option that is an expectation of what the true state-of-the-world must be. That expectation is better for some sender types and worse for others. This incentivizes the latter group to deviate and reveal the whole truth to the receiver. In standard models this unraveling leads inexorably to full revelation. This is not the case here and unraveling need not be complete.[19]

[18]When interactions are prescriptive, the stronger result holds that, if n is large enough, in any equilibrium, the sender always reveals the outcome of at least one non-default option. That is to say, it is not only that the sender cannot always say nothing, she can never say nothing (see Appendix B.4).

[19]The failure of no-advice as an equilibrium portends a more general difficulty in sustaining non-prescriptive equilibria. Indeed, it can be shown that, for any n, and any equilibrium, the probability that the sender communicates information is bounded away from zero, so that it is not possible to generate a sequence of equilibria in which the probability of no-advice

The option d^{na} provides a key benchmark in our model as it also represents the choice of the receiver in the absence of an expert. The influence of the sender, therefore, is measured by the degree that her advice moves the chosen option to the left and in her favor.

4 Receiver Optimal Equilibria

We begin by analyzing receiver-optimal equilibria. Exploring this set of equilibria illustrates how the amount of information supplied in equilibrium impacts the sender's power, all while holding receiver utility (approximately) constant.

4.1 Full Revelation is an Equilibrium

If the sender fully reveals her information to the receiver, the receiver becomes an expert, and the sender relinquishes her ability to sway the decision in her favor. Despite this, full revelation is an equilibrium strategy.[20]

Proposition 1 *An equilibrium that is fully revealing exists. In this equilibrium, the receiver chooses* $d^{\text{ro}} \in \arg\min_d |b - X(d)|$ *and, as* $n \to \infty$*,* $X(d_n^{\text{ro}})$ *converges to* b*.*

Full revelation supports an equilibrium because the sender is constrained by the strategy itself. Because full revelation is non-deceptive, the sender cannot deviate without being detected. This leaves her with only off-path deviations and, if the receiver is sufficiently suspicious, these are unprofitable. Full revelation ties the hands of the sender, and leaves her in a position in which she is compelled to reveal more information than she would like. As most of this information is referential, referential advice, in this case, acts as a weight around the sender's neck.

Full revelation is clearly optimal for the receiver. He obtains the best possible outcome in every state and, as the number of options grows large, the outcome converges to his ideal outcome b. The choice itself, however, need not converge. As n grows large, the outcome path can cross b at many points and d^{ro} will get arbitrarily close to one of these crossing points, although which one will vary along the sequence depending on the realization of outcomes. To illustrate, Figure 3 depicts a path with three crossing points. We denote these crossing points by $d^{\text{cr}} = \{d_1^{\text{cr}}, \ldots, d_k^{\text{cr}}\}$, where $k = 3$ in the figure.

goes to one as n grows large. We take up the broader class of non-prescriptive equilibria in

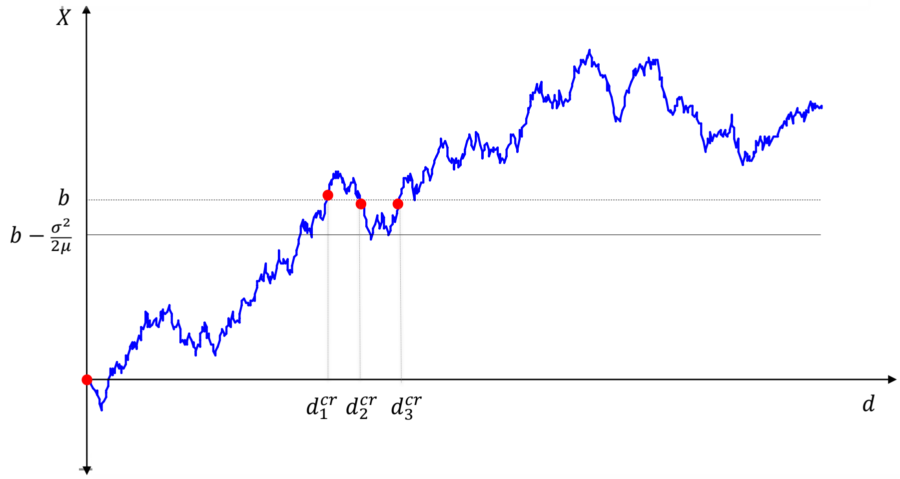

Figure 3: Full Revelation is an Equilibrium

4.2 Less than Full Revelation

In the full revelation equilibrium the receiver discards most of the information he receives, using only the information about his best outcome. This raises the question of whether the additional referential information is necessary to support equilibrium. Such a question is moot in standard models as receiver-optimality necessarily demands full revelation. In our richer setting it is possible for the receiver to be partially informed yet still have enough information to be satiated.

It is easy to see that full revelation is not necessary to ensure receiver optimality in equilibrium. Consider a prescriptive interaction in which the sender makes a recommendation d^{ro} and reveals all options to the right of d^{ro}, and the receiver chooses d^{ro}. In this case, the receiver cannot see directly that d^{ro} is his best option as he doesn't have all information. He can infer, however, that if a better option exists it must be to the left of d^{ro}, and in that case, the receiver would have incentive to reveal that option as it would move the decision to the left. This partially revealing strategy is an equilibrium strategy, therefore, and the corresponding equilibrium is equivalent to a fully revealing equilibrium.

This strategy can be amended, in turn, by having the sender reveal some

Section 5.5.

[20]See Hagenbach et al. (2014) for general conditions under which fully revealing equilibria exist in games with disclosure.

but not all of the information to the left of d^{ro}. In contrast to canonical models, there are many receiver optimal equilibria that differ in the amount of advice provided. Nevertheless, Proposition 2 establishes a lower bound on the amount of information. Although information to the left of d^{ro} is optional in equilibrium, the information to the right is not, and in every receiver-optimal equilibrium the sender must reveal the outcome of d^{ro} and all options to the right. For receiver-optimality, therefore, it is necessary that advice be referential.

Proposition 2 *In all receiver-optimal equilibria, with probability one, the sender reveals all the options to the right of the equilibrium decision.*

The only receiver-optimal equilibria that are non-deceptive is when the state is fully revealed with probability one. All other receiver-optimal equilibria are deceptive. For instance, the sender can deviate to the option to the right of d^{ro} that is best for the receiver and reveal all options to the right of that without the receiver detecting this as a deviation. The receiver's beliefs in these other equilibria are, therefore, non-neutral by Lemma 2. However, the inference the receiver draws from the indirect informational spillover is that the outcomes of these options are further from b than is the outcome of d^{ro} and so they are unappealing to him and, thus, the deviation is not profitable for the sender.

It is instructive to understand why so much information is required to sustain receiver-optimal equilibria. To see this, suppose the information supplied was reduced to a single option such that advice were conative. To be receiver-optimal, this option must be d^{ro} and the receiver must implement it in equilibrium. If the receiver were to rubber-stamp the sender's recommendation, however, there would be the temptation for the sender to make a different recommendation to the left that is more favorable to her. As the receiver does not know how good is her best option, this deviation is not detectable. The receiver would then implement the new recommendation and the deviation would be profitable.

This conative interaction fails as an equilibrium because the receiver cannot credibly reveal only d^{ro}. The ability to deceive the receiver—to deviate without detection—gives the sender too much freedom for an equilibrium to be supported. The failure of this equilibrium exposes why referential advice matters. The same recommendation can be supported in equilibrium if more information is provided. By providing referential advice, the receiver can be sure he is learning d^{ro} and choose it with confidence. Referential advice matters, therefore, because it provides credibility. By tying the hands of the sender, referential advice gives her credibility, and that credibility can support outcomes that are not obtainable with a recommendation alone.

Given the clear failure of conative advice as a receiver-optimal equilibrium, it is interesting that an equilibrium exists that in the limit is conative and converges on receiver-optimal. To see how, observe that there were two problems with implementing d^{ro} through conative advice. The first is that the best option is a relative standard, and the second is that non-detectable deviations exist to the left of the recommendation called for by the strategy. The following strategy rectifies both of these difficulties.

Definition 5 *The sender follows a* first-point conative strategy *when, for some* $\Delta > 0$*, 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.*

The first-point conative strategy replaces the relative standard of best alternative with an absolute standard—the outcome must be in a band around the receiver's ideal. In this way, the receiver knows whether the sender is telling him what she is supposed to be telling him. Such a standard does not rule out undetectable deviations, however, and the second key part of the strategy is that the sender reveals the *first* point that meets the absolute standard. It is deviations to the left that are profitable (given the sender's preferences) and this requirement rules them out.[21]

A first-point conative equilibrium is clearly not receiver-optimal. If an outcome falls within the band, the receiver will be getting a good outcome but potentially not the best. The receiver's loss is, however, decreasing in the narrowness of the band. As the band collapses around b (i.e., $\Delta \to 0$), the sender either reveals a single option with outcome arbitrarily close to b or she reveals all the options.

Proposition 3 shows that first-point conative equilibria exist and constructs a sequence of them such that $\Delta \to 0$ as $n \to \infty$ and the outcome approaches the receiver's ideal. Moreover, it establishes that this sequence can be constructed in such a way that Δ approaches 0 sufficiently slowly that the probability an outcome falls within the band is increasing and approaches one in the limit. Thus, the equilibrium is, in the limit, simultaneously conative and receiver optimal.

[21] It is important that the band is symmetric around b for the case of quadratic receiver utility (or any utility function symmetric around the receiver's ideal). Were it not, then with finite n it may be that the best option for the receiver, d^{ro}, is to the left of the first option to fall in the band, such that the deviation to full revelation is profitable. In Appendix B.3 we extend the notion of first-point conative strategy to the case of general—and possibly asymmetric—utility functions.

Proposition 3 *There exists a sequence of first-point conative equilibria such that interaction is conative in the limit and outcomes converge to the receiver's ideal.*

The first-point conative strategy is deceptive and the receiver's beliefs are not neutral. However, this is the case only to the left of the recommendation. To the right beliefs are neutral as the strategy is constructed in such a way that it conditions only on information to the left of the recommendation. Although the recommendation is almost surely not b, and the receiver cannot be sure he is seeing the global optimum, he is nevertheless willing to accept the recommendation as it is "good enough," and further experimentation is not worth the risk-reward trade-off.[22] We will develop the logic of conative equilibria more fully in Section 5.

Although this equilibrium is receiver-optimal in the limit, the sender strictly prefers it over full revelation. This is because this strategy converges on a particular decision in the limit, whereas full revelation doesn't. The limiting behavior is evident in the Brownian path depicted earlier in Figure 3. The conative equilibrium decision converges almost surely on the first point to cross b, d_1^{cr}, whereas full revelation does not, instead delivering some mixture of all the crossing points. This implies that full revelation and, indeed, all of the exactly receiver optimal equilibria in Proposition 2, are Pareto dominated in the limit by the conative equilibrium.[23]

The difference between the conative and referential equilibria of this section is one of expectations. In the referential equilibria, the receiver expects to receive a lot of information, and if he doesn't, and that deviation is detectable, he forms suspicious beliefs. The expectation of referential advice constrains the sender. The more referential information that is expected in equilibrium, the fewer on-path deviations are available to the sender, and this increases the set of outcomes supportable in equilibrium. By providing less information, even down to a single piece of information, the sender reduces this expectation. This increases the set of possible on-path deviations and from this freedom the sender can avoid revealing more information than she wants.

It may be tempting to extract from this the lesson that referential advice is necessarily bad for the sender, that it constrains her in leveraging her expertise

[22]This establishes the possibility of receiver optimality in the limit and is by no means the only equilibrium that does so, even among interactions that are conative.

[23]If instead we employed an alternative stochastic process with a positive probability of ties—such as a binary random walk—the receiver's optimal action along the sequence would be non-unique. In such a setting, multiple fully-revealing equilibria exist differing in how the receiver responds to his indifference, although the best of these equilibria would be no better than the conative equilibrium and the rest strictly dominated.

and that she should reveal her information as sparingly as possible. That intuition is wrong. We will show that whilst it will not be possible for the sender to avoid the expectations and constraints of referential advice, she can be better off with the constraints in place as, when constructed in just the right way, the constraints can be exploited to her advantage.

4.3 Refining Equilibrium

A long-standing tradition in signaling games (of which this is one) is to ask: Which beliefs are reasonable? What is meant by reasonable is subjective, and many refinements have been offered to sharpen equilibrium prediction. Our rich-information setting gives rise to its own questions of reasonableness, and we do not hope to provide a thorough treatment of the refinement question. Instead, we offer a simple, intuitive refinement that builds on the hard information feature of the model and that accords with classic ideas from the refinement literature.

One recurring theme in the refinements literature is the idea that beliefs are inherently uncertain. This manifests in the idea that equilibria should be robust to perturbations. This idea has particular force in our setting as some knowledge held by the receiver is hard information whereas other knowledge is only inferred. Specifically, it is possible in equilibrium for the receiver to believe with such confidence that an unrevealed option produces a certain outcome that he is indifferent between that option and another option for which the very same outcome has been revealed.

A minimal way to refine beliefs is to distinguish these possibilities, privileging known knowledge, so to speak, over inferred knowledge. We refer to this as the Bird-In-The-Hand (BITH) refinement and formalize it as follows.

Definition 6 *An equilibrium (M, D, B) satisfies the* ε-BITH refinement *if, for every off-path message m where $D(m)$ is not disclosed in m, and every option d disclosed in m,*

$$\mathbb{E}^{B(m)}[u_{\mathrm{R}}(X(D(m))] \geq \mathbb{E}^{B(m)}[u_{\mathrm{R}}(X(d))] + \varepsilon.$$

An equilibrium (M, D, B) satisfies the BITH refinement *if it satisfies the ε-BITH refinement for some $\varepsilon > 0$.*

This says that for an off-path message, m, if the receiver chooses an option $D(m)$ that is not revealed, it must be that the receiver believes the outcome of this option will deliver at least ε more utility than does the best option that is revealed. He only chooses an unrevealed option, therefore, if he believes it to strictly dominate the revealed options.

Within the class of receiver-optimal equilibria, the BITH refinement has a dramatic impact. Full revelation fails the refinement, as does any prescriptive equilibrium that is exactly receiver optimal. In contrast, every first-point conative strategy supports an equilibrium when off-path beliefs are chosen appropriately.[24] Moreover, any sequence of equilibria that satisfy the refinement and are receiver-optimal in the limit must approach the same decision as the first-point conative equilibrium. In the limit, therefore, the refinement pins down the equilibrium decision uniquely at the first crossing point of b.

Proposition 4 *No receiver-optimal prescriptive equilibrium satisfies the BITH refinement. However, there exists a sequence of first-point conative equilibria that satisfy the BITH refinement, that are conative in the limit, and in which the outcomes converge to b.*

The effect of BITH on the set of receiver-optimal equilibria is striking yet it follows from a clean intuition. Upon seeing with his own eyes an outcome arbitrarily close to his ideal outcome, the receiver asks himself: Why would I choose anything else? In particular, why would I engage in a complicated inference over unrevealed options, let alone have such confidence that I would choose one of those options?

The full revelation equilibrium relies on just this logic. Even if the sender deviates and reveals a single option with outcome arbitrarily close to b, the receiver nevertheless forms suspicious beliefs, beliefs that in this case must be so suspicious that he decides with almost certain confidence that the right-most option in the space, d_n, produces an outcome even closer to b. The BITH refinement suggests that such a high confidence is unlikely, and that the receiver would instead simply accept the recommendation, thereby breaking down the full revelation equilibrium.[25]

The same logic breaks down all other receiver-optimal equilibria. It does not break down the first-point conative equilibria though. These equilibria survive as there is never an option better for the sender at which he can make a take-it-or-leave-it offer to the receiver that will be accepted in this way. Full revelation is often taken as the benchmark equilibrium in models of hard

[24]The refinement rules out some off-path suspicious beliefs.

[25]Proposition 4 includes the caveat that equilibria be prescriptive. We can show that a receiver-optimal equilibrium cannot be non-prescriptive. The remaining possibility is a hybrid equilibrium in which the interaction is prescriptive for some states and non-prescriptive for others. Such an equilibrium must be deceptive and the additional requirements are highly complex, although we can't rule them out.

information. The BITH refinement, as well as Pareto efficiency, point away from this perspective and instead toward less information transmission as a more appropriate equilibrium benchmark.

5 Expert Power: Can the Sender do Better?

In the equilibria of the previous section the receiver extracted all of the surplus from expertise. Not only did he benefit from the removal of uncertainty over the outcome, but the outcome itself shifted to his ideal point. 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.

5.1 Conative Advice

In the conative equilibrium of the previous section, the sender was able to reveal an option that wasn't the receiver's ideal but that he accepted as good enough.[26] This raises the question of what qualifies as "good enough"? Theorem 1 characterizes exactly what is possible in equilibrium. It shows that the sender can do better and extract some of the value of her expertise, but that there is a tight bound on what is possible.

Establishing this result 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 $\sigma^2/(2\mu)$ from the receiver's ideal outcome.

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

Many conative equilibria can be supported because the strategy creates expectations for the receiver. If the receiver expects the revealed option to

[26] To simplify the language we abuse it slightly in referring to these as "conative equilibria" as it is only as n grows large that the first-point conative equilibria formally become conative.

meet a certain absolute standard he will be suspicious if it does not or if the sender reveals additional information, and so the expectation constrains the sender. Most of these equilibria are not receiver optimal, although the receiver still does well. By listening to the sender's advice he not only benefits from the removal of risk, but in all but the boundary case the outcome shifts toward his ideal relative to what he would get in expectation without an expert, and in that boundary case the outcome is exactly the same.

The limitation of Theorem 1 is striking and leads to the question of why. Why can't the sender leverage her expertise for her own gain by recommending an option with a lower outcome that corresponds to a smaller option? There do exist more favorable recommendations that still give the receiver the same or more utility than he would get without an expert, yet these recommendations cannot be supported in equilibrium. The answer is the inevitability of informational spillovers. The sender may be able to minimize the amount of information she reveals to the receiver, but she cannot stop the receiver from using that information for his own advantage.

For a strategy to yield a conative interaction in the limit, it must necessarily be deceptive.[27] Yet, as we saw in the previous section, the construction of a first-point conative strategy is able to avoid indirect spillover to the right. Direct informational spillover is unavoidable, however, and this inevitability bounds the sender's ability to persuade. To see this, suppose the sender follows a first-point conative strategy with $\Delta \in (\sigma^2/(2\mu), b)$, such that the outcome is between the default outcome at 0 and $b - \sigma^2/(2\mu)$. Although an outcome at $b - \Delta$ delivers more utility to the receiver than the default option, and possibly more than he would get without an expert, it also increases the utility the receiver expects from other, unrevealed options. The sender's problem is that once the recommendation is made, the receiver's outside option is no longer no-advice but rather it is what he can gain by using the informational spillover from the advice itself. The persuasiveness of the advice is undermined by the advice itself.

This can be seen mathematically by recalling the receiver's optimal choice in the absence of advice (Equation (6)). As the receiver's beliefs are neutral, his decision problem differs only in that the recommendation rather than the default provides the anchor point for beliefs. This means that the receiver still wishes to experiment, even as the recommendation gets closer and closer to $b-\sigma^2/(2\mu)$. The certainty in the recommendation spills over to nearby alternatives, such

[27] Non-deceptive strategies do exist that yield conative interactions—such as in Example 2—yet such strategies cannot guarantee an outcome in $[0, 2b]$ and, thus, guarantee being accepted by the receiver over the default option.

that the more attractive is the recommendation, the more attractive are nearby alternatives as well. This represents a positive complementarity between advice and the spillover of information. This process only ends, and the receiver is only satiated, when the first-point conative strategy provides a recommendation with outcome $b - \sigma^2/(2\mu)$ or better, and this defines the boundary for what the receiver deems "good enough."

This logic applies broadly to all conative interactions if the sender is using a first-point style strategy for all states. The richness in the theorem emerges when she is not. If the sender uses a first-point strategy for some states but not for others, the receiver's beliefs need not be neutral, even for a conative interaction. This is possible as the recommendation in one state using a first-point strategy may also be sent for a different state as part of a deceptive strategy. The non-neutrality of beliefs holds for all equilibria along the sequence as n grows large, such that establishing the bounds in Theorem 1 for *all* sequences of equilibria that are conative in the limit is non-trivial. Nevertheless, we show for these sequences that the receiver's beliefs must 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.

The intuition we have described depends on the sender using a first-point strategy on an increasing fraction of states, although, as we note above, this strategy could take a generalized form of that in Definition 5. The final piece of intuition for Theorem 1 is why the sender's strategy must come from this class. 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' with outcome $X(d') < b - \sigma^2/(2\mu)$ 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 based on information to the right of the recommendation. This is optimal only if these sender types recommend an option to the left of d'. But if all types who could send this message instead of d' do so, beliefs at d' 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.

The conative equilibria are good for the receiver but not so good for the sender. The first-point conative equilibrium that is optimal for the sender is the one with the widest band as this both delivers the smallest option when a point falls within the band and it minimizes the probability that the sender is compelled to fully reveal. We refer to it as the *sender-optimal conative equilibrium*. This equilibrium is defined by letting $\Delta = \Delta^{\max}$ in Definition 5, where:[28]

$$\Delta^{\max} = \frac{\sigma^2}{2\mu} + \frac{\mu}{2\sqrt{n}}. \tag{8}$$

In the limit the revealed outcome converges to $b - \sigma^2/(2\mu)$. To be sure, this is better for the sender than outcome b, yet it is not necessarily better than the expected outcome without her presence. For some paths, conative advice will pull decisions to the left of d^{na} (the choice with no advice) and make the sender better off, whereas for other realizations the decision will be pushed to the right. The net effect on the sender of providing conative advice is particularly stark With quadratic receiver utility. In expectation, the decision is exactly the same in the sender-optimal conative equilibrium as it is with no expert, and for conative equilibria that converge to a different outcome the option chosen is strictly larger.[29]

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.[30]

[28]This is the same threshold, and serves the same purpose, as in Equation (7).

[29]With general receiver utility this comparison is not strict, yet the sender's weakness in influencing the decision with conative advice persists. See Section 6.1 for a discussion.

[30]It is here that the failure of no-advice as an equilibrium has bite. Even if the conative equilibrium leaves the sender strictly worse off, she does not have the ability to commit to not communicate.

5.2 Referential Advice: Interval Equilibrium

The supply of referential information does not by itself deliver leverage to the sender. Every first-point conative equilibrium is equivalent to many referential equilibria. For instance, it is an equilibrium for the sender to follow a first-point conative strategy amended to reveal all options up to and including the recommendation, and this equilibrium is equivalent to the first-point conative equilibrium itself.

The addition of referential advice in this example 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 7 *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*

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

with $d^r = d_n$ *if no such option exists.*

The interval strategy represents a stopping rule. 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 that peak and ultimately equalling the peak. When it reaches that point—what can be thought of as an upper threshold—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 4 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. 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

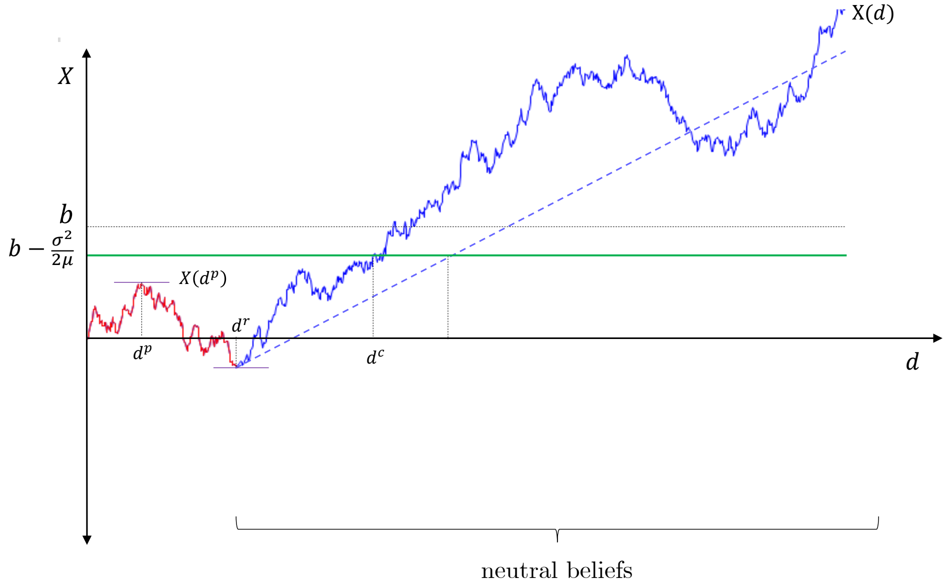

Figure 4: The interval strategy hitting the lower threshold

between information revealed 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, the more emboldened is the receiver to experiment.

The interval strategy turns this complementarity 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 further experimentation and accepting the recommendation (the right and the left-hand side of the expression in Definition 7, respectively). For the limiting case as n grows large, the analogue is given by:

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

where, recalling Equations 5 and 6, the first term on the right-hand side is the expected utility from a mean outcome at $b - \sigma^2/(2\mu)$ and the second term is the corresponding variance. This simplifies to:

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

The threshold increases in the peak and meets it at the upper threshold as at this point it is not profitable at all to experiment further to the right of d^r. This is exactly the same point as defines the sender-optimal conative equilibrium (Equation (8)) that in the limit converges to the now-familiar $b - \sigma^2/(2\mu)$, and marked in the figure by option d^c. If the upper threshold is reached, revelation stops and the recommendation is the right-most revealed option itself.

Proposition 5 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 equilibria. At best, 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 even the sender-optimal conative equilibrium.

Proposition 5 *The interval strategy supports an equilibrium. This interval equilibrium 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 always 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.

The referential advice allows the sender to paint a picture of the world that is unfavorable to the receiver, and the construction of the strategy is such that she can do so credibly. For a particular recommendation, the only sender types who can send the message are those with hard information that is consistent and for whom it is unfavorable for the receiver to experiment.

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.

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 $n \to \infty$.*

This implies that not just some sender types can find a peak and a downturn that deters experimentation, but that every sender type can. These recommendations may get close to the limit threshold of $b - \sigma^2/(2\mu)$, but one can always be found. This is possible because as the peak approaches $b - \sigma^2/(2\mu)$, the lower threshold converges on the peak and as the path nears this boundary the value of experimenting is itself small. As such, only a small downward step is needed to dissuade the receiver from experimenting.

That all types can dissuade in this way contrasts with existing intuitions. Typically if one set of types separate with an 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.

We can obtain further insight into the interval strategy numerically. Figure 5 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 $b - \sigma^2/(2\mu)$). 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 taking a very favorable decision. At the other extreme, the sender has little leverage because the outcome follows a

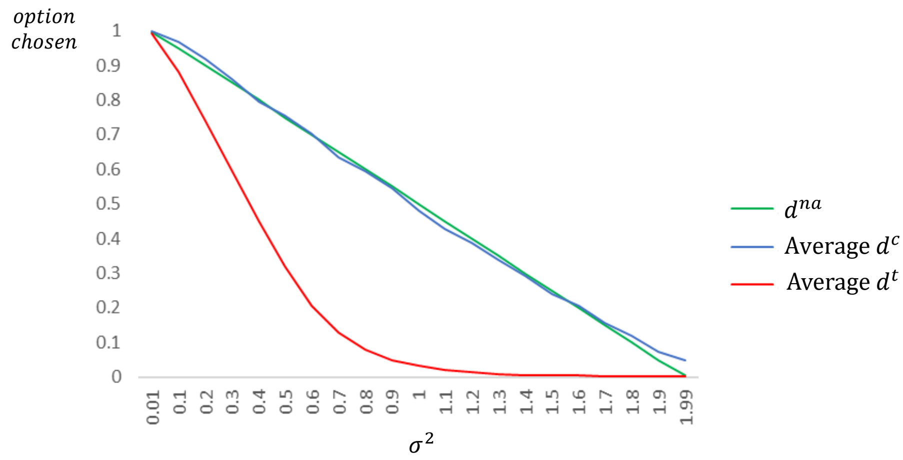

Figure 5: The receiver's decision as a function of σ^2 for no-advice (green), the sender-optimal conative equilibrium (blue) and the interval equilibrium (red), for parameter values $\mu = b = 1$, $n = 1000$, and sample size $10,000$ for every value of σ^2.

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 receiver can extract so much leverage. In accepting the recommendation, the receiver is accepting an outcome potentially far from his ideal, and in some cases he would get a better outcome by experimenting on his own. Using the same parameter values as in Figure 5, 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 her ideal. Going it alone, on the other hand, 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.

The effectiveness of the interval strategy raises two questions. Is the entire interval of options in the interval strategy necessary for equilibrium? And, can other referential equilibria be constructed? We turn to these two questions now.

5.2.1 How Much Referential Advice is Needed?

The key to the interval equilibrium is the peak and the end-point as these points persuade and dissuade the receiver, respectively. The additional information beyond these two points plays a role, although not all of it is necessary. 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 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 & 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 dismissible. 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. Although the outcome is different for the sender, this is the same reason why the receiver optimal equilibria of Proposition 2 require all information to the right of the optimum be revealed. 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 (such 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.

5.3 Other Referential Equilibria

The interval equilibrium is not the only way to persuade and dissuade the receiver with referential advice. Characterizing the full set of referential equilibria is a daunting task, and one that is beyond our capabilities at this time. The elegance of the interval strategy is that it achieves persuasion and dissuasion as a non-deceptive strategy and, thus, yields neutral beliefs. The space of deceptive strategies is unusually large in our setting, and with both direct and indirect informational spillovers, analyzing the beliefs to verify or dismiss a strategy as an equilibrium is challenging.

Nevertheless, some additional referential equilibria can be identified. The following deceptive sender strategy supports an equilibrium that is not equivalent to the interval equilibrium.

Example 3 *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.*[31] *For some $T^* < 0$, the sender:*

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

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

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

This strategy supports an equilibrium if T^* is negative enough such that, when the sender reveals q only, the best response of the receiver is to choose the default option. It differs from the interval strategy in that it seeks a low outcome at a particular option, q. The payoff is that should this occur, the receiver will choose the default option, d_0, which is the sender's optimal. The sender is credible despite not reporting the intermediate information because the option q is preset into the strategy.[32] Although this can achieve the sender's

[31] 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 $q\sqrt{n}$.

[32] Like the conative equilibria, this strategy demands an absolute rather than relative standard, and all deviations can only move the choice away from the default option.

most preferred option, it comes with a cost should the outcome of option q not dissuade sufficiently. If instead the outcome path rises quickly to cross $b - \Delta^{\max}$ to the left of q, the sender is compelled to reveal all the options and let the receiver implement his ideal option. If neither of these situations occurs, the sender proceeds with the sender-optimal conative equilibrium strategy.

The equilibrium that Example 3 generates neither dominates nor is dominated by the interval equilibrium. As n grows large, the probability of d_0 being chosen in the interval equilibrium approaches zero, whereas in Example 3 it remains strictly positive. At the same time, with a probability bounded away from zero, all options are revealed in the equilibrium of Example 3, making the sender strictly worse off than in the interval equilibrium.

The strategy in Example 3 is deceptive as the sender can deviate and use the first-point conative revelation strategy even when the outcome of option q meets the threshold. However, as with many equilibria in deceptive strategies, an equivalent equilibrium exists in non-deceptive strategies. In this case, an equivalent equilibrium that is non-deceptive is for the sender in case (iii) to reveal all options up to the conative threshold. Although this equivalence is common, it is not always possible. The following strategy provides an example.

Example 4 *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.*

This strategy supports an equilibrium if T^{**} is negative enough as then the receiver's best response is to choose the default option, d_0, when the sender reveals only one option; otherwise he chooses his ideal option.

As with first-point conative strategies, the sender's strategy in Example 4 requires that an absolute standard be met in terms of outcomes, and that the first point to meet this standard be revealed such that the only non-detectable deviations are unprofitable. However, unlike first-point conative strategies, this strategy does not produce an equivalent equilibrium if it is amended to a non-deceptive strategy in which all information is revealed up to the option that attains the threshold. If it were, the receiver would not choose d_0 should the mapping at first increase before it dips down. The difference between this and first-point conative strategies is that the revealed option is the recommendation in the conative strategy whereas here the revealed option is intended to dissuade. As the intermediate information must be revealed for non-detectable deviations to be precluded, the strategy in Example 3 cannot be amended to be non-deceptive and support an equivalent equilibrium.

The equilibria that are supported by Examples 3 and 4 lead to an increase in the variance in the receiver's choice relative to the interval equilibrium, pushing the choice for some states to the left and for others to the right. For a sender with concave preferences, this suggests the interval strategy is preferred, whereas with convex preferences this ordering may be reversed. We can turn to simulation to answer the question for specific functional forms and parameter values, but without guidance as to what those values should be, and without a full characterization of such equilibria, such an exercise is of unclear value.

5.4 Equilibrium Dominance

To understand the trade-offs across referential equilibria, we pursue a dominance result along the lines of the dominance of the interval equilibrium over all conative equilibria. The examples of the previous section neither dominate nor are dominated by the interval equilibrium. The examples are simple, however, and one may wonder whether a more elaborate extension could beat the interval equilibrium. In Example 3 for instance, the receiver's expectations reset for options to the right of $d_{q\sqrt{n}} = q$ and the sender could play the interval strategy in this range or repeat the same trick as with option q, at perhaps $2q$, and so on. As rich as these possibilities are, and despite our inability to characterize the set of equilibria fully, we can prove that any manipulations of information to the right of the recommendation involves a trade-off with the interval equilibrium of the type shown in the examples.

To focus on manipulations of information to the right of the recommendation, we define a *strongly prescriptive* interaction as one in which the sender reveals, at least, the options $d_0, d_1, \ldots, d^*$, where d^* is the receiver decision. An equilibrium is strongly prescriptive when interactions are strongly prescriptive for all states. Within this class of equilibria the sender is free to strategically provide information to the right of the recommendation whether non-deceptively or deceptively.

We prove that no equilibrium that is strongly prescriptive dominates the interval equilibrium. In fact, we show a stronger result. We show that if a strongly prescriptive equilibrium dominates the sender-preferred conative equilibrium, the interval equilibrium necessarily dominates it. Thus, any strongly prescriptive equilibrium that is better for the sender than is the interval equilibrium for a positive mass of states, must be worse for the sender for another positive mass of states and, in fact, must be worse on those states than is the sender's preferred conative equilibrium.

Theorem 3 *If Σ is an equilibrium whose interactions are strongly prescriptive with probability one and that weakly dominates the sender-optimal conative equilibrium, then the interval equilibrium weakly dominates Σ.*

The class of strategies that support strongly prescriptive interactions is broad. The interval strategy is in the class, as is full revelation. The class also includes deceptive strategies, such as the strategies in Examples 3 and 4.[33] Indeed, many strategies in which the interaction cannot be strongly prescriptive nevertheless can be amended so that it is and the equilibrium equivalent. An example is the first-point conative strategy amended to include all information to the left of the recommendation.

That the interval equilibrium is strongly prescriptive creates a connection between Theorem 2 and Theorem 3. Within the class of strongly prescriptive strategies, the two theorems establish that for a sender who wants an equilibrium that beats all conative equilibria for any strictly decreasing utility function, then the interval equilibrium is the best she can do.

The class of strongly prescriptive equilibria also has a tight connection to non-deceptive equilibria. In ruling out manipulation of information to the left of the recommendation, there is no room for deception. This is why the interval and many other non-deceptive strategies support strongly prescriptive interactions. This is not true for all non-deceptive strategies, however. For instance, if the sender always and only reveals the outcome of some option d_q, for $q \geq 2$, the interaction cannot be strongly prescriptive.

The strategy just described not only fails to support a strongly prescriptive equilibrium, it cannot support an equilibrium at all. This raises the question of whether non-deceptive equilibria can be found that are prescriptive but not strongly prescriptive. In Lemma 4 we answer this in the negative. We show that for a non-deceptive strategy to support an equilibrium, that equilibrium must be strongly prescriptive.

Lemma 4 *If a non-deceptive equilibrium is prescriptive, interactions are strongly prescriptive with probability one.*

The logic is familiar from the necessity of equilibrium advice in Lemma 3. A non-deceptive strategy that leaves a gap to the left of the recommendation leaves open the possibility that full revelation is a profitable deviation for some types. Because beliefs are neutral, there must be some sender type for whom the outcome path gets close to b within the gap and never gets as close to b again.

From the lemma it follows immediately as a corollary to Theorem 3 that the interval equilibrium is undominated by all non-deceptive equilibria. Any non-deceptive equilibrium that dominates the sender preferred conative equilibrium is, in turn, dominated by the interval equilibrium.

[33] These strategies call for gaps in some cases, although the gaps are to the right of the recommended option, which in both cases is d_0.

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. Throughout the 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 supply to the left of the recommendation is inter-dependent with information to the right. The strongly prescriptive 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 3 is easily extended to *all* prescriptive strategies. We cannot construct an interdependence such that a non-strongly-prescriptive equilibrium exists, let alone one that dominates the interval equilibrium, but we cannot, alas, rule it out.[34]

5.5 Non-Prescriptive Equilibria

We have until this point focused on prescriptive equilibria in which the sender conveys a recommendation. Non-prescriptive equilibria cannot be ruled out, although they must satisfy stringent properties that, seemingly, render them unappealing to the sender. The uncertainty in outcomes—as the receiver is choosing an unrevealed option—also implies that many non-prescriptive equilibria are Pareto inefficient, including all non-deceptive equilibria.

What makes the existence of non-prescriptive equilibria difficult is that the sender can always deviate and fully reveal. We saw earlier that full revelation undermines no-advice as an equilibrium and, in the same way, full revelation plagues all other non-prescriptive strategies. The issue is, once again, if a gap exists between the revealed options and the decision by the receiver as this creates space for this deviation.

One way to create non-prescriptive equilibria is by eliminating this gap. That is, by using a strategy that reveals up to the option next to the one that the receiver chooses. Constructing these equilibria are nonetheless involved and non-obvious. More importantly, they are unsatisfying as in the limit as

[34]This seemingly represents a technical rather than substantive difficulty. The purpose of information to the right of the recommendation is to dissuade, and Theorem 3 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.

n grows large, the option space grows dense and these equilibria effectively become prescriptive. We refer to this type of equilibria as *near-prescriptive.*

Definition 8 *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*.*

To see the possibilities for non-prescriptive strategies, it is helpful to return to the categories of deceptive and non-deceptive strategies. For non-deceptive strategies the receiver's beliefs are neutral and this implies that all outcomes are possible for unrevealed options. But if all outcomes are possible, and a gap exists between the revealed options and the receiver's intended choice, there must be a path that gets closest to b in the gap and for whom full revelation is profitable. Therefore, by the same logic as Lemma 4 above, a non-deceptive strategy can support a non-prescriptive equilibrium only if a gap doesn't exist and the strategy is near-prescriptive.

Lemma 5 *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.[35]

A further way to evaluate the reasonableness of non-prescriptive equilibria is to evaluate the reasonableness of the receiver's beliefs following off-path deviations. An interesting distinction that our setting creates is whether a deviation over or under-communicates relative to equilibrium. For instance, a deviation could include an equilibrium message plus additional information—what we refer to as *over-communication*—or it could include something less

[35]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 $d^* \leq d_{np}$, 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.

than an equilibrium message, possibly plus other information, which we call *under-communication.*

Over-communication is an interesting possibility as the sender shows the receiver what he would get in equilibrium and then adds more information. This raises the question of what the receiver would believe were the extra information to reveal a better outcome than he expects in equilibrium. Would he take the extra information at face value? That is, would he see the extra information as non-deceptive and incorporate the information only with direct informational spillovers? This would be a particularly tempting inference if the extra information was also beneficial for the sender and a Pareto improvement.[36]

One may construct an equilibrium refinement that follows this logic—what might be called an "over-communication" refinement.[37] If the sender over-communicates then the receiver is not suspicious at all, but rather takes the information at face value. It is easy to see that this refinement would eliminate all non-prescriptive equilibria that are Pareto inefficient.[38] In that case, mutually beneficial over-communication is possible. At the same time, the interval equilibrium and all first-point conative equilibria survive this refinement. The Pareto requirement 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.

This informal argument and Lemma 5 do not provide a definitive judgement against non-prescriptive equilibria. Indeed, the multiplicity of equilibria in our model goes hand-in-hand with the richness of expertise and the information structure. Nevertheless, we see these arguments, on top of the already difficult requirements for the existence of non-prescriptive equilibria, as suggestive of why we might more reasonably expect prescriptive equilibria to be those played in practice.

6 Robustness

The model we analyze is abstract and general. Nevertheless, restrictions have had to be imposed. In the appendices we consider extensions to allow for a richer

[36]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."

[37]The formalization of this refinement is straightforward, although in the interests of space, we do not include the details here.

[38]For instance, it would eliminate the potential equilibrium described in Footnote 35.

message space (Appendix B.2) and a partially informed sender (Appendix A). In this section we focus on two specific variants that are perhaps more aptly viewed as robustness rather than extensions.

6.1 Generalizing Receiver Utility

To simplify the presentation 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 are provided under general receiver utility and, to that end, we establish several intermediary results about this environment in Appendix B.3.

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 $b - \sigma^2/2\mu$, 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.[39]

Assumption 1. u_{R} is smooth to the fourth order and exponentially dominated.[40]

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 $(-\infty, b)$, $R' > 0$.

[39] 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.

[40] 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.

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.2), 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 $b > \sigma^2/(2\mu)$ that ensures the sender's problem is non-trivial.

The logic for conative equilibria is the same, although the condition describing the first-point conative equilibria is more involved (see Appendix B.3 for the extended definition). Regardless of the default outcome, however, there exists an outcome that the receiver is willing to accept risk-free 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 peak-and-trough 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 that conative threshold is reached.

The one result that does not extend to general receiver preferences is Corollary 1 that reveals a preccise 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. The appendix provides complete details for all of these results.

6.2 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. In this section, we describe these results informally for the case in which the number of options goes to infinity, leaving the formal statements to the appendix.

So suppose that the sender's preferences are strictly increasing in d, such as the linear form $u_S(d) = d$. Since the sender plays no role in the no-advice benchmark, this change does not affect the option the receiver would choose in the absence of any advice, which is still given by d^{na} that produces outcome $b - \sigma^2/(2\mu)$ in expectation.

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 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 6a.[41]

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 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. Figure 6b illustrates for the interval strategy. Rather than reveal from the left, the sender now reveals from the right, with all revealed

[41] The construction of this strategy now yields neutral beliefs to the left of the revealed option rather than to the right.

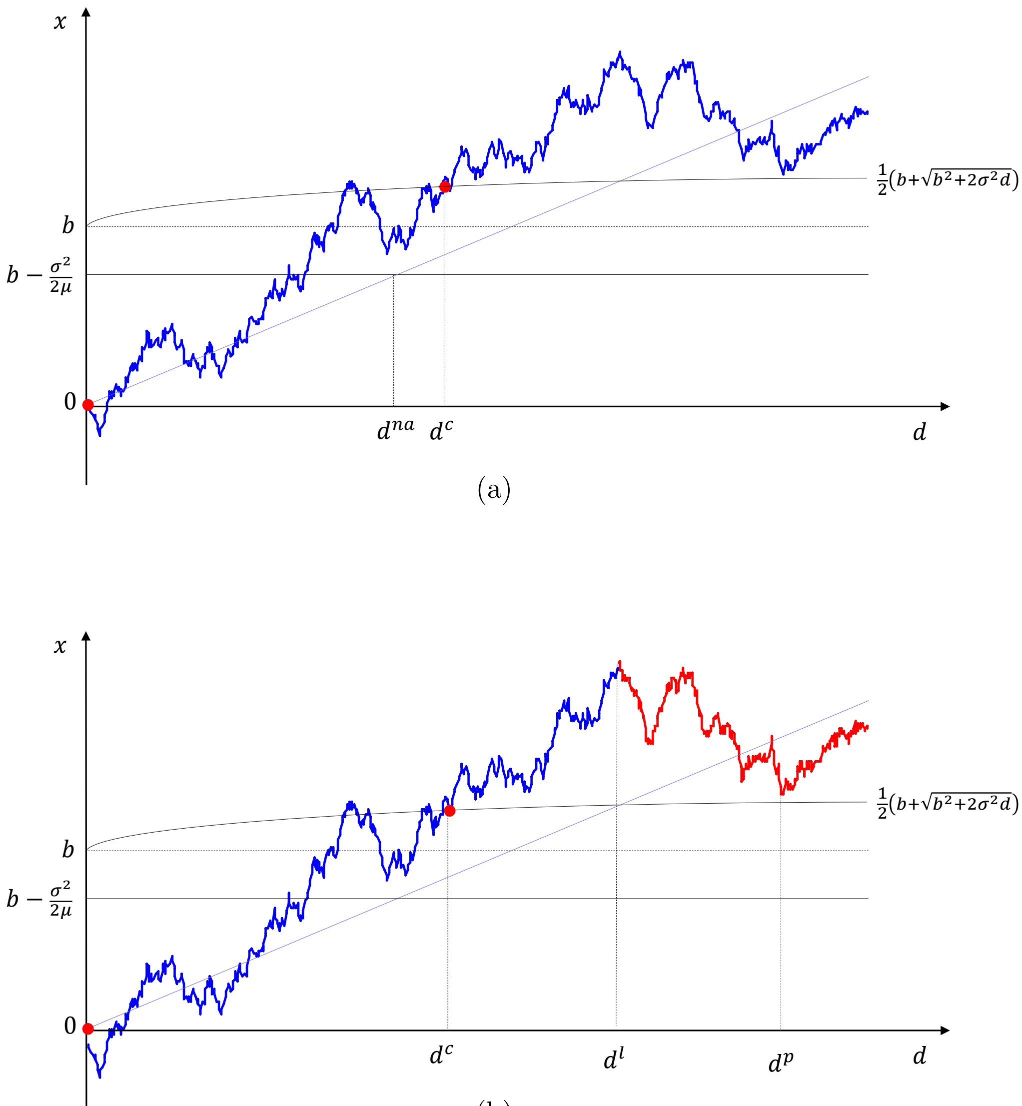

(a)

(b)

Figure 6: Strictly increasing sender preferences.

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 δ^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 needing 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.

7 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. CEO's 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.[42]

Beliefs play an important role in our analysis. The BITH and over-communication refinements are initial steps to explore which beliefs are natural and which are not. The standard refinement of never-weak-best-response (NWBR) can be applied here as well, although it eliminates none of the equilibria of interest and its ability to refine prediction is limited. Many more possibilities exist given the rich informational structure. An interesting direction suggested by the over-communication refinement is that of plain-spoken versus obtuse communication. The idea of that refinement is that an expert who over-communicates does so with good intent and, as such, the information should be taken at face value. In contrast, under-communication suggests deception and breeds mistrust. This logic reinforces the connection between neutral beliefs, non-deceptive strategies, and the absence of indirect informational spillovers. It argues for equilibria in non-deceptive strategies over those that rely on deception. It also points away from non-prescriptive equilibria in which the expert reveals information but strategically refrains from making a recommendation. Reinforcing this argument is that cognitive load of required for non-deceptive strategies and neutral beliefs is lower as the play is, in a sense, strategically simple. There are no subtle inferences that need to be drawn, and beliefs are simple extrapolation and interpolation from the hard information. Exploring these ideas in theory and particularly in the lab promise to yield more insight into the psychology and strategy of communication.

Our focus has been on equilibria when the number of options is large, although our results are characterized for finite n. Thus, a large set of options is not required for referential advice to be relevant, although there must be more than two as, with only two options, there is no space for referential advice. Our results show simply that the bigger the set, the more likely the opportunity will arise for referential advice to be able to influence the decisionmaker, and for referential to dominate conative equilibria.[43] A similar conclusion emerges

[42]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.

[43]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

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 the appendix 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 can also be tailored to fit more tightly to empirical applications. A remarkable feature of health care in the United States is how much variation exists in the incidence and quality of health care across the country, and considerable effort has been put into understanding the reasons why (Chandra et al. 2012). In the working paper version of the paper we develop an extension in which a patient—the receiver—cares both about outcomes (health) and the option chosen (cost and inconvenience). The amended interval equilibrium that results suggest a novel interpretation for the complicated interdependency between insurance coverage and medical care. In addition to the demand and supply side factors documented in the literature, the equilibrium suggests that differences in insurance produce different degrees of persuadability of patients, and that this may explain some of the variation in the data. In the equilibrium the patient who cares less about the option chosen—a patient with good insurance and lower copays—is more susceptible to persuasion and, as a result, will receive worse outcomes. Examining this connection more closely in the data, and separating it from other demand side effects, offers an interesting and promising direction to investigate.

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

task more difficult by ruling out that possible channel of influence.

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 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 extends the main model to the case of imperfectly informed senders. Appendix B includes the results omitted from the main text and several preliminary results used throughout the proofs. Appendix C describes and proves the formal results of Section 6.2. Finally, appendix D contains the proofs omitted from the main text.

A Sender is Imperfectly Informed

In this appendix, we 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 of 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:

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

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

$$\mathbb{E}[X(d) \mid Y(d')] = \mathbb{E}[X(d) \mid Y(d''), \forall d'' \leq d'] = Z(d') + (d-d')\mu, \tag{A.1}$$

and

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

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

$$-(Z(d)-b)^2 - \gamma\sigma^2 d. \tag{A.3}$$

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 $b - \sigma^2/(2\mu)$ (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.1 below.

Proposition A.1 *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 contraction that there exists a sequence of equilibria that are conative in the limit, which we write $\Sigma_1, \Sigma_2, \ldots$.

Let $\underline{z} = b - \sigma^2/(2\mu)$. 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 5.1 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 (A.3), 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.2 *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 $[b - \sigma^2/(2\mu), b + \sigma^2/(2\mu)]$. If the sender communicates the signal of only one option d^*, applying the receiver's beliefs given by (A.1) and (A.2) and using the same logic as in the second part of Theorem 1, the receiver is never strictly better off choosing a 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 $-\sigma^4/(4\mu^2)$. 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 A.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 5.1.

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

$$\max_{d \leq d^r} \mathbb{E}[u_{\mathrm{R}}(X(d)) \mid Y(d)] \geq \max_{d > d^r} \mathbb{E}[u_{\mathrm{R}}(X(d)) \mid Y(d^r)], \tag{A.4}$$

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.3 *The sender strategy just described is a prescriptive equilibrium strategy.*

Proof. Note that the left-hand side of Equation (A.4) 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.4) 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 A.1 ensures equilibrium existence. ■

Note that as ε vanishes, the above sender strategy converges to the interval strategy of Section 5.2. Thus, as ε vanishes and n grows large, the equilibrium that results strictly dominates the all-or-one equilibrium of Proposition A.2.

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 5.2).

B Auxiliary Results

This appendix includes the results omitted from the main text and preliminary results used throughout the proofs.

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 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 $m \in M(\Theta)$,

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

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 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 A.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 A.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 $\theta \in \Gamma(m)$,

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

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 A.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 A.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 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 $\mathcal{M}$ 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.4 *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 A.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.3 On General Receiver Utility

In this section, we provide auxiliary results to account for the case of general receiver utility. We assume the receiver utility meets Assumptions (1)–(5) of Section 6.1.

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), $\underline{x} \in (0, b)$, and the arguments made in Section 3.2 together with Lemma A.3 below imply that the sender cannot get her first best in equilibrium if n is large enough. Observe that, for the case of quadratic receiver utility in the main text, $\underline{x} = b - \sigma^2/(2\mu)$ and $\bar{x} = b + \sigma^2/(2\mu)$.

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 extend below the definition of first-point conative sender strategy.

Definition 9 *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.

In Lemmas A.2 and A.3 below, Z is an independent random variable that follows the standard normal distribution.

Lemma A.2 *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 the 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),

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

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

Lemma A.3 *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.2, 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.2. 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. ■

B.4 On No-Advice Equilibrium Messages

The goal of this section is to 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 A.4 *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. ■

C Formal Results of Section 6.2

Here we consider the case in which the sender prefers larger decisions as opposed to smaller decisions, letting her utility function be $u_{\mathrm{S}}(d) = +d$. The receiver's utility function remains unchanged.

Suppose the receiver knows $X(d^*) = x^*$, 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 5.1 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 $d^* < 4b^2/\sigma^2$, 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 $d^* > 4b^2/\sigma^2$, so that $\beta(d) = b$, then for $x^* < b + \beta(d^*) = 2b$, the receiver's optimal decision is d^*. For $x^* \in [2b, d^*\sigma^2/(2b)]$, the optimal decision is 0. Then, as x^* increases above $d^*\sigma^2/(2b)$, the optimal decision gradually increases towards d^*, and then 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.5 *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 a 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 $\operatorname{Var}[X(d) \mid X(d^*){=}x^*] = \sigma^2(d^* - d)d/d^*$. 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

$$-\left(b - \mathbb{E}[X(d) \mid X(d^*){=}x^*]\right)^2 - \operatorname{Var}[X(d) \mid X(d^*){=}x^*]$$
$$= -\left(b - \frac{dx^*}{d^*}\right)^2 - \sigma^2 \frac{(d^* - d)d}{d^*}$$

  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

$$\frac{d^* - d}{(d^*)^2}\big((d + d^*)x^* - dd^*\sigma^2 - 2bd^*x^*\big), \tag{A.5}$$

  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 (A.5) 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 A.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 $X(d^*) = x_C(d^*)$. ■

Lemma A.5 *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 proof of Lemma A.5 analogous to the proof of the first part of Theorem 1 and is omitted. The lemma implies that the sender's preferred conative equilibrium, in the limit, implements outcome $b + \beta(d^*)$ for equilibrium decision d^*.

The second result concerns the existence of a referential equilibrium that does better for the sender than the sender's preferred conative equilibrium, when the option set grows large.

Proposition A.6 *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 a 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 $d^l \geq 4b^2/\sigma^2$ and when $X(d^l)/d^l \leq \sigma^2/(2b)$. Second, if 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 $d^l < 4b^2/\sigma^2$ or $X(d^l)/d^l > \sigma^2/(2b)$ (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.5, the analog of Lemma A.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 $d^l < 4b^2/\sigma^2$,

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

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. ■

D Proofs

D.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.

D.2 Proof of Proposition 1

Proposition 1 follows immediately from the existence of suspicious beliefs that can be established, for example, by Corollary A.1.

D.3 Proof of Proposition 2

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 1. 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 a contradiction.

D.4 Proof of Proposition 3

We prove Proposition 3 for the case of quadratic receiver utility, its extension to the case of general receiver utility of Section 6.1 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 $d_i \leq d_{n-1}$, $X(d_{i+1}; \theta) - X(d_i; \theta) > 2\Delta_n$.

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 $d_i \leq d_{n-1}$,

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

For n large enough,

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

and

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

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 easy 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 5.1, and the general case follows from the second part of Theorem 1.

D.5 Proof of Proposition 4

In this proof, the receiver utility function can be assumed to be quadratic or to satisfy the general assumptions of Section 6.1.

First we show that no receiver-optimal prescriptive equilibrium satisfies the BITH refinement. We proceed by contradiction and focus on the special case $n=2$, which immediately extends to $n>2$. Suppose that an equilibrium (M,D,B) is receiver-optimal, prescriptive, and also satisfies the ε-BITH refinement for some $\varepsilon>0$. Consider the message m_1 that reveals that the outcome of d_1 is y, where y solves $u_{\mathrm{R}}(y)=u_{\mathrm{R}}(b)-\varepsilon/2$. This message must be an off-path message as it were on path, the receiver would choose d_1 upon receiving m_1—as the equilibrium is prescriptive—and the sender would then find it strictly profitable to deviate and send m_1 for all states θ compatible with m_1 but with $u_{\mathrm{R}}(X(d_2;\theta))>u_{\mathrm{R}}(y)$. In addition, and for the same reason, we must have $D(m_1)=d_2$. Thus, applying the ε-BITH refinement, we must have

$$\mathbb{E}^{B(m_1)}[u_{\mathrm{R}}(X(d_2))] \geq u_{\mathrm{R}}(y) + \varepsilon,$$

and since $u_{\mathrm{R}}(y) + \varepsilon = u_{\mathrm{R}}(b) + \varepsilon/2$, no belief $B(m_1)$ can satisfy this inequality.

Second, we argue that every first-point conative sender strategy with parameter $\Delta > 0$ that supports an equilibrium also supports an equilibrium that satisfies the BITH refinement. That fact, together with Proposition 3, implies the second part of Proposition 4 since it then suffices to take the sequence of equilibria used to prove Proposition 4.

Consider any first-point conative sender strategy M with parameter $\Delta > 0$, and assume that M supports an equilibrium, which is equivalent to $\Delta \leq \Delta^{\max}$ as defined in Equation (8) in Section 5.1. Let $\varepsilon = (u_{\mathrm{R}}(b) - u_{\mathrm{R}}(b - \Delta))/2 > 0$. Assume further and without loss of generality that $\Delta < b$.

For the receiver, consider the following strategy D: Upon receiving any on-path message, the receiver chooses the option that is optimal among the options disclosed (and in case of ties, chooses the sender-preferred option). Upon receiving an off-path message m that discloses at least one option whose outcome belongs to $[b - \Delta, b + \Delta]$, the receiver again chooses the option that is optimal among the options disclosed, following the same sender-preferred tie-breaking rule. Then, set $B(m)$ such that the beliefs of any undisclosed option has a very high variance—a variance high enough to make the receiver's decision consistent with $B(m)$.

Finally, if an off-path message m discloses either nothing or options with outcomes outside of the range $[b - \Delta, b + \Delta]$, consider two cases:

(1) For some i, message m discloses $d_i, \dots, d_n$, and the receiver is strictly better choosing one of these options than choosing any other option disclosed in m.

(2) There exists an option $d^\dagger$ not disclosed in m and such that, for every option disclosed in m to the right of $d^\dagger$ (if any), there exists one option disclosed in m to the left of $d^\dagger$ and whose outcome is at least weakly better for the receiver.

In Case (1), as above, set $B(m)$ such that the beliefs of any undisclosed option has a very high variance and set $D(m)$ to be the option that is optimal among the options disclosed, and in case of ties, chooses the sender-preferred option.

In Case (2), set $D(m) = d^\dagger$ and set $B(m)$ such that the mean of $X(d^\dagger)$ is b under $B(m)$ and the variance just high enough so that the expected receiver utility of $X(d^\dagger)$ under $B(m)$ is equal to $(u_{\mathrm{R}}(b) + u_{\mathrm{R}}(b - \Delta))/2$. Note that choosing $d^\dagger$ makes the receiver strictly better off than choosing any other option disclosed in m, for which receiver utility is less than $u_{\mathrm{R}}(b-\Delta)$. Also set $B(m)$ such that the variance of any undisclosed option except $d^\dagger$ is large enough such that the receiver is never weakly better off choosing any undisclosed option except $d^\dagger$.

It is clear that (M, D, B) just defined is an equilibrium: The sender strategy is always optimal by construction, the receiver's decision is optimal upon receiving an off-path message (also by construction) and is optimal upon receiving an on-path message observing that beliefs are then the same as in any first-point conative equilibrium with parameter Δ.

Finally, if m is an off-path message where $D(m)$ is not disclosed in m, then this off-path message corresponds to Case (2) above. Then, for any d revealed in m,

$$\mathbb{E}^{B(m)}[u_{\mathrm{R}}(X(D(m)))] = u_{\mathrm{R}}(b-\Delta) + \varepsilon \geq \mathbb{E}^{B(m)}[u_{\mathrm{R}}(X(d)] + \varepsilon,$$

and hence (M, D, B) satisfies the BITH refinement.

D.6 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 6.1, of which quadratic utility is a special case. We define $\underline{x}$ and $\bar{x}$ as in Section B.3.

Theorem A.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 A.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 6.1). 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 B.3. 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.3, 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 A.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 3 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 A.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

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

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 $\Pr[X(d^*;\theta) \leq y \mid X(d^*;\theta) = x^*]$.[44] Consequently, the expected utility of the 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.3 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 1, 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 1. Therefore, as n grows large, with probability converging to 1, the sender reveals a single option which is not the largest

[44]In general, if, for every n, A_n and B_n are two non-null events and $\Pr[B_n] \to 1$, then $|\Pr[A_n|B_n] - \Pr[A_n]| \to 0$.

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.

D.7 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 6.1).

As explained in Section 3, 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 $\underline{x} = b - \sigma^2/(2\mu)$.

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 $\underline{x}/\mu$.

D.8 Proof of Proposition 5

In this proof, the receiver utility function can be assumed to be quadratic or to satisfy the general assumptions of Section 6.1.

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 A.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 0, in which case it strictly dominates sender-optimal conative equilibrium.

D.9 Proof of Theorem 2

Throughout this proof, we consider the case of general receiver utility under the assumptions of Section 6.1, of which quadratic utility is a special case. We define $\underline{x}$ and $\bar{x}$ as in Section B.3.

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 $x \leq b - \sigma^2/(2\mu)$, 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 $x = b - \sigma^2/(2\mu)$, then the maximum is reached for $\Delta = 0$, and if $x > b - \sigma^2/(2\mu)$, 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

$$b - X(d^r) = \frac{\mu}{\sigma^2} \min\{(b - X(d))^2 : d \in [0, d^r]\} + \frac{\sigma^2}{4\mu}. \tag{A.6}$$

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 $b - \sigma^2/(2\mu)$, then d^r satisfies Equation (A.6). 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.2, this maximizer exists and is unique.

The proof utilizes Lemmas A.6 and A.7 below.

Lemma A.6 *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.2,

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

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.2. As f is strictly concave in Δ by Lemma A.2, $\Delta^{\text{opt}}(x)$ is the only value of Δ satisfying $g(\Delta) = 0$; if $\Delta < \Delta^{\text{opt}}(x)$ then $g(\Delta) > 0$; if $\Delta > \Delta^{\text{opt}}(x)$, then $g(\Delta) < 0$. Since $g(|x - \underline{x}|/C) < 0$ if $|x - \underline{x}|$ is small enough, $\Delta^{\text{opt}}(x) < |x - \underline{x}|/C$ is $|x - \underline{x}|$ is small enough, which concludes the proof. ■

Lemma A.7 *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 A.6, 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.2, 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

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

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,

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

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,

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

which by choice of the constant K is positive is δ is small enough, concluding the proof of Lemma A.7. ■

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 A.7. 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.

D.10 Proof of Theorem 3

Throughout this proof, we consider the case of general receiver utility under the assumptions of Section 6.1, of which quadratic utility is a special case. We define $\underline{x}$ and $\bar{x}$ as in Section B.3. 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 $\underline{x}^n$ is guaranteed by Lemmas A.2 and A.3. 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 (8) of Section 5.1.

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 A.8 *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 A.8. 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 Theorem 3. 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 Theorem 3 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 $N \leq n-1$, if, with probability one, M reveals at least the options $d_0, \ldots, d_N$, then the interval equilibrium weakly dominates Σ.*

Observe that Theorem 3 is included in the case $N = 0$. If $N = n-1$, then, with probability one, the sender communicates a simple message, and by Lemma A.8, 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 strongly prescriptive, 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 strongly prescriptive, 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 (A.7)–(A.9) below:

$$m_0 \notin M(\Theta), \tag{A.7}$$
$$m_0(d_i) < \underline{x}^n \quad \text{for all} \quad i \leq N-1, \tag{A.8}$$
$$0 < \Pr[M(\theta) \in \mathcal{Q} \mid \theta \in \Gamma(m_0)], \tag{A.9}$$

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

Equation (A.7) is implied by Lemma A.8.

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 (A.8) 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 A.8, 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 (A.9).

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 (A.9), 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, $\underline{x}^n \leq m_1(d_{K-1}) < b$.

Lemma A.9 *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 $d_0, \ldots, d_{N-1}$ in Σ, interactions are strongly prescriptive, 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 (A.7)), $m^\dagger$ 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 strongly prescriptive, and by monotonicity of the outcomes over options $d_N, \ldots, d_{K-1}$, $m^\dagger$ must be a simple message that reveals exactly the options $d_0, \ldots, d_k$ for $k \leq K-1$.

Suppose $k < K-1$. Observe that for all $i = d_N, \ldots, d_{K-2}$, 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 A.8, 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 $M(\Theta) \cap \mathcal{M}_\delta \neq \emptyset$. ■

Finally, observe that if $m_1 \in M(\Theta) \cap \mathcal{M}_\delta$, then $D(m_1) = d_{K-1}$.

Step 3: In this step, we establish the inequality

$$\Pr[M(\theta) \in \mathcal{Q} \mid \theta \in \Gamma(m_1)] = \Pr[M(\theta) \in \mathcal{Q}_K \mid \theta \in \Gamma(m_1)] > 0. \tag{A.10}$$

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,

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

Observing that the (Gaussian) conditional distributions of $(\theta_K, \ldots, \theta_n)$ 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 (A.10) 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

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

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.

D.11 Proof of Lemma 4

Lemma 4 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 strongly prescriptive.

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 strongly prescriptive, 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 a contradiction.

D.12 Proof of Lemma 5

Lemma 5 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 a contradiction.

References

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

Bendor, J., S. Taylor, and R. Van Gaalen (1985). Bureaucratic expertise versus legislative authority: A model of deception and monitoring in budgeting. *The American Political Science Review 79*(4), 1041–1060.

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.

Chandra, A., D. Cutler, and Z. Song (2012). Who ordered that? The economics of treatment choices in medical care. In M. V. Pauly, T. G. Mcguire, and P. P. Barros (Eds.), *Handbook of Health Economics*, Volume 2, Chapter Six, pp. 397–432. Elsevier.

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.

Gruber, J. and M. Owings (1996). Physician financial incentives and cesarean section delivery. *The RAND Journal of Economics 27*(1), 99–123.

Hagenbach, J., F. Koessler, and E. Perez-Richet (2014). Certifiable pre-play communication: Full disclosure. *Econometrica 82*(3), 1093–1131.

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.

Levitt, S. D. and C. Syverson (2008). Market distortions when agents are better informed: The value of information in real estate transactions. *The Review of Economics and Statistics 90*(4), 599–611.

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. and J. Roberts (1988). Economic theories of the firm: Past, present, and future. *The Canadian Journal of Economics 21*(3), 444–458.

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.