Imperfect Perception and Vagueness*

Timothy Lambie-Hanson[†] & Giri Parameswaran[‡]

July 16, 2024

Abstract

We present a model that locates the source of vagueness as the speaker's inability to perfectly perceive the world. We show that the agents will communicate clearly about the world as the sender perceives it. However, the implied meaning about the actual world will be vague. Vagueness is characterized by probability distributions that describe the degree to which a statement is likely to be true. Hence, we provide micro-foundations for truth-degree functions as an equilibrium consequence of the sender's perception technology and his optimal, non-vague communication in the perceived world — connecting the epistemic and truth-degree approaches to vagueness.

*We are grateful to the editor and two anonymous referees for their helpful feedback. Neither author has any financial or other conflicting interests to declare.

[†]Department of Economics, Haverford College, Haverford, PA 19041. Email: tlambiehan@haverford.edu

[‡]Department of Economics, Haverford College, Haverford, PA 19041. Email: gparames@haverford.edu

1 Introduction

Vagueness is a common feature of natural languages. A message is vague when the receiver of a message cannot be certain of which 'states' or outcomes the sender of the message had intended to invoke. For example, though we routinely describe people or things as 'tall', 'heavy', 'fast', and so on, we would, in most cases, struggle to identify the boundary between tall and short, or between fast and slow. Indeed, vagueness is often associated with a 'blurring of the boundaries' between the meanings of words.[1]

Many theories seek to explain the nature and source of vagueness. In this paper, we investigate epistemicism — the idea that vagueness arises because agents perceive the world imperfectly, and so cannot describe it in a way that is crisp — as a source of vagueness. To do so, we present a formal model of communication that is in the spirit of (but not identical to) the framework in Williamson (1994). We begin with the observation that communication in the presence of imperfect perception is different from communication with perfect perception; imperfect perception limits the scope of communication to statements about the perceived world rather than about the actual world. We show that imperfect perception does not necessitate that communication about the perceived world be vague. We additionally show that, if this perceived-world communication is extended to the actual (objective) world — i.e. if we try to give meaning to what is based on claims about what appears to be — then meaning will be vague in the actual world. However, this vagueness is metaphysical (it inheres in the objects being described) rather than epistemic, and is closest in spirit to the truth-degree (continuum-valued logic) approach. We thus provide micro-foundations for truth-degree functions as the consequence of the optimal, non-vague communication in the perceived world, and the sender's perception technology, thereby connecting the epistemic and truth-degree approaches to vagueness.

Communication, it has long been recognized, is facilitated when agents are coordinated on a common language (see Lewis, 1969). The sender's choice of message depends on his belief about how that message will be interpreted by the receiver, and the receiver's interpretation will depend on how she expects the sender to use the available messages. The meaning of words and messages, then, are not fully exogenous, but arise as a consequence of a communication game between sender and receiver. To capture this dynamic, we develop a formal model of communication in the spirit of Crawford and Sobel (1982). The sender observes an

[1]A distinct but related concept is 'ambiguity', which arises when a given word has applications in distinct contexts. For example, to say that 'John went to the bank' is ambiguous, in that it is unclear whether John has gone to a financial institution or to the edge of a river. Accounts of ambiguity are not intrinsically tied up with notions of 'degree' in the way that our model of vagueness will be.

informative, but imperfect, signal about a state of the world. The sender transmits a message to the uninformed receiver, who then takes an action that affects both parties. To highlight the effect of epistemic uncertainty, we abstract from other frictions (such as preference disagreement), that may induce vagueness through other channels. Our model characterizes the optimal *use* of messages by the sender, given the anticipated response by the receiver. The *meaning* of those messages is pinned down in equilibrium, given the sender's use. Optimal communication will be *imprecise* if multiple states (when communicating about the actual world) or signals (when communicating about the perceived world) are associated with the same message. By contrast, if communication is *vague*, then (additionally) multiple messages will be associated with the same signal or state.

As we previously noted, the nature of communication differs between perfect and imperfect perception environments. Whereas a perfectly perceiving sender can make objective statements of the form 'John is tall', an imperfectly perceiving sender can only claim that 'John *appears* tall (*to me*)'. Imperfect perception relegates communication to the world of subjective claims, even if the sender renders a statement in a seemingly objective way. We show that imperfect perception alone is insufficient to generate vagueness in the realm of statements about apparent truths. Even if the sender is uncertain that what he perceives is true, this should not prevent him from clearly indicating what he has perceived. To do so, we construct an equilibrium in which the sender optimally partitions the set of perceived states so that each apparent state is associated with precisely one message. For example, there will be a threshold that partitions the set of apparent heights into those that appear tall and those that appear short.

It is common, however, to interpret subjective statements as objective ones — e.g. we often take the statement 'John is tall' to actually be an objective claim about John's height. And we are inclined to do so, even knowing that the sender does not have privileged access to objective truths. Reflecting this tendency, we extend the perceived world language to the actual world, and analyze its meaning in this new space. To be clear, the sender's use cannot be different in the extension, since the sender does not observe the true state when choosing his message. However, if the true state is revealed *ex post*, we may observe the mapping between states and messages, after the fact.[2] Given that the sender's perception is imperfect, there will be borderline cases where the sender classifies persons of the same height as tall in some instances, and short in others. In this *ex post* sense, messages become vague,

[2]This is consistent with Williamson's (1994) account in his motivating example, where a sender claims that there are 30,000 people at a sporting arena, when in reality the actual number is close to, but not exactly 30,000. At the time of his utterance, Williamson's sender does not know the true number of attendees, although we may be able to determine this *ex post*, by counting ticket sales (for example).

since multiple predicates are associated with the same state of the world. Communication that was well defined in the subjective world becomes vague when extended to the objective world. We show that the sender's ideal message use (in the subjective space) combined with the technology that governs perception, induces a probability distribution that describes the likelihood that each predicate is ascribed to a given object. As long as perception is imperfect, this probability distribution will be non-degenerate over a range of 'borderline' outcomes. The characteristic feature of vagueness arises when the probability distribution is degenerate over many states (for which the appropriate predicate to use is clear) but becomes non-degenerate over a subset of 'boundary' states.

In its extension to the objective world, our model has strong similarities to the continuum-valued logic approach to vagueness. This approach rejects the principle of bivalence and instead conceives of statements as having 'truth-degrees' that range from zero (definitely false) to one (definitely true). Since use determines meaning, and our model determines the likelihood of using a particular message to describe a given state, we, in effect, provide micro-foundations for the assignment of truth-degrees. Our model, therefore, unites two distinct approaches to explaining vagueness that are predominant in the literature. We use epistemic theory to provide the causal mechanism that generates the descriptive features of the truth-degrees approach. Truth-degrees are determined *in equilibrium*, given the properties of the optimal communication strategies in the subjective realm and the properties of the technology that governs perception. Similar to canonical truth-degree models, our induced truth-degree functions respect comparisons over ordered objects — if John is taller than Mary, then the truth-degree assigned to John being tall will be at least as large as the truth-degree for Mary being tall. However, in contrast to many truth-degree theories, our truth-degree functions are not truth-functional. Instead, our truth-degree functions satisfy the axioms of probability (which still permits the assignment of truth-degrees to compound statements if the joint-probability distribution is known) as well as standard rules of logic such as the Law of the Excluded Middle.

2 Literature

2.1 Philosophical Accounts

There are many accounts of the sources and characteristics of vagueness. These can be broadly categorized into four approaches: metaphysical, semantic, epistemic, and psychological (see Smith, 2008; Schiffer, 2000). Metaphysical accounts attribute vagueness directly to properties of the object being described. For example, whilst it is clear that a sky-scraper is tall, it is unclear whether a ten-story building ought to be described as tall or not. This clarity, or lack thereof, arises directly from the building's height, and is inherent to the object being described. As a matter of logic, metaphysical accounts are typically forced to reject the principle of bivalence — it may be neither (clearly) true nor (clearly) false that the building is tall — in favor of multi-valued logics (see Halldén, 1949). The continuum-valued logic approach (see Zadeh, 1975; Smith, 2008, amongst others), is a particular instantiation of this approach, which replaces the binary notions of truth or falsity with 'truth-degrees' which can take any value from 0 (clearly false) to 1 (clearly true). These truth-degrees are typically taken as primitives of the language. A useful property for truth-degrees is truth-functionality — the property that the truth degree of a compound statement can be determined purely from the truth degrees of the constituent simple statements. Indeed, most truth-degree proponents endow truth-degrees with this property (Edgington, 1997, being a notable exception). However, as Fine (1975) demonstrates, truth-functionality causes truth-degrees to be inconsistent with the laws of probability and standard results of logic, such as the Law of the Excluded Middle.

Our model departs from the canonical truth-degree approach in two ways. First, truth-degrees are not primitives in our model. Rather, they are determined in equilibrium by more primitive features, such as the sender's perceptive faculties. Second, we construct truth-degrees as probability measures, making them consistent with standard results in logic. We do so at the cost of truth-functionality, although with a complete specification of the joint probability distribution of events, we can still assign truth degrees to compound statements.

Semantic accounts attribute vagueness to indeterminacy in the way messages are used by different speakers. Plurivaluationism (see Smith, 2008), captures this idea that different speakers may describe the same object differently. It is closely related to, although distinct (as Smith, 2008, takes pains to argue) from, Supervaluationism (see Dummett, 1975; Fine, 1975; Keefe, 2000), which posits that messages are vague when its extension to indeterminate cases admits multiple interpretations. Under this approach, vagueness arises because of an inability by the community to coordinate on a common use of language.

The epistemic account (see Williamson, 2002; Sorensen, 2001) locates vagueness in the limitations of human perception. Proponents of this approach insist that vagueness is neither metaphysical (it does not inhere in objects) nor semantic (it is not a consequence of who is communicating). Language itself is well-defined and characterized by sharp thresholds, and a perfectly informed speaker would use messages in ways that are consistent with their meanings. Vagueness arises because imperfect perception prevents agents from precisely comparing the true state of affairs against these thresholds. A challenge for the epistemic theorist is that, by this account, these thresholds are seemingly determined independently of the speaker's usage, thereby severing the link between use and meaning.[3]

Our model lays bare this challenge. The speaker's use of messages is determined by what he perceives, and his expectation of how any given message will be interpreted. The receiver, in turn, interprets the meaning of messages based on her expectation of how the sender uses each message. This implies that, if the parties communicate optimally, message use will be characterized by firm thresholds in the subjective world. However, meaning necessarily cannot be governed by firm thresholds when extended to the actual world, since the same actual state may be mapped onto multiple signals that are associated with different messages. If thresholds exist that delineate meaning in the actual world, they must be generated by some mechanism other than the sender's use of messages.

Other explanations of vagueness fall into the category of psychological accounts. The idea of typicality and the degree to which an item belongs to a particular group has long been discussed by philosophers (see Murphy, 2004, for a thorough review). More recent work (e.g. see Hampton and Jönsson, 2012) discusses how even items that may be well-defined when considered alone, may be become vague in combination. In contrast to our model, these accounts are not typically concerned with gradable predicates, which makes a direct comparison to our analysis, difficult.

However, the concept of vagueness-related partial beliefs (VRPBs), introduced in Schiffer (2000) and further developed in Schiffer (2003), provides an avenue for connecting psychological accounts of vagueness in gradable predicates to our approach. In these accounts, when confronted with a borderline case of a property, the receiver associates a vagueness-related partial belief $\in (0, 1)$ with the particular property. The innovation is that VRPBs do not operate under the laws of probability. For example, the law of the excluded middle may be violated by one's VRPBs. Our account of vagueness adheres to classical probability theory

[3]Williamson (1994) argues that the mechanism linking use and meaning may be complicated and unknown to the philosopher — but that this in no way refutes that the former determines the latter. We find this account difficult to sustain. Plainly, if use determines meaning, it cannot be that meaning is determined by factors inaccessible to the speaker when choosing which words to use. The mapping from use to meaning should be readily determined by simply observing how the sender uses his words. We understand the meaning of the word 'tall' by observing all the instances in which we describe an object as tall. To say that use determines meaning isn't to merely suggest that there is *some* mechanism that links the meaning of words to their use. Rather, it is the stronger claim that use is itself that mechanism. Smith (2008) provides a more detailed critique.

and so at that primitive level seems in conflict with the psychological account. However, our model endogenously identifies borderline cases and derives a probability that each of these borderline cases will be described by a given message. Analogously, we can use our model to consider the probability that a given message refers to a borderline case, something that Schiffer (2010) considers seriously in the world of vagueness-related partial beliefs. Parikh (2019) builds a model that incorporates the psychological account and discusses psychological mechanisms by which beliefs might shift in ways that violates the laws of probability. While not entirely complementary, the commonalities between this approach and ours are notable.

2.2 Economic Accounts

There is a long literature on the economics of communication dating back to the canonical models of persuasion (see Grossman, 1981; Milgrom, 1981) and cheap talk (see Crawford and Sobel, 1982). Crawford and Sobel (1982) study a communication game between an informed sender and an uninformed receiver who must take an action that affects both parties. The paper provides two important insights. First, it demonstrates the (equilibrium) relationship between use and meaning; the sender's use is determined by the meaning ascribed to each message by the receiver, and these ascribed meanings are in turn determined by the sender's use. Second, differences in preferences between a sender and receiver generate incentives for the sender to not fully reveal his information to the receiver, thereby rationalizing imprecise communication.[4] Qing and Franke (2014), building on the model in Lassiter and Goodman (2014), present a variant of this analysis in which the communicants' strategies are probabilistic, reflecting satisficing rather than perfectly maximizing behavior.

While some economists have directly applied the cheap talk framework to linguistics and message meaning (see Jäger, Metzger and Riedel, 2011), until recently, economic models of communication did not feature messages that could be construed as vague (Lipman, 2003, 2009). Blume and Board (2013*a*), with what they term 'message indeterminacy', are amongst the first to study what we term vagueness. They do so by assuming uncertainty about language competence (which roughly corresponds to the richness of vocabulary). 'Message

[4] Imprecision is to be distinguished from vagueness. A message is imprecise if the sender associates multiple sates with that message, thereby preventing the receiver from exactly learning the true state. If a message is vague, it will additionally be unclear which set of states the message seeks to invoke. For example, it is imprecise to say that 'John's height is at least 6 feet', since saying so provides the receiver with a range of possible heights for John, rather than his actual height. (We implicitly assume that John's actual height is salient to the receiver.) However, the message is not vague — it clearly delineates the set of heights that John may have.

indeterminacy' arises when the receiver is uncertain about the sender's language competence. Lambie-Hanson and Parameswaran (2016) study a communication game in which message use is ideally modified to suit different 'contexts'. (For example, 'tall' would be used differently in the context of the town of Lilliput as compared to the town of Brobdingnag.) In such a game, meaning will be clear as long as the sender correctly perceives the receiver's belief about the context. Vagueness arises when the sender's message use fails to coincide with the receiver's belief about how the sender is communicating. As an example, a sophisticated Lilliputian who knows to use the word 'tall' differently when communicating in the Brobdingnag context, will be able to communicate without misunderstanding. Vagueness arises, not because different contexts *per se*, but because of a lack of common knowledge about how each communicant is modifying his use to suit the context at hand. Both these theories are *semantic* in that they locate the source of vagueness in differences in (expected) message use between the communicants. Importantly, in both cases, speakers are not intentionally vague. The sender always transmits a message with well-defined meaning; vagueness arises when the meaning inferred by the receiver and the meaning intended by the sender, diverge.

A different approach locates the source of vagueness in frictions in the communication technology itself, which may cause messages to become 'garbled' during transmission. Blume and Board (2014) demonstrate that garbling may provide an incentive for the sender to be intentionally vague. Rick (2013) similarly shows that there may be deliberate miscommunication in the presence of garbling, and that this may improve outcomes for both parties.

3 Model

Let $X = [0, 1]$ be the set of pay-off relevant states. There is a partially-informed sender (he) who observes a noisy signal y about the true state x, and an uninformed receiver (she) who has no information about the state. Both communicants share common prior beliefs over the likely realization of the state, represented by a distribution function F with associated continuously differentiable density f. We assume that F has full support on $[0, 1]$, so that $f(x) > 0$ for all $x \in [0, 1]$. The receiver must take an action $a \in X$ that affects both the sender and receiver. We abstract from cases where the sender has incentives to hide information from the receiver (such as the standard models of cheap talk) by assuming that the sender and receiver have identical preferences. This enables us to focus attention on the effect of perceptive limitations on communication. For simplicity, we assume that both agents have state-dependent preferences represented by the utility index: $u(x,a) = -(x-a)^2$. (Our results can be generalized to accommodate any convex loss function.) Intuitively, the agents seek to match the action to the realized state of the world, and suffer increasingly larger losses as the action deviates from the true state.[5]

The sender observes signal $y \in Y = [0,1]$. The signal technology is represented by a (conditional) distribution function $Q(y|x)$ with support on a convex subset of $[0,1]$, that admits a continuously differentiable density $q(y|x)$. The density $q(y|x)$ is the 'likelihood' that the sender observes signal y given that the true state is x. We make two additional assumptions about the signal technology. First, we assume that for each $y \in [0,1]$, there is some $x \in [0,1]$ such that $q(y|x) > 0$. This ensures that the (unconditional) signal distribution has full support on $[0,1]$. Second, we assume that the signal is informative in the sense that a higher signal statistically indicates a higher true state. This property is formalized by assuming that Q satisfies the monotone likelihood ratio property. I.e. if $x_1 > x_0$ then $\frac{q(y_0|x_1)}{q(y_0|x_0)} \leq \frac{q(y_1|x_1)}{q(y_1|x_0)}$ whenever $y_1 > y_0$.

The monotone likelihood ratio property is standard in signaling games. It implies that if the true state is high, then the sender will be more likely to observe a high signal than a low signal, than if the true state is low. Apart from this restriction, the perception technology is quite general, and there is considerable scope for variation in precision and bias. For example, the technology is consistent with a sender who perceives very tall and very short buildings quite accurately, but is more error prone when observing buildings of intermediate height. Similarly, it is consistent with a sender who systematically misperceives buildings as taller than they are.

After observing a given signal, the sender can make inferences about the true state of the world, according to Bayes' Rule. Let $F(x|y)$ denote the sender's posterior belief about the true state after observing signal y.[6] The monotone likelihood ratio property implies that these posterior beliefs respect first-order stochastic dominance.[7] After observing a higher signal, the sender rationally infers that the true state is more likely high than low.

Upon observing the signal, the sender can send a message $m \in M = \{m_1, ..., m_K\}$ to the

[5]Under our framework, all states are pay-off relevant, and so knowledge of the state is salient to the receiver. We could imagine a more expansive state-space, in which subsets of states were pay-off irrelevant, for example because they encoded information that was incidental to the agents' utility. If so, the receiver might not care to distinguish between pay-off irrelevant states; a message that pooled pay-off irrelevant states might reasonably be considered to be precise. This complication does not arise in our model, since all states are pay-off relevant.

[6]These posterior beliefs are well defined, since every signal $y \in [0,1]$ can be generated.

[7]I.e. $y_1 > y_0$ implies $F(x|y_1) \leq F(x|y_0)$ for every $x \in X$.

receiver. The set of messages is finite, capturing the idea that the communicants' share a limited vocabulary, and ordered, capturing the grades of an adjective. We can think of each message as behaving similarly to a first-order predicate. To transmit message m_k is to ascribe to a subject the k^{th} degree of a gradable adjective (with K possible degrees). In this section's motivating example, we had three possible messages, with m_1 indicating that the subject is 'short', and m_3 indicating that the subject is 'tall'.[8]

A strategy $\mu : Y \to M$ for the sender assigns a message $\mu(y)$ to each signal $y \in Y$. A strategy $\alpha : M \to X$ for the receiver assigns an action $\alpha(m)$ to each message received. Let $F(x|y)$ be the sender's posterior belief about the true state after receiving signal y, and let $G(x|m)$ be the receiver's posterior belief about the true state after observing message m. Let $f(x|y)$ and $g(x|m)$ be the associated densities. A *Perfect Bayesian equilibrium* consists of a strategy μ for the sender, a strategy α for the receiver, and a pair $(f(\cdot|y), g(\cdot|m))$ of belief functions which satisfy:

1. For each signal $y \in Y$, the sender chooses the message which maximizes his expected utility, given his posterior beliefs and the equilibrium strategy of the receiver:

$$\mu(y) = \arg\max_{m \in M} \left\{ -\int_X (x - \alpha(m))^2 f(x|y)\, dx \right\}$$

2. For each message $m \in M$, the receiver chooses the action which maximizes her expected utility, given her posterior beliefs:

$$\alpha(m) = \arg\max_{a \in X} \left\{ \int_X (x - a)^2 g(x|m)\, dx \right\}$$

3. The communicants' posterior beliefs are determined according to Bayes' Rule, given their common prior beliefs and the equilibrium strategies:

$$\begin{aligned} f(x|y) &= \frac{f(x)\, q(y|x)}{\int_{z \in X} f(z)\, q(y|z)\, dz} \\ g(x|m) &= \frac{f(x) \left[\int_{y \in Y} q(y|x)\, \mathbf{1}_{\mu(y)=m}[y]\, dy \right]}{\int_{z \in X} f(z) \left[\int_{y \in Y} q(y|z)\, \mathbf{1}_{\mu(y)=m}[y]\, dy \right] dz} \end{aligned}$$

[8] Many of our results would continue to hold if the sender had access to a larger set (e.g. a continuum) of messages. We limit our attention to finite message spaces for two reasons. First, we think it is reasonable and realistic in the context of gradable adjectives. Second, as we discuss on p.22, the properties of truth-degree functions under finiteness more closely match the defining characteristic of vagueness — non-degeneracy of truth-degrees (or 'blurring') at the boundaries.

for every message m that is transmitted with positive probability.

Given strategies (μ, α), the *ex ante* expected loss from communication is:

$$\ell(\mu, \alpha) = \int_{x \in X} \left[\int_{y \in Y} (x - \alpha(\mu(y)))^2 q(y|x) dy \right] f(x) dx$$

This measure of loss captures the imperfections in communication that result in the receiver taking an action different from the one best suited to the given state. We say that an equilibrium is *optimal* if there is no feasible perturbation — that either adds to, or removes from, the set of distinct actions that the receiver will take — that causes the *ex ante* expected loss from communication to decrease.[9] A sub-optimal equilibrium may arise if the sender's strategy involves redundancies, where multiple messages induce the receiver to take same action. An example is the 'babbling' equilibrium, in which, after any signal, the sender transmits a message at random, and the receiver, understanding that messages are uninformative, takes the same *ex ante* optimal action regardless of the message received. Though equilibrium consistent, such strategies forgo opportunities for the sender to transmit valuable information to the receiver. In this paper, we focus on equilibria that are optimal, since we are in a common values setting where the parties are strongly incentivized to communicate as effectively as possible,

Even limiting attention to optimal equilibria, it is well known that communication games of this sort typically admit multiple equilibria. In particular, for any given equilibrium, a related equilibrium can be constructed by simply permuting the messages. For example, the sender may use the word 'tall' to describe a person who appears small-heighted and use the word 'short' to describe a person who appears large-heighted. This is an equilibrium provided that the receiver understands the sender's usage. An equilibrium of a communication game is simply a commonly understood code; the same information could be transmitted, if the code mapping signals to messages were reversed (say).

As previously mentioned, in this paper, we assume that there is a pre-existing, exogenous ordering over messages that associates m_1 with the lowest adjective grade and m_K with the highest. This requires that, though the meanings of messages are pinned down in equilibrium, messages enter the model with *some* exogenous content.[10] For example, the parties

[9]If, fixing the number of distinct messages transmitted, equilibria were unique, then the optimal equilibrium is the one that minimizes the *ex ante* loss from communication.

[10]See Blume (2021) for a detailed discussion of the necessity of messages having exogenous content to facilitate communication.

will expected that the message ‘tall’ will be used to refer to large-heighted people and that ‘short’ will refer to short-heighted people. We focus on equilibria that respect this natural exogenous ordering over messages. Consistent with the exogenous understanding, the equilibrium analysis then determines more concretely which apparent heights are associated with each message.

This paper uses a formal model to analyze the effect of imperfect perception on the nature of communication. As is standard in all formal models, we make several assumptions that keep the model simple and tractable. Our goal is to focus on factors relevant to the issue of interest; namely the effect of imperfect perception on communication. To this extent, we abstract from other factors that may be salient in their own right, but are not crucial to the epistemic story. Before proceeding to the equilibrium analysis, we briefly comment on some of our modeling choices.

First, though our notion of equilibrium requires that both sender and receiver to understand and best respond to the strategic environment, we are sensitive to the objection that humans are often not nearly so sophisticated. We do not dispute that agents often operate according to some ‘exogenous’ view of how language works, based on their experience, intuition and internalized ‘rules of thumb’. (Indeed, our assumption that messages are naturally ordered requires that messages be endowed with *some* exogenous content[11].) Rather, in invoking the idea of equilibrium, our point is that if this ‘exogenous’ language is not equilibrium consistent, then either the sender or receiver (or both) will have incentives to use or interpret messages differently from the exogenous understanding. By contrast, if the exogenous language is equilibrium consistent, then no such incentive to ‘mis-use’ will arise. Our notion of equilibrium can thus to be understood as a situation where use and meaning are stable.[12]

Second, our analysis is confined to instances where the agents seek to communicate about a property that exists on a grade (such as height). We acknowledge that such communication does not typically occur in isolation; the parties will also need to use language to establish any number of other salient facts, including the relevant context, the identity of the subject, amongst others. In our analysis, we abstract from all these other processes, confining our analysis to communication about the object’s grade. To this end, our use of the Crawford and Sobel (1982) framework is appropriate for our task. But, of course, that framework

[11]Several models, including Crawford (2003), Franke (2014), and Blume (2021), specify the existence of an exogenous ‘level-0’ language, with commonly understood meaning, which agents may use and interpret strategically. For our purpose, fully specifying an exogenous language is not necessary — it suffices to note that messages will used in their natural order.

[12]This interpretation of equilibrium — as the long run stable play of agents — is common in economics (see Osborne et al., 2004).

need not be applicable in modelling other aspects of communication, for which alternative frameworks have been developed (see Parikh, 2019).

Third, though we assume that messages are used according to their natural ordering, we are agnostic as to the particular labels attached to those messages. For example, the sender may describe a building's height as either 'tall', 'neither-tall-nor-short', or 'short'. Or in describing preferences on a political spectrum, the sender may characterize a politician as belonging to the 'left', 'center', or 'right', but also to the 'center-left' or 'center-right'. For our purposes, these will understood as distinct messages. And the fact that the labels attached to the some messages are hybrids (of sorts) of the others, does not affect how meaning is attached to any of the messages in equilibrium, except that meaning must respect the natural ordering over messages. Of course, the inclusion in the message space of messages with hybrid labels requires the presence of messages with the more basic labels. And the presence of certain messages enables the inclusion of other messages whose labels are hybrids of the former. Our framework easily accommodates these details. Accordingly, we will understand the message space to include all of the messages that the sender may wish to use, including those whose labels may be hybrids of others.[13]

Finally, a simplifying feature of our model is that both communicants share a common prior over the likely state of the world, and that all aspects of the model (other than the true state) are commonly known by the players. In particular, we assume that the receiver understands the nature of the sender's perception technology. (Since the sensory abilities of humans are roughly similar, we think it is not unreasonable to assume that the receiver can predict how the sender may misperceive the world.) Again, we acknowledge the strength of these assumptions, and the reality that the communicants' beliefs about these objects may not perfectly align. However, whilst such differences may affect the nature of communication, they are not intrinsically linked to the problem of imperfect perception. To the extent that these features induce vagueness, they do so through the channel of interpersonal differences between the communicants, and therefore more properly represent a *semantic* source of vagueness rather than an epistemic one (see Blume and Board, 2013*b*; Lambie-Hanson and Parameswaran, 2016; Körner, 1962). Since these forces would continue to operate even if the sender's perception were perfect, their abstraction does not pose a threat to understanding the epistemic account of vagueness.

[13]It may be objected that such a construction is not well defined, since for any set M, one can always construct a larger set M' with $M \subset M'$, where M' includes additional messages whose labels are hybrids of the labels associated with message set M. But, in practice, there is a limit to how far this process can be pushed. For example, it is rare to find hybrids of hybrids. Though we are happy to squeeze 'center-left' between 'center' and 'left', we would typically not further squeeze 'center-center-left' between 'center' and 'center-left', nor would we describe a building as 'between short and neither-tall-nor-short'.

Similarly, we assume that there is common knowledge about the set of messages that available for use. Blume and Board (2013a) show that a lack of common knowledge in this dimension is another channel through which vagueness (what they call 'message indeterminacy') can arise. Since that channel has already been explored, we abstract from it here, though we discuss the differences between our channel and various others in Section 4.4. Abstracting from message indeterminacy also highlights key insights of this paper: that vagueness can arise even when there is common knowledge about the message space, and that imperfect perception operates as an independent source of vagueness to these others.

4 Analysis

Recall that to transmit message m_k is to ascribe to the subject the k^{th} degree of the gradable adjective. But what precisely is the subject? We previously distinguished between subjective messages that described what the sender perceives from objective messages which describe what is. Subjective message are of the form: 'The state appears to have property m_k'. Since such statements are conditioned upon the signal received by the sender, we say that they live in Y-space. By contrast, objective statements are of the form: 'The state actually has property m_k'. Since these statements are about the true state, we say that they live in X-space. From herein, we use X-space and Y-space as a shorthand for indicating objective and subjective claims, respectively.

With these distinctions in mind, we turn to solving the model. Our analysis is in two parts. First, we characterize an optimal equilibrium of the communication game between the imperfectly informed sender and the uninformed receiver. The equilibrium determines how messages will be used by the sender to describe the world as it appears to him. We then analyze the extension of this language to claims about the objective world.

4.1 Equilibrium and Properties of 'Apparent' Statements

We being by characterizing an optimal equilibrium of this game:

Proposition 1. *There exists an optimal equilibrium of the communication game. In any optimal equilibrium, the sender will use a threshold strategy and utilize all K messages. An optimal equilibrium is characterized by a vector $(s_0, ..., s_K) \in Y^{K+1}$ with $0 = s_0 < ... < s_K = 1$ such that:*

1. *The sender transmits message*

$$\mu(y) = \begin{cases} m_1 & \text{if } y \in [s_0, s_1] \\ m_k & \text{if } y \in (s_{k-1}, s_k] \text{ for } k = 2, .., K \end{cases}$$

2. *The receiver takes action* $\alpha(m_k) = \int_0^1 xg(x|m_k)\,dx$ *after receiving message* m_k*; and*

3. *The communicants' belief functions satisfy:*

$$\begin{aligned} f(x|y) &= \frac{f(x)\,q(y|x)}{\int_0^1 f(z)\,q(y|z)\,dz} \\ g(x|m_k) &= \begin{cases} \frac{f(x)\int_{s_{k-1}}^{s_k} q(y|x)dy}{\int_0^1 f(z)\left[\int_{s_{k-1}}^{s_k} q(y|z)dy\right]dz} & \text{if } y \in (s_{k-1}, s_k] \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Additionally, if the signal technology is unbiased (so that $E[x|y] = y$ *for each signal* y*) and the unconditional distribution of signals is uniform, then the optimal equilibrium is unique.*

In an optimal equilibrium, the sender partitions the signal space into K disjoint intervals, such that each interval is associated with a given message. The sender transmits message m_k whenever the received signal is contained within the k^{th} interval. Since $s_{k-1} < s_k$, each message is transmitted with positive probability. Several properties of the equilibrium are worth noting.

First, given the signal technology, both communicants form beliefs about the likely true state according to Bayes' Rule. For example, after observing signal y, the sender's conditional belief about the true underlying state x is given by the density $f(x|y)$. Let $S_S(y) = \{x \in X|\, f(x|y) > 0\}$ be the support of the sender's conditional beliefs, which is the set of possible true states that the sender cannot rule out. This corresponds to the *'margin for error'* in Williamson (1994). The receiver similarly forms beliefs about the likely true state. Although she doesn't observe the sender's signal, she can make inferences given the message she receives and her knowledge of the sender's optimal communication strategy. Upon receiving message m, the receiver's belief that the true state is x is given by the conditional density $g(x|m)$. In determining their optimal choices, both sender and receiver use these updated beliefs about the true state; i.e. both players take into account the possibility that the sender misperceives when making their choices.

Second, upon receiving message m_k, the receiver's choice of optimal action $\alpha(m_k)$ simply reflects her best guess about the true state, given her information. If the receiver knew the state perfectly, she would choose the action that precisely matched the state. Since she does not, she chooses the action that matches the state in expectation, given her updated beliefs. It should thus be clear that the modeling fiction of the receiver taking an action simply serves to capture the process of information transmission between sender and receiver.

Third, the sender's optimal strategy assigns a single message to each signal. To see why, note that the receiver chooses a different action for each different message received. If so, the sender will generically not be indifferent between transmitting each of the available messages, but will rather have a strict incentive to send the message that induces the action that is closest to the sender's expectation of the true state. Hence, there will generically be a unique message associated with each signal. In spite of the sender's uncertain perception, he will typically be certain about the message he wishes to send, given what he perceives and given the receiver's anticipated response.

Fourth, by partitioning the signal space into K disjoint intervals, the sender transmits a more informative message to the receiver than would be the case if either he used fewer than K messages, or if the intervals associated with different messages overlapped. It is in this sense that equilibria are optimal. The Proposition verifies that an optimal equilibrium exists and is in threshold strategies.

The baseline assumptions outlined in Section 2 do not guarantee that the communication game will admit a unique optimal equilibrium. The final part of Proposition 1 provides sufficient conditions for the equilibrium to be unique, analogous to Theorem 2 in Crawford and Sobel (1982). As with their result, the sufficient conditions that we provide are quite strong.[14] However, as Example 1 demonstrates, and as Crawford and Sobel themselves note in their remarks following Theorem 2, there may be unique equilibria even when the sufficient conditions are not met.

Since the sender's optimal strategy partitions the signal-space, we can find thresholds $\{s_0, ..., s_K\}$, which delineate the intervals and determine which message is sent. For example, there will be some threshold perceived height, such that the sender will report that the 'building appears tall' whenever his signal of the building's height exceeds this threshold.[15]

[14] Though Crawford and Sobel specify their sufficient conditions slightly differently, in the case of quadratic preferences, their conditions reduce to the requirement of a uniform prior over states.

[15] Though the existence of a sharp threshold may seem stark, to a first order, we think it well captures how agents intend (or attempt) to communicate. When assessing the temperature of water that is being heated, one starts by reporting 'warm' and then switches to reporting 'hot' when the perceived temperature is sufficiently high, consistent with threshold behavior.

Two comments about this threshold strategy are worth noting.

First, we stress that the threshold strategy arises as an equilibrium result, rather than as an assumption of the model. Nothing in our model compels the sender to use a threshold strategy. Rather, if the receiver chooses different actions after different messages, then the sender will want to partition the state space. And, as long as the sender partitions the state-space, the receiver will want to take different actions after different messages.

Second, we address the common objection[16] that the location of thresholds is arbitrary and so threshold strategies ought to be impermissible. It is indisputable that in drawing thresholds, we distinguish seemingly similar states which just happen to fall on opposite sides of the threshold. Taken in isolation, such distinctions do indeed appear arbitrary. Nevertheless, when considered globally, these thresholds are in fact located optimally. Threshold strategies are a consequence of the agents' limited vocabulary. If there were no limit on the number of degrees that we could express, we would associate a separate message with every possible signal, thereby appropriately acknowledging every nuance and distinction between signals. Since we make the reasonable assumption that our vocabulary is limited, we are forced to 'pool' several states into the same message. An unavoidable consequence of pooling is that some pairs of states will be treated identically when pooled together even though they are distinct, whilst other seemingly similar pairs of states will be treated differently by virtue of not being pooled together. The more states that are pooled into the same message, the less informative that message will be. The challenge for optimal communication is to pool states together in the way that best facilitates information transfer. In our model, the location of the thresholds $\{s_0, ..., s_K\}$ have the property of minimizing the expected (square) deviation between the true state and the receiver's expectation of the state. It should be clear then, that these thresholds are not located arbitrarily. Instead, their location depends on the global properties of the system, anticipating the agents' likely communication needs.

An important consequence of the above proposition is that language is not vague in the Y-space. The sender's communication strategy is characterized by an unambiguous mapping from signals to messages, and this is understood by the receiver. Given the perceived height of any building, the receiver knows whether the sender will describe it as tall or not. Although imperfect perception may leave the sender with some doubt about the true state, it does not prevent him from clearly indicating the signal that he has perceived. This is consistent with

[16] For example, the Sorites Paradox, which is commonly associated with the problem of vagueness, arises precisely because of a rejection of threshold behavior. With threshold behavior, the induction argument that generates the Sorites series would not hold globally.

the critique in Lipman (2009), that a speaker should not be intentionally vague in his use messages.

To demonstrate the features of optimal equilibria, we construct the following stylized example:

Example 1. Suppose the state x is drawn from a uniform distribution on $[0,1]$, and that, conditional upon the realized state x, the sender observes a signal y, which is itself drawn from a uniform distribution on $\left[\frac{x}{1+2\varepsilon}, \frac{x+2\varepsilon}{1+2\varepsilon}\right]$.[17] The signal precision, or 'margin for error' is parametrized by $\varepsilon > 0$, where a larger ε implies more imperfect perception. (We assume $\varepsilon < \frac{3}{8}$ for technical convenience.) For any true state, the sender's signal is contained within a band of uniform width. The size of this band indicates how accurately the sender perceives the world.

Suppose $K = 3$, so that the sender has access to three messages (e.g. small, medium and large). Then, there is an equilibrium characterized by thresholds: $s_1(\varepsilon) = \frac{1}{6} + \frac{4\varepsilon+\sqrt{1+4\varepsilon^2}}{6(1+2\varepsilon)}$ and $s_2(\varepsilon) = \frac{5}{6} - \frac{4\varepsilon+\sqrt{1+4\varepsilon^2}}{6(1+2\varepsilon)}$. The sender transmits message 1 whenever he observes a signal y in the interval $[0, s_1(\varepsilon)]$, he transmits message 2 whenever he observes a signal in the interval $(s_1(\varepsilon), s_2(\varepsilon)]$, and he transmits message 3 whenever he observes a signal in the interval $(s_2(\varepsilon), 1]$. The receiver's optimal action after each message are: $a_1(\varepsilon) = \frac{\sqrt{1+4\varepsilon^2}}{3} - \frac{1}{6}$, $a_2(\varepsilon) = \frac{1}{2}$ and $a_3(\varepsilon) = \frac{7}{6} - \frac{\sqrt{1+4\varepsilon^2}}{3}$. We provide a full characterization of this equilibrium, including the equilibrium belief functions, in the Appendix.

As a benchmark, note that if $\varepsilon = 0$, so that the sender perfectly perceives the world, then $(s_1, s_2) = \left(\frac{1}{3}, \frac{2}{3}\right)$ and $(a_1, a_2, a_3) = \left(\frac{1}{6}, \frac{1}{2}, \frac{5}{6}\right)$. The sender would partition the state-space into three equally sized intervals, and the receiver implements the action which corresponds to the expected state in each interval. Partitioning the state space into equally sized intervals ensures that message sent is equally informative, no matter which state of the world is realized.

When $\varepsilon > 0$, we notice that both the sender's thresholds and the receiver's optimal actions, are responsive to the signal precision ε — both communicants are aware that the sender imperfectly perceives the world, and they adjust their use and understanding of messages accordingly. Importantly, an imperfectly perceiving sender's use may systematically vary from the perfect-perception benchmark; the imperfectly informed sender doesn't simply apply the

[17]The signal structure is not as complicated as it may seem. What we have in mind is a signal with conditional distribution $y \sim U[x-\varepsilon, x+\varepsilon]$. But this produces signals that lie outside the assumed signal space $[0,1]$. We simply do a linear re-scaling of signals to ensure $y \in [0,1]$.

perfectly-perceiving sender's thresholds to the imperfect signals that he observes.[18] □

In summary, we have constructed an equilibrium in which it is optimal for an imperfectly perceiving sender to partition the signal space into disjoint intervals (characterized by thresholds), and to associate a distinct message with each interval. Accordingly, we have shown that imperfect perception alone is not sufficient to cause messages to be vague. As we noted previously, vagueness may still arise in this environment if we introduce additional frictions to communication. For example, we could introduce multiple 'types' of receivers, where a 'type' may capture differences in the receiver's beliefs about the prior distribution of the state, the perception technology, the size of the sender's vocabulary, and so on. Blume and Board (2013*a*) and Lambie-Hanson and Parameswaran (2016) demonstrate that introducing higher order uncertainty about these types is sufficient to induce semantic vagueness. And, of course, the agents could make mistakes and misapply equilibrium strategies. However, since our focus in this paper is on the effect of imperfect perception, and these effects operate independently, we can safely abstract from those other channels.

4.2 Properties of Statements about Actualities

We now turn our attention to statements about actualities, rather than appearances. We begin by noting that, if the sender could directly observe the true state x, we could simply repeat the above exercise and characterize optimal communication strategies in X-space. This ideal communication in X-space would retain all of the characteristics of the optimal Y-space strategies, including that the use of messages is delineated by sharp thresholds. But for the sender's fallible perception, these would be the (epistemic theorist's ideal) thresholds that governed our use of messages. We caution, however, that, if we take impediments to perfect perception seriously, such an exercise is no more than a thought experiment. Whilst we can conceive of such a language, it is not actually available for use by an imperfectly-perceiving an sender. Moreover, whilst we can insist on the meaning generated by such a thought experiment as being 'correct' or 'ideal', doing so necessarily severs the relationship between actual use (by imperfectly perceiving senders) and meaning. If use is to determine meaning, then use cannot be conditioned upon information to which the sender lacks access.

[18] Indeed, as the signal imprecision ε increases, the sender will be more likely to transmit messages m_1 and m_3, and less likely to transmit m_2. With greater imprecision, the receiver recognizes that a given signal is consistent with a larger range of true states. Then, if the thresholds did not change, the average state which generated message m_1 would be higher, and the average state which generated m_3 would be lower — causing the receiver to choose higher α_1 and lower α_3, respectively. But, this feeds back into the sender's choice, making him less inclined to transmit m_2.

Instead, we take the following approach: We retain the assumption that the sender imperfectly perceives the world, and instead ask how to give meaning to those messages as claims about the actual world. Recall that, when the sender transmits message m_k, this has the unambiguous meaning in Y-space, that the sender's signal is contained in the interval $(s_{k-1}, s_k]$. The receiver is able to determine this meaning because she understands which signals are associated with which messages. Similarly, to determine meaning in X-space, the receiver must understand the mapping from states to message. Of course, we have just argued that the sender can only condition his message on his signal, and not the true state. Since a given state can generate multiple signals, a state may be associated with multiple messages. The mapping from states to messages need not be unique. This is the channel through which vagueness arises.

We formalize this idea. Suppose the true state is x, and that the sender receives signal y which is consistent with the signal technology Q. Given the above discussion, we know that the sender will transmit a unique message $\mu(y)$ which depends on the signal y. The communication strategy is not random or probabilistic. However, from the perspective of an external observer who can perfectly perceive the state, but not the sender's signal, the sender's communication strategy may appear random — since the sender may send different messages after receiving different signals in the same state. (One way to conceive of such an observer is to suppose that the true state is revealed, *ex post*, and agents keep track of the frequency with which the sender transmits each message in each state.) Let $\phi_k(x)$ denote the probability that state x is characterized as having property m_k. We know that:

$$\phi_k(x) = \int_{\{y \in Y \,|\, \mu(y) = m_k\}} q(y|x)\, dy = \int_{s_{k-1}}^{s_k} q(y|x)\, dy$$

The function $\phi_k(x)$ indicates the probability that property m_k is ascribed to state x by an imperfectly perceiving sender (who only observes a noisy signal y of x). By the laws of probability, we know that $\sum_k \phi_k(x) = 1$. We can think of $\phi_k(x)$ as the membership function which determines the application of predicate m_k to state x. Since use determines meaning, these probabilities have the natural interpretation as the *degree of truth* that x has property m_k. If the sender is more likely to associate state x with property m_k than m_l, then it is natural (in the sense of use determining meaning) to assign a higher degree of truth to x having property m_k than m_l. We stress that these truth degree functions are not primitives of the model (i.e. they are not taken as exogenous facts). Rather, they are determined in equilibrium by the sender's optimal Y-space message use and the signal technology which determines the sender's perception.

Before characterizing the properties of truth-degrees, we briefly digress to make the following observation: The process of extending an optimal Y-space communication to X-space is analogous to the procedure that determines truth (or super-truth) under the supervaluationist approach. For each actual state x, we consider every signal y that could potentially be observed by the sender (given the perception technology Q), and ask which message the sender would transmit, given his observed signal. The state x is then *definitely* characterized by property m_k if, under every possible signal-realization, the sender would transmit message m_k. By contrast, if different signal realizations result in different messages being transmitted, then the property associated with state x is indefinite. We note that our model directly determines which extensions are admissible in generating the supervaluation, and that, in particular, admissibility is governed by the perception technology Q and the sender's optimal Y-space strategy. Our model also demonstrates a connection between the supervaluation and truth-degrees approaches to vagueness. Whereas the supervaluation approach enumerates the possible interpretations under admissible extensions, truth-degrees describe how likely these interpretations are.

Following the literature on fuzzy sets (see Zadeh, 1975), we define the *support* of message m_k, $S(m_k) = \{x \in X | \, \phi_k(x) > 0\}$, as the set of states that are associated with message m_k with positive probability. Similarly, the *core* of message m_k, $C(m_k) = \{x \in X | \, \phi_k(x) = 1\}$, is the set of states that are definitely associated with message m_k. Naturally $C(m_k) \subseteq S(m_k)$. If $x \in C(m_k)$, then the sender's use will be definite and non-random in state x. If $x \in S(m_k)$ then the receiver understands the meaning of m_k to convey the possibility that the true state is x. Communication will be non-vague if every state is associated with a unique message — i.e. the core and support coincide for every message. By contrast, if messages are vague, there must be some state which is associated with multiple messages, which implies that the supports of at least two messages must overlap. The set of states for which the supports overlap are precisely the 'borderline regions' that characterize vagueness.[19]

For concreteness, we return to the example in the previous section, and characterize its extension to X-space. As before, the full characterization can be found in the Appendix. Let $\varepsilon^* = \frac{3-\sqrt{3}}{8}$. Figure 1 shows the equilibrium truth-degree functions and the support and core sets for the messages in the above example. There are two cases to consider. The left panel illustrates the case when the 'margin for error' is relatively small, $\varepsilon < \varepsilon^*$. The right

[19]Seemingly implicit in the idea that boundary cases characterize vagueness is that there are also some states that are not on the boundary — i.e. there are some messages with non-empty core. If the core of every message is empty, then every state is a boundary case. But then, meaning would be indeterminate in a very different manner than we typically associate with vagueness. Thus, the requirement that some states are associated with multiple messages is necessary, but not sufficient, for vagueness.

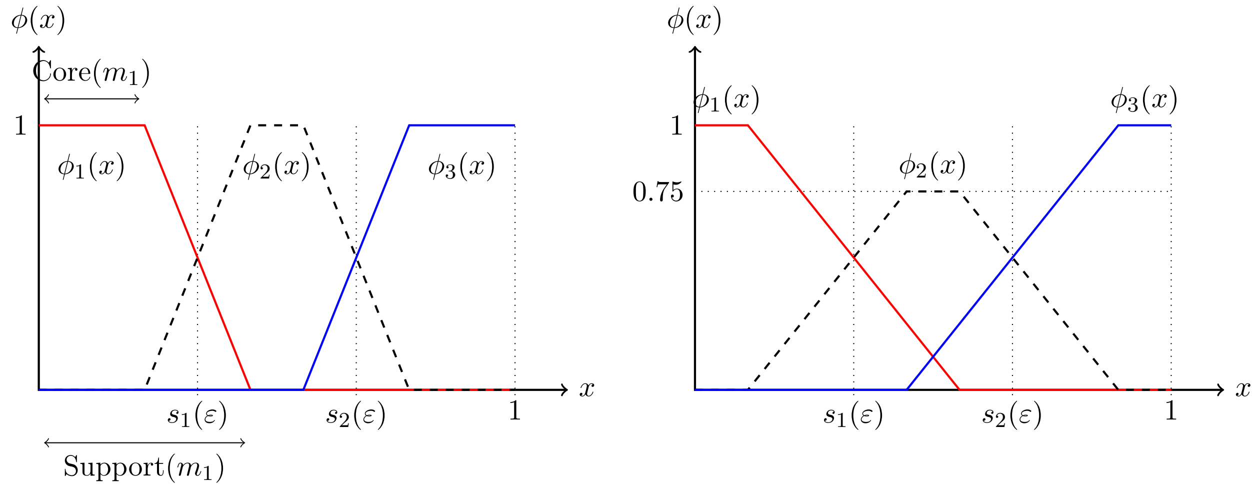

Figure 1: Truth degrees implied by the equilibrium communication strategies in Example 1. The left panel illustrates a situation with a small 'margin for error' ($\varepsilon < \varepsilon^*$). The right panel illustrates the situation with a large 'margin for error' ($\varepsilon \in \left(\varepsilon^*, \frac{3}{8}\right)$).

panel illustrates a scenario with a larger margin for error.

Consider the left panel, with a small 'margin for error'. The left solid (red) line is the truth-degree function $\phi_1(x)$ for message m_1. Similarly, the dashed line and the right solid (blue) line are the truth-degree functions ϕ_2 and ϕ_3 associated with messages m_2 and m_3, respectively. (The linearity of the truth-degree functions is an artifact of the state and signal both being drawn from uniform distributions. Generically, these functions will be 'curved'.) The core and support sets are indicated for message 1. The core of m_1 contains states that are sufficiently small, such that even if the sender receives a (conditionally) above-average signal, this signal will still be low enough that he is guaranteed to transmit m_1. Likewise, the core of m_3 contains states that are sufficiently large that the sender is guaranteed to receive a signal above threshold s_2, and so is guaranteed to transmit m_3. By contrast, there are a range of states for which the sender reports multiple messages with positive probability. Intuitively, such states will be 'close' to a threshold of the Y-space language (in this example, if it lies within $\frac{\varepsilon}{1+2\varepsilon}$ of a threshold), so that the induced signal will lie on one side of the threshold in some instances, and on the other side in other instances. In such regions, truth-degrees are positive for multiple messages. (Consistent with the theory, these truth-degrees must always sum to one, since it is certain the sender will transmit *some* message.) These regions cannot be contained in the core of any message — they are the 'borderline cases' that are characteristic of vagueness.

Next, consider the right panel, where the 'margin for error' is relatively large. This case is distinguished from the previous one in two ways. First, the core of m_2 is empty — there is no state for which we can be certain the sender will transmit m_2. The margin for error is sufficiently large that for any intermediate state there is always the possibility that the sender may occasionally perceive it as being small or large. The core of m_1 and m_3 remain non-empty, although we note that these sets are smaller than in the previous case. By contrast, the supports of all messages are larger; as the margin for error increases, so does the range of states that may be associated with a given message. Second, there exists a range of states for which the sender reports all three messages with positive probability. The margin for error is sufficiently large that, when the true state takes an intermediate value, the sender will sometimes perceive it as being (sufficiently) small and other times perceives it as being (sufficiently) large.

Notice that the set of boundary cases — those states with truth degrees strictly between 0 and 1 — are themselves a function of the perception technology. As this technology becomes more precise (i.e. as ε decreases), the set of boundary cases will narrow as well, and more states will be clearly associated with one message or other.

Finally, as the example makes clear, vagueness is inherent to certain states and not to others. And, this is true even though the perception technology behaved identically across all states. It is not simply the case that meaning is vague in regions of the state space where the perception technology is particular noisy, and clear in regions where the perception technology is more precise. Instead, meaning becomes vague in regions of the state-space for which the generated signals will straddle the optimal Y-space thresholds. Whilst this does depend on the precision of the perception technology, it is also depends in a far more basic sense on the state itself, and its 'location' relative to the threshold. As such, vagueness is necessarily *metaphysical* — it is inherent in the boundary cases themselves.

The preceding discussion also illustrates the role that finiteness (or more generally, the requirement that there be more states than messages) plays in our model. To see this, consider how truth degrees would be different in the example if the communicants had access to a continuum of messages $m \in [0, 1]$. If so, the sender could perfectly reveal her signal to the receiver. However, since every state can be mapped onto multiple signals (given the sender's imperfect perception), it must then follow that every state can be mapped onto multiple messages.[20] The core of every message will be empty. All states will be boundary cases.

[20] Formally, let $\phi(x, y)$ denote the truth density assigned to state x having the y^{th} degree of the relevant property being ascribed. ϕ is a density, rather than a probability mass, in the sense that $\phi(x, y)\Delta y$ is the truth degree associated with state x being assigned a property in the neighborhood Δy of y. Clearly $\phi(x, y) = q(y|x)$, and in our example, $\phi(x, y) = \frac{1}{2\varepsilon}$ for each $y \in \left[\frac{x}{1+2\varepsilon}, \frac{x+2\varepsilon}{1+2\varepsilon}\right]$.

Such dynamics do not sit well with typical accounts of vagueness. After all, there is no disagreement that a man with no hair is bald, nor that a man with a thick head of hair is 'not-bald'. The trouble lies in identifying the boundary between these predicates. Thus vagueness is typically characterized by messages with non-empty cores, but overlapping supports, creating a subset of indeterminate 'boundary cases'. As Example 1 demonstrates, finiteness does not guarantee that every message will have non-empty core. However, as long as the number of messages is small relative to the imprecision in the sender's perception, there will likely be at least some messages with non-empty core.

4.3 Properties of Truth Degree Functions

We now characterize the properties of truth-degree functions. We demonstrate that truth-degrees satisfy three properties. First, they are continuous. Second, they are monotone. Third, they are not truth-functional, but instead satisfy the axioms of probability.

Lemma 1. *Suppose $Q(y|x)$ is continuous in x for every $y \in Y$. Then, for every $k = 1, ..., K$, the truth-degree function $\phi_k(x)$ is continuous in x.*

Loosely speaking, continuity is the property that small changes in the inputs of a function cannot cause dramatic changes in outputs. Continuity of the perception technology formalizes the intuitive assumption that small changes in the underlying state (what is being observed) should not dramatically change the sorts of signals that are generated. For example, if the sender systematically misperceives a person whose true height is 1.7 meters as being much shorter than he actually is, the sender should not then systematically misperceive a slightly taller person as being much taller than is actually the case.[21] Lemma 1 shows that a continuous perception technology causes truth-degrees to be continuous.

An important consequence of continuity is that truth-degree functions respect the desideratum that we treat similarly situated states similarly, *ex ante*. Since the sender cannot easily distinguish between similarly situated states, we should not expect the probability of the sender ascribing property m_k to vary dramatically across those states. And yet this continuity seems to be in stark contrast to the threshold strategy that we derived in the previous section, which necessarily makes stark distinctions between similar objects. In fact, these

[21] To be clear, we are not saying that whenever two states are similar, the sender will perceive them as being similar. As we have argued repeatedly, the *same* state may be perceived differently in different instances. Rather, our claim is that likelihood of the sender (mis)perceiving in some way or other should be similar in the two states.

features are perfectly consistent with one another. It is true that small changes in the signal can dramatically affect which message is transmitted. But, truth degrees are constructed as if by an external agent who observes the true state; and small changes in the state can only generate small changes in the likelihoods of signals that will result in different messages being transmitted. Hence, the external observer's beliefs will change in a gradual fashion, even though the sender's message choice may change starkly in any given instance.

Another consequence of continuity is that use and meaning cannot be characterized by firm thresholds in X-space. Indeed, the fact that use and meaning are optimally characterized by threshold behavior in Y-space precludes the possibility that they respect thresholds in X-space, since the mapping between the two spaces is stochastic. Things which look determinate in Y-space must necessarily seem probabilistic in X-space. As we noted at the beginning of this subsection, it is certainly possible to define a communication strategy over X-space that is characterized by thresholds — however, such a language cannot respect the requirement that use determines meaning, since the sender does not have access to the appropriate information to use messages in the required way. Accordingly, and in contrast to Williamson (1994), we demonstrate that when subject to imperfect perception, optimal communication cannot be characterized by distinct thresholds with respect to statements about actualities. Instead, we show that as the true state increases, there is a gradual and continuous transition in which messages are sent — what Williamson (1994) describes as a 'smear' — which renders meaning vague.

A second property of truth-degree functions is that they are monotone in the ordering over the message-space. Recall, the set of messages was ordered so that m_1 indicated the lowest degree of the gradable adjective and m_K denoted the highest degree. Monotonicity captures the idea that that higher ranked states will be more likely to be described using messages of higher (rather than lower) degree. To make this notion precise, let $\Phi_k(x) = \sum_{j=1}^{k} \phi_j(x)$ denote the truth-degree assigned to state x having property m_k or *lower*. It is easily verified that $\Phi_k(x) = Q(s_k|x)$. We have the following Lemma:

Lemma 2. *Suppose $x_0 < x_1$. Then for every k, $\Phi_k(x_0) \geq \Phi_k(x_1)$.*

Suppose that building A is (actually) taller than building B. Since the sender perceives the world imperfectly, it may be that in some instance, he categorizes B as tall and A as not. However, Lemma 2 demonstrates that this cannot be systematically true. Lemma 2 is a consequence of the monotone likelihood ratio property. This implied that when the true state is high, the sender must be more likely to receiver a higher signal, than when the true state is low. Whilst the sender's classification of objects may be imperfect, it must be statistically consistent with the true grading of objects. Higher graded objects cannot on average be described by lower degrees of the adjective. The assignment of truth-degrees must accord naturally with the use of comparatives.

We stress that the monotonicity property is with respect to cumulative truth-degree functions, rather than individual ones. To make this clear, return to the example of building A which is taller than building B, and suppose the sender can describe these using one of three terms — short, medium and tall. Although 'medium' expresses a higher degree than 'short', it need not be that the sender is more likely to describe building A as medium-heighted than building B. If building A is a sky-scraper, he may be certain to describe it as 'tall', whilst he may well describe building B as medium-heighted in some instances. However, it will be true that the sender is more likely to describe building A as either 'medium heighted' *or* 'tall', than he is to describe building B as such.

Finally, we note that, in strong contrast to most truth-degree proponents, our equilibrium truth-degrees are not truth-functional.[22] A simple example makes this clear. Suppose an object is equally likely to be ascribed each of the properties 'small', 'medium', 'large', and 'enormous', so that the truth-degree of each is $\frac{1}{4}$. Then the truth degrees associated with the ascriptions 'either small or medium-sized', 'either medium-sized or large', and 'either large or enormous' will each be $\frac{1}{2}$. (This follows from the axioms of probability, which the truth-degree functions obey, since they are probability measures by construction.) We have constructed compound statements using the disjunction, and thus far, the truth-degree of the disjunction appears to simply be the sum of the truth-degrees of the disjuncts. However, now consider the ascriptions 'either small, medium-sized or large' and 'either small, medium-sized, large or enormous'. The former is the disjunction of the first and second compound statements above, whilst the latter is the disjunction of the first and third compound statements. If our truth degrees are truth functional, then the truth degrees of these final two sentences must be the same. However, by the laws of probability, the truth-degree of the former is $\frac{3}{4}$, whilst the truth-degree of the latter is 1. Clearly, we cannot universally construct the truth-degrees of compound statements, from the truth-degrees of the constituent statements. (In one of the cases above, we needed some additional information, namely the truth-degree of the conjunction.)

Truth functionality can be a valuable property in a world where truth-degree functions are taken as primitive. Absent truth-functionality, the truth-degree proponent must spec-

[22]Truth-degrees are *truth-functional* if the truth degree of compound sentences can be determined directly from the truth degree of each component sentence.

ify truth-degrees for every conceivable sentence that can be constructed, no matter how long or cumbersome. Such a burden is evidently onerous. Truth-functionality alleviates this need, by reducing all truth-degrees down to the truth-degrees of the underlying simple statements. However, the benefit of truth-functionality is less important in a world where truth-degrees are not primitive, but determined by other known features of the model — in our case, the sender's perception technology and the equilibrium communication strategy. This information (which we used to construct truth-degrees for simple statements in the first place) suffices to construct truth-degrees for any conceivable statement. Truth-functionality provides no additional benefit. Indeed, since the truth-degree functions in our model are probability measures, we can use the laws of probability to map truth-degrees of simple statements onto truth-degrees of compound statements, and vice versa. Thus, the purpose of truth-functionality is preserved. Of course, excepting for special cases, this mapping will not be truth functional, reflecting the idea that the joint distribution of random variables cannot generically be constructed from the marginal distributions alone.

Truth functionality is, of course, not without its own problems. For example, Fine (1975), Williamson (1994) and Edgington (1997) (amongst others, although see Smith (2008) for a defense), note that truth-functionality necessitates that truth-degrees violate standard results in classical logic, including the *Law of the Excluded Middle*. Our truth-degrees-as-probability-measures approach avoids these pitfalls, which provides additional support for this approach to measuring degrees of truth. Although truth functionality may generically be desirable, these benefits vanish in the presence of a well-defined probability measure that can consistently assign truth-degrees.

4.4 Discussion

We conclude our analysis with a brief comparison of the mechanism that generated vagueness in our model from those in other studies. At its core, vagueness arises in our model, despite the sender trying to communicate in a non-vague manner, because an exogenous source of randomness caused the same state to be mapped onto different messages. In our model, we locate that exogenous source of randomness in the imperfect perception of the sender.

Other models locate it elsewhere. Blume and Board (2013*a*) present a model where the sender perceives perfectly, but the receiver is uncertain about the sender's 'language competence'. For example, when communicating about heights, does the sender limit himself to a small vocabulary (e.g. 'short' and 'tall') or a larger one (e.g. 'short', 'medium-height', and ‘tall’). The receiver’s uncertainty about the sender’s vocabulary renders messages vague in equilibrium, even though (as in our model), each type of sender communicates according to clear thresholds. The vagueness arises because there will be a set of boundary cases which the low-vocubulary-type sender would describe as ‘short’ but the high-vocabulary type would describe as ‘medium-height’. Being uncertain about the sender’s type, the receiver associates those objects with both messages.

Other work (see Blume and Board, 2014; Blume, Board and Kawamura, 2007) explore models where the sender perceives perfectly and there is no uncertainty about his language competency, but errors in the transmission technology cause messages to occasionally be rendered incorrectly (or ‘garbled’). Here again, the sender’s communication strategy uses clear thresholds, but the receiver will be uncertain about whether she received the correct message or not. Garbling causes the same state to be associated with multiple different messages, thus generating vagueness.

Under the hood, the mechanism underlying each of these accounts of vagueness is much the same. And yet, the particular details of how the source of randomness is introduced will have different implications for the characteristics of the equilibrium and the properties of truth degrees. Here, we outline a few differences that arise between our model and these variants.

First, consider the language-competency model of Blume and Board (2013*a*), and for concreteness, take the above example of a sender who either uses two or three messages. Since only the high-vocabulary type uses the middle message, its use perfectly reveals the sender’s type. There will be no uncertainty about which states are associated with that message. By contrast, since both types use messages ‘short’ and ‘tall’, but for different intervals of states, these messages will be rendered vague in equilibrium. As noted above, there will be a set of boundary cases, where the low-vocabulary sender uses the message but not the high-type. Notice the contrast to our imperfect perception model (and also the model with garbling), in which, generically, every message will be associated with some boundary cases.

Second, the behavior of the truth degree functions is quite different than under our approach. For example, again consider the language-competency model of Blume and Board (2013*a*), in which the sender may either use two or three messages, and let ρ be the probability that the receiver assigns to the sender being the high-vocabulary type.[23] There are set of states

[23]Following Blume and Board (2013*a*), we assume that the high-vocabulary type is aware that they may be mis-perceived as a low-type, but not vice versa. Keeping the same preferences as in our model, and assuming that $x \sim U[0,1]$, the following is an equilibrium: The low-vocabulary sender transmits ‘short’ if

for which both types of senders will transmit 'short'. These messages will be in the core of the message 'short', and will have associated truth-degree 1. However, there will also be a set of states for which the low-type will transmit the message 'short', while the high type will transmit 'medium-height'. These states are in the support of the message 'short', and all states in this range will have truth degree $1 - \rho$. Truth degrees in the garbling model will behave similarly, except that, generically, the core of every message will be empty (assuming every message may be possible miscommunicated).

Note the difference to our model with imperfect perception, where truth degrees are continuous and gradually decline from 1 (for states near the core) to 0 (for states near the edge of the support). This gradualism is a hallmark of vagueness. As we remove more and more hair from a person's head, we should become increasingly confident in describing them as 'bald'. That dynamic arises in our model. By contrast, under the alternate approach, the receiver will be equally uncertain about whether to describe the person as bald or not, over a range of cases. The reason for this difference lies in how the exogenous randomness operates. Under the imperfect perception approach, this randomness is a function of the state, which causes truth degrees to vary across states. By contrast, under the other approaches, the randomness is statistically unrelated to the underlying state.

Similar differences arise in other approaches. For example, consider the 'Intentional Vagueness' model Blume and Board (2014), whose setup includes a common-interest game as a special case.[24] That model departs from ours in other ways; for example it inverts the structure between states and messages. In the 'Intentional Vagueness' model, there are only two states but a continuum of messages. By contrast, in the typical Sorites paradox setup, there are a large number (approximated by a continuum) of states and only two messages. States on the extremes are easily identified with one message or other; however there is a blurring of the boundaries between the messages, so that a subset of intermediate states are occasionally associated with either message. Such a dynamic does not arise in the 'Intentional Vagueness' setup. With only two states and garbling, there are no states that are identified with a single message; all states are boundary states. Moreover, we cannot generate a Sorites series by

$x \in [0, \frac{1}{2}]$ and 'tall' otherwise. The high-vocabulary type transmits 'short' if $x \in [0, \frac{3-\rho}{8-2\rho}]$, 'medium-height' if $x \in (\frac{3-\rho}{8-2\rho}, \frac{5-\rho}{8-2\rho}]$, and 'tall' if $x \in (\frac{5-\rho}{8-2\rho}, 1]$. The receiver's actions are $(a_L, a_M, a_H) = (\frac{2-\rho}{8-2\rho}, \frac{1}{2}, \frac{6-\rho}{8-2\rho})$. For this equilibrium, truth degrees for the message 'short' are: $\phi_1(x) = 1$ for $x \in [0, \frac{3-\rho}{8-2\rho}]$, $\phi_1(x) = 1 - \rho$ for $x \in [\frac{3-\rho}{8-2\rho}, \frac{1}{2}]$, and $\phi_1(x) = 0$ for $x \in [\frac{1}{2}, 1]$. Truth degrees for the remaining messages are computed similarly.

[24] In that model, the sender perfectly observes a binary state $x \in \{0, 1\}$, and transmits a message $m \in [0, 1]$ to the receiver. (In equilibrium, the sender will transmit $m_0 = 0$ in state 0 and $m_1 = 1$ in state 1.) The message becomes garbled during transmission, so that the receiver observes message $q \sim N(m, \sigma^2)$. The receiver then takes an action a. As in our model, both agents have common quadratic loss preferences, and so the receiver's equilibrium action is simply her posterior belief that the state is 1.

gradually increasing the state and asking what message would be attached to it. We *could* gradually increase the message and ask which state likely generated that message, but that provides the answer to the question: 'how does the posterior belief change as the message increases?', which is different from the question 'how does the likelihood of a state being associated with a particular predicate change as the state increases?'

Of course, there are apparent similarities between the 'Intentional Vagueness' model and our own. For example, the receiver's posterior belief a looks like a 'smear', increasing from close to 0 (when the receiver's message is low) to almost 1 (when the receiver's message is high).[25] However, the receiver's posterior belief is not the truth degree function for the proposition 'the state is 1'. Nor was it in our model. Instead the truth degree (at least as perceived by the receiver) is simply given by the density of the $N(1, \sigma^2)$ distribution — which does not behave like a smear.

The construction of truth degrees themselves becomes problematic in the 'Intentional Vagueness' framework. In our model, there is an 'objective' mapping between states and messages that defines truth-degrees. By contrast, in 'Intentional Vagueness', due to the garbling, the mapping between states and messages is different for the sender and receiver. If we were to associate this mapping with truth degrees, then truth degrees would be degenerate for the sender (who associates a single message with each state) but non-degenerate for the receiver. But truth degrees are not usually taken to be 'subjective' in this way — they are properties of messages/statements and not contingent on the speaker/listener. Of course, the fact that different speakers may understand the same message differently may explain why a proposition has a truth degree strictly between zero and one; but that is a different thing from the proposition having speaker-depending truth degrees.

5 Conclusion

This paper examined imperfect perception as a source of vagueness. We developed a model of communication in which an imperfectly informed sender may transmit a message to an uninformed receiver, who must take an action that affects both parties. Both agents share identical concave preferences, which incentivizes complete information transfer between the parties. To focus attention most cleanly on the effect of imperfect perception, we abstract

[25] Blume & Board endow the garbling technology with unbounded support, so any message can by generated by either state. With bounded support, the receiver would be able to perfectly identify certain messages with only one state or other, and so posterior beliefs would be exactly zero or one in some instances.

from other features that may independently cause messages to be vague. Although our framework is stylized, we are able to shed light on the several properties of vagueness.

Our analysis begins with the recognition that, in a world with imperfect perception, we must distinguish two sorts of statements — subjective statements which convey what the sender perceives, and objective statements which convey what actually is. Since the sender only observes signals about the world — and not the actual world itself — he can only transmit subjective statements, and his messages ought to be interpreted as such. Nevertheless, we can attempt to imply meaning about the actual world from the sender's message about the perceived world.

Our analysis also recognizes that a sender's optimal choice about which message to transmit will depend on his belief about how the receiver will interpret messages, and that the receiver will optimally interpret messages according to her expectation of when and how the sender transmits each available message. *Use* and *meaning* are, as such, jointly determined in equilibrium, given the agents' communication needs and the sender's perception technology.

We first characterized an optimal equilibrium when the communication is understood to be about what the agent perceives. We showed that, notwithstanding the sender's imperfect information, communication in this world is non-vague and characterized by firm thresholds that demark the use of words. Intuitively, although the sender understands that his perception of the world may not be accurate, this does not prevent him from clearly communicating what he has perceived. Moreover, since the receiver can rationally understand the sender's communication strategy, she can clearly infer what the sender has perceived. Hence, communication is not vague.

We then consider how to assign meaning to statements, if they are to be interpreted as being about actualities. (As we argue, the sender must continue to transmit messages based on what he perceives.) Given the sender's imperfect perception, he may in different instances ascribe different properties to the same state of the world. Although, the sender's use is determined in any given instance by his signal, his use may appear random or probabilistic to an external agent who observes the state but not the signal. We interpret these probabilities as truth-degrees, since they capture the likelihood that the sender will ascribe a particular property to a given state, *ex ante*. As such, we provided micro-foundations for truth-degrees as a consequence of the equilibrium communication (about what is perceived) and the sender's perception technology, thus connecting two distinct theories of vagueness. Indeed, epistemic theory provided the mechanism that enabled us to generate the descriptive features of the truth-degrees approach.

A useful feature of our approach is that, since truth-degrees are determined (in equilibrium) rather than assumed, we can investigate the properties that truth-degree functions are likely to satisfy. We derived three features of truth-degree functions that were predicated upon standard assumptions about the perception technology. First, we showed that truth-degrees are continuous, capturing the natural idea that senders will, on average, similarly describe similarly situated states. We showed that continuity of truth-degrees and gradualism was a feature that naturally arose in our imperfect perception framework, but was unlikely to obtain under other mechanisms. An important consequence of continuity is that communication about the objective world cannot simultaneously satisfy the use-determines-meaning criterion and be characterized by thresholds. This result stands in strong contrast to epistemic theory, which insists that the underlying language is not inherently vague. Additionally, we showed that truth-degree functions are monotone, and therefore accord naturally with the use of comparatives. Finally, we demonstrated that our induced truth-degree functions were not truth functional. However, we argued that this was not problematic, since truth-degrees of compound statements could still be discerned from the truth-degrees of simpler statements, given the laws of probability. Hence, the essential benefit of truth-functionality is preserved. Moreover, we showed that truth-degrees-as-probabilities avoid some of the more problematic features inherent to truth-functional truth-degrees, such as inconsistency with classical results in logic, such as the law of the excluded middle.

References

Blume, Andreas. 2021. "Meaning in Communication Games." Available at SSRN: https://ssrn.com/abstract=3893109.

Blume, Andreas and Oliver Board. 2013*a*. "Language barriers." *Econometrica* 81(2):781–812.

Blume, Andreas and Oliver Board. 2014. "Intentional vagueness." *Erkenntnis* 79(4):855–899.

Blume, Andreas and Oliver J Board. 2013*b*. "Higher-Order Uncertainty About Language.".

Blume, Andreas, Oliver J Board and Kohei Kawamura. 2007. "Noisy talk." *Theoretical Economics* 2(4):395–440.

Crawford, Vincent P. 2003. "Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions." *American Economic Review* 93(1):133–149.

Crawford, Vincent P and Joel Sobel. 1982. "Strategic information transmission." *Econometrica: Journal of the Econometric Society* pp. 1431–1451.

Dummett, Michael. 1975. "Wang's paradox." *Synthese* 30(3-4):301–324.

Edgington, Barbara. 1997. Truth by Degrees. In *Vagueness: A Reader*, ed. Rosanna Keefe and Peter Smith. Bradford.

Fine, Kit. 1975. "Vagueness, truth and logic." *Synthese* 30(3):265–300.

Franke, Michael. 2014. "Pragmatic reasoning about unawareness." *Erkenntnis* 79:729–767.

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

Halldén, Sören. 1949. "The logic of nonsense.".

Hampton, James A and Martin L Jönsson. 2012. "Typicality and composition a lity: The logic of combining vague concepts.".

Jäger, Gerhard, Lars P Metzger and Frank Riedel. 2011. "Voronoi languages: Equilibria in cheap-talk games with high-dimensional types and few signals." *Games and economic behavior* 73(2):517–537.

Keefe, Rosanna. 2000. *Theories of vagueness.* Cambridge University Press.

Körner, Stephan. 1962. "The Philosophy of Mathematics, an Introduction.".

Lambie-Hanson, Timothy and Giri Parameswaran. 2016. "A Model of Vague Communication.".

Lassiter, Daniel and Noah D Goodman. 2014. Context, scale structure, and statistics in the interpretation of positive-form adjectives. In *Semantics and Linguistic Theory.* pp. 587–610.

Lewis, David. 1969. *Convention: a philosophical study*. Cambridge, Harvard university press.

Lipman, Barton L. 2003. Language and economics. In *Cognitive Processes and Economic Behaviour*, ed. N. Dimitri, M. Basili and I. Gilboa. Routledge.

Lipman, Barton L. 2009. "Why is language vague?" *Boston University* .

Milgrom, Paul R. 1981. "Good news and bad news: Representation theorems and applications." *The Bell Journal of Economics* pp. 380–391.

Murphy, Gregory. 2004. *The big book of concepts.* MIT press.

Osborne, Martin J et al. 2004. *An introduction to game theory.* Vol. 3 Oxford university press New York.

Parikh, Prashant. 2019. *Communication and content.* Language Science Press.

Qing, Ciyang and Michael Franke. 2014. Gradable adjectives, vagueness, and optimal language use: A speaker-oriented model. In *Semantics and Linguistic Theory.* Vol. 24 pp. 23–41.

Rick, Armin. 2013. "The Benefits of Miscommunication.".

Schiffer, Stephen. 2000. "Vagueness and partial belief." *Philosophical Issues* 10:220–257.

Schiffer, Stephen. 2003. *The things we mean.* Clarendon Press.

Schiffer, Stephen. 2010. Vague Properties. In *Cuts & Clouds: Vagueness, its Nature, and its Logic*, ed. Richard Dietz and Sebastiano Moruzzi. Oxford University Press pp. 109–130.

Smith, Nicholas JJ. 2008. *Vagueness and degrees of truth.* Oxford University Press.

Sorensen, Roy. 2001. *Vagueness and contradiction.* Oxford University Press.

Williamson, Timothy. 1994. *Vagueness.* Psychology Press.

Williamson, Timothy. 2002. *Vagueness.* Routledge.

Zadeh, Lotfi A. 1975. "Fuzzy logic and approximate reasoning." *Synthese* 30(3-4):407–428.

Appendices

Proof of Proposition 1. We prove the proposition in three stages. First, we show that an equilibrium exists in threshold strategies. Second, we show that any optimal equilibrium must utilize all messages, and thus be in threshold strategies. Third, we show that the equilibrium is unique under certain conditions.

Existence. The proof of existence is itself in several parts. We first conjecture a particular threshold strategy for the sender. Taking this strategy as given, we compute the receiver's optimal action for each message (which requires that we first compute the receiver's posterior belief about the true state given the transmitted message). Having computed the receiver's strategy, we then find the sender's best response. The strategy profile is an equilibrium if the sender's best response to the receiver coincides with the original conjecture of the sender's strategy (which informed the receiver's strategy). Using a fixed point argument, we show that there exists a threshold strategy for the sender that is equilibrium consistent.

Let $\Sigma = \left\{ s \in Y^{K+1} | s = (s_0, ..., s_K) \text{ with } 0 = s_0 \leq ... \leq s_K = 1 \right\}$. Take some $s \in \Sigma$. Suppose the sender uses the strategy: $\mu(y) = m_k$ provided that $y \in (s_{k-1}, s_k]$.

Step 1: The agents' beliefs. First, let us compute the players' beliefs about the state. Using Bayes' Rule:

$$f(x|y) = \frac{f(x)\, q(y|x)}{\int_0^1 f(z)\, q(y|z)\, dz}$$

and:

$$g(x|m_k) = \frac{f(x) \left[\int_{s_{k-1}}^{s_k} q(y|x)\, dy \right]}{\int_0^1 f(z) \left[\int_{s_{k-1}}^{s_k} q(y|z)\, dy \right] dz}$$

provided that $s_{k-1} < s_k$. If $s_{k-1} = s_k$, we set $g(x|m_k) = f(x|s_k)$. (I.e. we assume that the signal must have been $y = s_k$ if the out of equilibrium message m_k is ever transmitted.) Note that since $\int_{s_{k-1}}^{s_k} q(y|x)\, dy$ is continuous in the sender's communication strategy s, so is $g(x, m_k)$.

Two properties of the belief functions will prove useful. First, by the monotone likelihood ratio property, the sender's posterior beliefs respect first order stochastic dominance. I.e. $F(x|y_1) \leq F(x|y_0)$ whenever $y_1 > y_0$. (We show this in the proof of Lemma 2, below.) Second, and relatedly, the receiver's posterior beliefs also respect first order stochastic dominance. I.e. $G(x|m_k) \leq G(x|m_{k'})$ whenever $k > k'$.

To see this latter property, notice that:

$$
\begin{aligned}
G(x|m_k) &= \frac{\int_0^x f(z)\left(\int_{s_{k-1}}^{s_k} q(y|z)dy\right)dz}{\int_0^1 f(z)\left(\int_{s_{k-1}}^{s_k} q(\gamma|z)d\gamma\right)dz} \\
&= \int_{s_{k-1}}^{s_k}\left(\frac{\int_0^1 f(z)q(y|z)dz}{\int_{s_{k-1}}^{s_k}\left(\int_0^1 f(z)q(\gamma|z)dz\right)d\gamma}\right)\cdot\frac{\int_0^x f(z)q(y|z)dz}{\int_0^1 f(z)q(y|z)dz}dy \\
&= \int_{s_{k-1}}^{s_k} h(y|m_k)\cdot F(x|y)dy
\end{aligned}
$$

where $h(y|m_k) = \frac{\int_0^1 f(z)q(y|z)dz}{\int_{s_{k-1}}^{s_k}\left(\int_0^1 f(z)q(\gamma|z)dz\right)d\gamma}$ is the (conditional) density over the signals that induce the sender to transmit message m_k. By construction $h(y|m_k) > 0$ for all $y \in [s_{k-1}, s_k]$ and $\int_{s_{k-1}}^{s_k} h(y|m_k) = 1$. Now, since $k' < k$, we have $s_{k'} \le s_{k-1}$. Then, since $F(x|y)$ respects first order stochastic dominance, we have:

$$
G(x|m_k) = \int_{s_{k-1}}^{s_k} h(y|m_k)F(x|y)dy \le F(x|s_{k-1}) \le F(x|s_{k'}) \le \int_{s_{k'-1}}^{s_{k'}} h(y|m_k)F(x|y)dy = G(x|m_{k'})
$$

Step 2: The receiver's strategy. Next, we compute the receiver's optimal strategy, given the beliefs induced by the sender's conjectured strategy. The receiver's expected utility from choosing action $a \in [0,1]$ after receiving message m_k is $-\int_{x\in X}(x-a)^2 g(x|m_k)\,dx$. The sender chooses a_k to maximize her expected utility. We can compute this maximizer by taking first-order conditions:

$$
-2\int_0^1 (x-a)\,g(x|m_k)\,dx = 0
$$

$$
a_k^* = \int_0^1 xg(x|m_k)\,dx
$$

since $\int g(x|m_k)\,dx = 1$. For each $k = 1,..,K$, let $A_k(s) = a_k^*$, and note that $A_k(s)$ is continuous in s. Let $A(s) = (A_1(s),...,A_K(s))$. Additionally, since the receiver's belief functions respect first order stochastic dominance, it must be that: $A_1(s) \le ... \le A_K(s)$.

To see this, note that, suppose $k' < k$. Then:

$$\begin{aligned} a_k^* = \int_0^1 xg(x|m_k)dx = [xG(x|m_k)]_0^1 - \int_0^1 G(x|m_k)dx \\ = 1 - \int_0^1 G(x|m_k)dx \\ \geq 1 - \int_0^1 G(x|m_{k'})dx \\ = \int_0^1 xg(x|m_{k'})dx = a_{k'}^* \end{aligned}$$

Step 3: The sender's strategy. Next, we must verify that the sender's strategy is optimal, given the receiver's action profile and the sender's beliefs. The sender chooses the message that maximizes his *ex ante* utility, anticipating the receiver's action, and given his beliefs about the true state. We have: $\mu(y) = \arg\max_{m_k \in M} \left\{-\int_{x \in X} (x - A_k(s))^2 f(x|y)\, dx\right\}$.

Fix some $y \in [0,1]$, and suppose m_k is an optimal message. For any k' s.t. $A_{k'}(s) = A_k(s)$, it must that $m_{k'}$ is also an optimal message. Suppose there is a $k' < k$ s.t. $A_{k'}(s) < A_k(s)$. Then, since m_k is optimal, we must have:

$$\begin{aligned} \int_{x \in X} (x - A_k(s))^2 f(x|y)\, dx \;&\leq\; \int_{x \in X} (x - A_{k'}(s))^2 f(x|y)\, dx \\ \int x f(x|y)\, dx \;&\geq\; \frac{1}{2}(A_k(s) + A_{k'}(s)) \end{aligned}$$

Since this must be true for every such $k' < k$, and $A_1 \leq ... \leq A_K$, we have: $\int x f(x|y)\, dx \geq \frac{1}{2}(A_k(s) + A_{k^-}(s))$, where $k^- = \max\{k'|A_{k'} < A_k\}$ (and $k^- = k-1$ if $A_{k-1} < A_k$). By a similar argument, it must be that:

$$\int x f(x|y)\, dx \leq \frac{1}{2}(A_k(s) + A_{k^+}(s))$$

where $k^+ = \min\{k'|A_{k'} > A_k\}$ (and $k^+ = k+1$ if $A_k < A_{k+1}$). Let $\psi(y) = \int x f(x|y)\, dx$ denote the sender's assessment of the expected state, given signal y. We have $\frac{1}{2}(A_k(s) + A_{k^-}(s)) \leq \psi(y) \leq \frac{1}{2}(A_k(s) + A_{k^+}(s))$.

By the continuity in the setup, it must be that $\psi(y)$ is continuous in y. Additionally, the fact that $F(x|y)$ respects first-order stochastic dominance implies that $\psi(y)$ is strictly increasing. Hence it is a best response to report m_k if $y \in \left[\psi^{-1}\left(\frac{1}{2}(A_k(s) + A_{k^-}(s))\right), \psi^{-1}\left(\frac{1}{2}(A_k(s) + A_{k^+}(s))\right)\right]$.

Notice that if $A_1 < A_2 < \cdots < A_K$, then the best response intervals are non-overlapping; the sender's optimal strategy partitions the signal space. If $A_k = A_{k'}$, then the intervals for which m_k and $m_{k'}$ are best responses coincide. Each signal in the common interval could be assigned to either message, or the sender could randomize between the message. To this extent, the sender's strategy need not partition the signal space.

However, the sender is free to select an assignment of signals to messages that does partition the signal space. In particular, let $S_k(s) = \psi^{-1}\left(\frac{1}{2}\left(A_k(s) + A_{k+1}(s)\right)\right)$ for $k = 1, ..., K-1$, and let $S_0(s) = 0$ and $S_K(s) = 1$. Then $S_0 \leq S_1 \leq \cdots \leq S_k$. Note that S_k is continuous in s. Given the preceding analysis, it must be that message m_k is optimal if $y \in [S_{k-1}(s), S_k(s)]$. Let $S(s) = (S_0(s), .., S_K(s))$.

Step 4: Consistency. We need to show that we can conjecture a threshold strategy s that is consistent with optimal behavior by the sender. I.e. we need to show that there exists an s such that $s = S(s)$ — i.e. that s is a fixed point of the mapping $S : \Sigma \to \Sigma$. Since this mapping is continuous over a compact space, Brouwer's fixed point theorem ensures that it admits a fixed point. Hence, an equilibrium exists.

Finally, we must show that $s_{k-1} < s_k$ in any equilibrium with fixed point strategies. (This guarantees that each message is transmitted with positive probability.) Suppose not. I.e. suppose there is a fixed point of the mapping for which $s_{k-1} = s_k$. Then $A_k(s) = \psi(s_k)$. Next note that

$$\int_{x \in X} (x-a)^2 f(x|y)\, dx = (\psi(y) - a)^2 + \int (x - \psi(y))^2 f(x|y)\, dx$$

which is simply the usual mean-square error decomposition, since $\psi(y) = E_{X|Y=y}[X]$. For message k with $a = A_k(s)$, this expression is minimized when $y = s_k$. Moreover, for any k' s.t. $A_{k'}(s) \neq A_k(s)$, this expression is strictly larger when $\mu = m_{k'}$. Then by continuity m_k must be optimal for $y \in (s_{k-1}, s_{k-1} + \varepsilon)$ where $\varepsilon > 0$ is small enough. Hence $s_k > s_{k-1}$, and so $A_k(s) > A_{k-1}(s)$ for each k.

Optimal Equilibria. Next, we show that any optimal equilibrium must induce the agent to take K distinct actions. To do so, take an equilibrium (μ, α), not necessarily in threshold strategies, that induces $L < K$ distinct actions $\{a_1, \ldots, a_L\}$. Let a_L denote the highest action taken in such an equilibrium. By the argument in steps 2 and 3 above, it must be that $a_L < 1$. Additionally, by step 3, it must be that $\mu(y) = m_L$ whenever $y > \psi(a_L)$.

It suffices to find a different set of strategies (not necessarily an equilibrium) that induces $L+1$ distinct actions, and which achieves a lower *ex ante* expected loss than the equilibrium (μ, α). Given the common values framework, if a set of such strategies exists, then there must be an equilibrium that induces $L+1$ distinct actions that does better as well. Consider the strategies: (μ', α') where: (i) the induced actions are $\{a_1, \ldots, a_{L+1}\}$ with $a'_k = a_k$ for $k = 1, \ldots, L$ and $a'_{L+1} = 1$, and (ii) $\mu'(y) = \mu(y)$ whenever $y \leq \psi^{-1}\left(\frac{a_L+1}{2}\right)$ and $\mu'(y) = m_{L+1}$ whenever $y > \psi^{-1}\left(\frac{a_L+1}{2}\right)$. Recall that $\ell(\mu, \alpha)$ is the *ex ante* expected loss under strategy (μ, α). Then:

$$\begin{aligned}
\ell(\mu', \alpha') &= \ell(\mu, \alpha) + \int_0^1 \int_{\psi^{-1}\left(\frac{a_L+1}{2}\right)}^1 \left[(x-1)^2 - (x-a_L)^2\right] f(x|y) q(y) dy dx \\
&= \ell(\mu, \alpha) - (1-a_L) \int_{\psi^{-1}\left(\frac{a_L+1}{2}\right)}^1 \left[\int_0^1 \left[x - \frac{a_L+1}{2}\right] f(x|y) dx\right] \cdot q(y) dy \\
&= \ell(\mu, \alpha) - (1-a_L) \int_{\psi^{-1}\left(\frac{a_L+1}{2}\right)}^1 \left[\psi(y) - \frac{a_L+2}{2}\right] q(y) dy \\
&< \ell(\mu, \alpha)
\end{aligned}$$

where the final inequality makes use of the fact that $\psi(y) > \frac{a_L+1}{2}$ whenever $y > \psi^{-1}\left(\frac{a_L+1}{2}\right)$. Then since $\ell(\mu, \alpha) < \ell(\mu', \alpha') < 0$, the expected loss under (μ', α') is smaller.

In Step 4, we showed that an equilibrium exists in which K distinct actions are induced, and that such an equilibrium will be in threshold strategies. (Since the sender chooses from a finite set of K messages, it cannot be that more than K actions are induced in equilibrium.) By the preceding analysis, we showed that an optimal equilibrium induces K distinct actions. Hence, an optimal equilibrium exists, and it is in threshold strategies.

Uniqueness. The proof of uniqueness is adapted from Crawford and Sobel (1982). Let $q(y) = \int_0^1 q(y|x) f(x) dx$ denote the unconditional density of the signal y. We can verify that $g(x|m_k) = \int_{s_{k-1}}^{s_k} \frac{q(y)}{\int_{s_{k-1}}^{s_k} q(\gamma) d\gamma} f(x|y) dy$. In equilibrium, we know that the receiver will choose a_k to satisfy:

$$\begin{aligned}
a_k = \int_0^1 x g(x|m_k) dx &= \int_{s_{k-1}}^{s_k} \frac{q(y)}{\int_{s_{k-1}}^{s_k} q(\gamma) d\gamma} \left(\int_0^1 x f(x|y) dx\right) dy \\
&= E\left[y | y \in (s_{k-1}, s_k)\right] \\
&= \frac{s_{k-1} + s_k}{2} \qquad (1)
\end{aligned}$$

where the second line uses the assumption that $\int_0^1 xf(x|y)dx = E[x|y] = y$, and the third line uses the assumption that the unconditional distribution of y is uniform. Similarly, we know that the sender will choose s_k to satisfy: $\int_0^1 (a_k - x)^2 f(x|s_k)dx = \int_0^1 (a_{k+1} - x)^2 f(x|s_k)dx$. Since $a_{k+1} \neq a_k$, this simplifies to

$$s_k = \frac{a_k + a_{k+1}}{2} \tag{2}$$

where again we use the assumption that $E[x|s_k] = s_k$.

Let $s_0 = 0$ be given and specify some $a_1 > s_0$. Let $(s_0, \ldots, s_K)$ and $(a_1, \ldots, a_K)$ be sequences that satisfy (1) and (2). Given s_{k-1}, (1) determines s_k as a function of a_k, i.e. $s_k = 2a_k - s_{k-1}$; and given a_k, (2) determines a_{k+1} as a function of s_k, i.e. $a_{k+1} = 2s_k - a_k$. This implies that:

$$s_{k+1} = 2a_{k+1} - s_k = 2[2s_k - a_k] - s_k = s_k + 2(s_k - a_k)$$

Using this fact and (2), we have $s_{k+1} - a_{k+1} = s_k - a_k$, and so by induction $s_k - a_k = s_1 - a_1$. Also, by (1) and since $s_0 = 0$, $s_1 = 2a_1$. Hence, by induction, we have $s_k = 2ka_1$. Notice that the s_k's are monotonically increasing in a_1. (Crawford and Sobel refer to this as Condition (M).) Hence the choice of a_1 pins down all the other terms. Additionally, since $s_K = 1$, we must have $a_1 = \frac{1}{2K}$. But this implies that the communication game admits a unique equilibrium. □

Example in Detail. We construct the example in four steps. First, we determine the agents' belief functions. Second, we determine the receiver's optimal action, given her beliefs for an arbitrary strategy by the sender. Third, we determine the sender's optimal strategy given her beliefs, anticipating the receiver's optimal response. These three steps together characterize the equilibrium. Additionally, we construct the truth degree functions.

Step 1: The agents' beliefs. Suppose $x \sim U[0,1]$ and let the perception technology generate a signal $y \sim U\left[\frac{x}{1+2\varepsilon}, \frac{x+2\varepsilon}{1+2\varepsilon}\right]$, where $\varepsilon < \frac{3}{8}$ denotes the signal precision. Notice that this implies $y \in [0,1]$.

Start with the sender's posterior belief about the true state after receiving signal y. By Bayes' rule, we have the following: If $y < \frac{2\varepsilon}{1+2\varepsilon}$, then

$$f(x|y) = \begin{cases} \frac{1}{(1+2\varepsilon)y} & \text{if } x \in [0, (1+2\varepsilon)y] \\ 0 & \text{otherwise.} \end{cases}$$

If instead $y \in \left[\frac{2\varepsilon}{1+2\varepsilon}, \frac{1}{1+2\varepsilon}\right]$, then

$$f(x|y) = \begin{cases} \frac{1}{2\varepsilon} & \text{if } x \in [(1+2\varepsilon)y - 2\varepsilon, (1+2\varepsilon)y] \\ 0 & \text{otherwise.} \end{cases}$$

and finally if $y > \frac{1}{1+2\varepsilon}$, then

$$f(x|y) = \begin{cases} \frac{1}{(1+2\varepsilon)(1-y)} & \text{if } x \in [(1-2\varepsilon)y - 2\varepsilon, 1] \\ 0 & \text{otherwise.} \end{cases}$$

Now turn to the receiver's beliefs after receiving a given message. Suppose $K = 3$, so that 3 messages are available, and let s_1 and s_2 be the thresholds that delineate the use of these messages. We use Bayes' Rule to compute the receiver's beliefs after receiving message $i \in \{1, 2, 3\}$. For concreteness, consider her beliefs after receiving message m_1 (which implies that the signal y was in $[0, s_1]$). By Bayes' Rule, we have:

$$g(x|m_1) = \frac{\int_0^{s_1} q(y|x)\, dy}{\int_{z \in X} \left[\int_0^{s_1} q(y|z)\, dy\right] dz}$$

Then, using the fact that $q(y|x) = \frac{1}{2\varepsilon}$ for $y \in \left[\frac{x}{1+2\varepsilon}, \frac{x+2\varepsilon}{1+2\varepsilon}\right]$ (and 0 otherwise), we have:

$$\int_0^{s_1} q(y|x) dy = \begin{cases} 1 & \text{if } x < (1+2\varepsilon)s_1 - 2\varepsilon \\ \frac{(1+2\varepsilon)s_1 - x}{2\varepsilon} & \text{if } x \in [(1+2\varepsilon)s_1 - 2\varepsilon, (1+2\varepsilon)s_1] \\ 0 & \text{if } x > (1+2\varepsilon)s_1 \end{cases}$$

Assume that $s_1 > \frac{2\varepsilon}{1+2\varepsilon}$. (Since $\varepsilon < \frac{3}{8}$ and $s_1 < \frac{1}{2}$, we know that $(1+2\varepsilon)s_1 < 1$.) This implies that:

$$\begin{aligned} \int_{z \in X} \left[\int_0^{s_1} q(y|z)\, dy\right] dz &= \int_0^{(1+2\varepsilon)s_1 - 2\varepsilon} dx + \int_{(1+2\varepsilon)s_1 - 2\varepsilon}^{(1+2\varepsilon)s_1} \left(\frac{(1+2\varepsilon)s_1 - x}{2\varepsilon}\right) dx \\ &= (1+2\varepsilon)s_1 - \varepsilon \end{aligned}$$

and so, we have:

$$g\left(x|m_1\right)=\begin{cases}\frac{1}{(1+2\varepsilon)s_1-\varepsilon} & x\in[0,(1+2\varepsilon)s_1-2\varepsilon]\\ \frac{(1+2\varepsilon)s_1-x}{2\varepsilon[(1+2\varepsilon)s_1-\varepsilon]} & x\in[(1+2\varepsilon)s_1-2\varepsilon,(1+2\varepsilon)s_1]\\ 0 & x\in[(1+2\varepsilon)s_1,1]\end{cases}$$

We can construct $g\left(x|m_2\right)$ and $g\left(x|m_3\right)$ similarly.

Step 2: The receiver's optimal actions. Let a_i be the action chosen by the receiver following the receipt of message m_i. Since the problem is perfectly symmetric, we know that $s_2 = 1 - s_1$, that $a_3 = 1 - a_1$ and $a_2 = \frac{1}{2}$. Hence, it suffices to characterize a_1 as a function of s_1.

Following message m_1, the receiver chooses the action which corresponds to the expected true state conditional upon the message received. We have:

$$\begin{aligned}a_1(s_1) &= \int_X xg\left(x|m_1\right)dx\\ &= \int_0^{(1+2\varepsilon)s_1-2\varepsilon}\frac{x}{(1+2\varepsilon)s_1-\varepsilon}dx+\int_{(1+2\varepsilon)s_1-\varepsilon}^{(1+2\varepsilon)s_1}x\frac{(1+2\varepsilon)s_1-x}{2\varepsilon[(1+2\varepsilon)s_1-\varepsilon]}dx\\ &= \frac{3[(1+2\varepsilon)s_1-\varepsilon]^2+\varepsilon^2}{6[(1+2\varepsilon)s_1-\varepsilon]}\end{aligned}$$

Step 3: The sender's optimal messages. Let $\psi\left(y\right)=\int xf\left(x|y\right)dx$, which is the sender's assessment of the expected state, after receiving signal y. Take $y\in\left[\frac{2\varepsilon}{1+2\varepsilon},\frac{1}{1+2\varepsilon}\right]$. Then we know that:

$$\psi\left(y\right)=\int_{(1+2\varepsilon)y-2\varepsilon}^{(1+2\varepsilon)y}\frac{x}{2\varepsilon}dx=(1+2\varepsilon)y-\varepsilon$$

By construction, the sender must be indifferent between sending either message 1 or message 2 after receiving signal s_1. Hence, s_1 satisfies $\psi\left(s_1\right)=\frac{1}{2}\left(a_1+a_2\right)$. Then, assuming $s_1\in\left[\frac{2\varepsilon}{1+2\varepsilon},\frac{1}{1+2\varepsilon}\right]$, we have:

$$\begin{aligned}(1+2\varepsilon)s_1-\varepsilon &= \frac{1}{2}\left[\frac{3[(1+2\varepsilon)s_1-\varepsilon]^2+\varepsilon^2}{6[(1+2\varepsilon)s_1-\varepsilon]}+\frac{1}{2}\right]\\ s_1 &= \frac{(1+6\varepsilon)+\sqrt{1+4\varepsilon^2}}{6(1+2\varepsilon)}=\frac{1}{6}+\frac{4\varepsilon+\sqrt{1+4\varepsilon^2}}{6(1+2\varepsilon)}\end{aligned}$$

We are left to confirm that s_1 so defined indeed satisfies $s_1\in\left[\frac{2\varepsilon}{1+2\varepsilon},\frac{1}{1+2\varepsilon}\right]$. It is easily confirmed that this will be the case provided that $\varepsilon < \frac{3}{8}$.

In summary, we have: $(s_1^*, s_2^*) = \left(\frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}, \frac{5+6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}\right)$ and $(a_1^*, a_2^*, a_3^*) = \left(\frac{\sqrt{1+4\varepsilon^2}}{3} - \frac{1}{6}, \frac{1}{2}, \frac{7}{6} - \frac{\sqrt{1+4\varepsilon^2}}{3}\right)$.

Step 4: Characterizing the truth-degree functions Finally, we characterize the core and support of each message, and the associated truth-degree functions. The support sets are:

$$
\begin{aligned}
S_R(m_1) &= \left[0, \frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}\right] \\
S_R(m_2) &= \left[\frac{1-6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}, \frac{5+6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}\right] \\
S_R(m_3) &= \left[\frac{5-6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}, 1\right]
\end{aligned}
$$

The core sets are:

$$
\begin{aligned}
C_R(m_1) &= \left[0, \frac{1-6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}\right] \\
C_R(m_3) &= \left[\frac{5+6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}, 1\right]
\end{aligned}
$$

If $\varepsilon > \frac{3-\sqrt{3}}{8}$, then $C_R(m_2) = \emptyset$; else $C_R(m_2) = \left[\frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}, \frac{5-6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}\right]$. Finally, the truth-degree functions are generally given by $\phi_i(x) = \int_{s_{k-1}}^{s_k} q(y|x)\,dx$. Hence, we have:

$$
\phi_1(x) = \begin{cases} 1 & x \in \left[0, \frac{1-6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}\right] \\ \frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}-6x}{12\varepsilon} & x \in \left[\frac{1-6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}, \frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}\right] \\ 0 & x \in \left[\frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}, 1\right] \end{cases}
$$

$$
\phi_3(x) = \begin{cases} 0 & x \in \left[0, \frac{5-6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}\right] \\ \frac{6x-\left(5-6\varepsilon-\sqrt{1+4\varepsilon^2}\right)}{12\varepsilon} & x \in \left[\frac{5-6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}, \frac{5+6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}\right] \\ 1 & x \in \left[\frac{5+6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}, 1\right] \end{cases}
$$

If $\varepsilon < \frac{3-\sqrt{3}}{8}$, then we have:

$$\phi_2(x) = \begin{cases} 0 & x \in \left[0, \frac{1-6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}\right] \\ \frac{6x-\left(1-6\varepsilon-\sqrt{1+4\varepsilon^2}\right)}{12\varepsilon} & x \in \left[\frac{1-6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}, \frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}\right] \\ 1 & x \in \left[\frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}, \frac{5-6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}\right] \\ \frac{\left(5+6\varepsilon-\sqrt{1+4\varepsilon^2}\right)-6x}{12\varepsilon} & x \in \left[\frac{5-6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}, \frac{5+6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}\right] \\ 0 & x \in \left[\frac{5+6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}, 1\right] \end{cases}$$

By contrast, if $\varepsilon \in \left[\frac{3-\sqrt{3}}{8}, \frac{3}{8}\right]$, then we have:

$$\phi_2(x) = \begin{cases} 0 & x \in \left[0, \frac{1-6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}\right] \\ \frac{6x-\left(1-6\varepsilon-\sqrt{1+4\varepsilon^2}\right)}{12\varepsilon} & x \in \left[\frac{1-6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}, \frac{5-6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}\right] \\ \frac{3-\sqrt{1+4\varepsilon^2}}{6\varepsilon} & x \in \left[\frac{5-6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}, \frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}\right] \\ \frac{\left(5+6\varepsilon-\sqrt{1+4\varepsilon^2}\right)-6x}{12\varepsilon} & x \in \left[\frac{1+6\varepsilon+\sqrt{1+4\varepsilon^2}}{6}, \frac{5+6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}\right] \\ 0 & x \in \left[\frac{5+6\varepsilon-\sqrt{1+4\varepsilon^2}}{6}, 1\right] \end{cases}$$

□

Proof of Lemma 1. Follows immediately from the continuity of $Q(y|x)$ in x. Recall that $\phi_k(x) = \int_{s_{k-1}}^{s_k} q(y|x)\,dy = Q(s_k|x) - Q(s_{k-1}|x)$. Then since $Q(\cdot|x)$ is continuous in x, so is ϕ_k. □

Proof of Lemma 2. Follows as a well known consequence of the monotone likelihood ratio property. To see this, first note that $\Phi_k(x) = \sum_{j=1}^{k} \phi_k(x) = \sum_{j=1}^{k} \int_{s_{j-1}}^{s_j} q(y|x)\,dy = Q(s_k|x)$. Take $x_1 > x_0$. It suffices to show that $Q(s_k|x_1) \leq Q(s_k|x_0)$ for every k.

By the monotone likelihood ratio property, we know that $\frac{q(y_0|x_1)}{q(y_0|x_0)} \leq \frac{q(y_1|x_1)}{q(y_1|x_0)}$ whenever $y_1 > y_0$, which we can rewrite as: $q(y_0|x_1)\,q(y_1|x_0) \leq q(y_1|x_1)\,q(y_0|x_0)$. This implies:

$$\begin{aligned} \int_0^{y_1} q(y_0|x_1)\,q(y_1|x_0)\,dy_0 &\leq \int_0^{y_1} q(y_1|x_1)\,q(y_0|x_0)\,dy_0 \\ Q(y_1|x_1)\,q(y_1|x_0) &\leq Q(y_1|x_0)\,q(y_1|x_1) \\ \frac{Q(y|x_1)}{Q(y|x_0)} &\leq \frac{q(y|x_1)}{q(y|x_0)} \end{aligned}$$

Similarly, we have:

$$\begin{aligned}
\int_{y_0}^{1} q\left(y_0|x_1\right) q\left(y_1|x_0\right) dy_1 &\leq \int_{y_0}^{1} q\left(y_1|x_1\right) q\left(y_0|x_0\right) dy_1 \\
q\left(y_0|x_1\right)\left[1-Q\left(y_0|x_0\right)\right] &\leq q\left(y_0|x_0\right)\left[1-Q\left(y_0|x_1\right)\right] \\
\frac{q\left(y|x_1\right)}{q\left(y|x_0\right)} &\leq \frac{1-Q\left(y|x_1\right)}{1-Q\left(y|x_0\right)}
\end{aligned}$$

Combining these gives:

$$\frac{Q\left(y|x_1\right)}{Q\left(y|x_0\right)} \leq \frac{1-Q\left(y|x_1\right)}{1-Q\left(y|x_0\right)}$$

for every y. This implies that $Q\left(y|x_1\right) \leq Q\left(y|x_0\right)$ for every y, which completes the proof.

To show that $F(x|y)$ respects first-order stochastic dominance, it suffices to show that the conditional densities $f(x|y)$ have the monotone likelihood ratio property. (If so, we can simply repeat the above method.) To show that MLRP is satisfied, take $x_1 > x_0$ and $y_1 > y_0$, and notice that:

$$\frac{f(x_1|y_1)}{f(x_0|y_1)} = \frac{f(x_1)q(y_1|x_1)}{f(x_0)q(y_1|x_0)} \geq \frac{f(x_1)q(y_0|x_1)}{f(x_0)q(y_0|x_0)} = \frac{f(x_1|y_0)}{f(x_0|y_0)}$$

where the middle inequality follows from the MLRP of q. □