CORA Cork Open Research Archive
Cartlann Taighde Oscailte Chorcaí

Title	Belief distributions and the measure of social norms
Authors	Wang, Cuizhu
Publication date	2023-07-31
Original Citation	Wang, C. (2023) 'Belief distributions and the measure of social norms', in Viciana, H., Gaitán, A. and Aguiar, F. (eds.) Experiments in Moral and Political Philosophy, pp. 125-143. Abingdon-on-Thames, Oxfordshire, England: Routledge. doi: 10.4324/9781003301424-8
Type of publication	Book chapter
Link to publisher's version	https://www.routledge.com/Experiments-in-Moral-and-Political-Philosophy/Viciana-Gaitan-Aguiar/p/book/9781032293905 - 10.4324/9781003301424-8
Rights	© 2023, the Author. All rights reserved. This is an Accepted Manuscript of a book chapter published by Routledge in Viciana, H., Gaitán, A. and Aguiar, F. (eds.) Experiments in Moral and Political Philosophy. Available online: https://doi.org/10.4324/9781003301424-8
Download date	2026-11-03 09:12:45
Item downloaded from	https://hdl.handle.net/10468/14285

UCC
University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

Belief Distributions and the Measure of Social Norms

Cuizhu Wang
Department of Philosophy
School of Society, Politics, and Ethics
University College Cork

dawncz_wang@hotmail.com
September 2022

Abstract

This chapter has the goal of advocating for the Quadratic Scoring Rule as the indispensable tool in the operationalization of one of the most popular concepts of social norms in the literature. Section 1 reviews the concept of social norms developed by Bicchieri (2006, 2017) and argues that this theory offers an operational definition of social norms which focuses on beliefs (expectations) and preferences and allows for the possibility of their empirical investigation. Section 2 introduces the well-known Intentional Stance and Revealed Preference Theory as the best interpretive framework for the intentional concepts applied in Bicchieri's conceptual analysis of social norms. On the basis of limitations arising from experimental work operating Bicchieri's theory, section 3 introduces and defends the Quadratic Scoring Rule (QSR) for eliciting subjective belief distributions. Three different types of arguments are given to support the suitability of this method in the measuring of social norms.

Keywords: Social Norms, Belief Distribution Elicitation, Intentional Stance, Revealed Preference Theory, Quadratic Scoring Rule, Savage

1. *Bicchieri's Conceptual analysis of Social Norms*

Cristina Bicchieri's analysis of social norms follows the game-theoretic tradition, and defines a social norm broadly as a Nash equilibrium in behavioral strategies. As a social norm is characterized as an equilibrium, it therefore is "supported by self-fulfilling expectations in the sense that players' beliefs are consistent, and thus the actions that follow from players' beliefs will validate those very beliefs."[1] Bicchieri defines a social norm as "a rule of behavior such that individuals *prefer* to conform to it on condition that they *believe* that (a) most people in their reference network conform to it (empirical *expectation*), and (b) that most people in their reference network believe they ought to conform to it (normative *expectation*)."[2]

There are two key features in this definition. First, there is the central role of mutually consistent expectations in individuals' "belief system"[3] in supporting a norm. Second, there is the existence of conditional preferences for norm compliant behavior.

1.1 *Social Expectations*

A social expectation is the type of expectation that can support the existence of social norms, according to Bicchieri (2006, 2017). It is the first key concept in Bicchieri's conceptualization of social norms. They are *social* in the sense that they are based on what we believe about the prevalence and distributions of others' behaviors and beliefs.

Bicchieri's model distinguishes between empirical expectations and normative social expectations. Empirical expectations are what people think "others *will* do". Empirical expectations can support a social norm if *enough* other people in a population follow the norm in question. Empirical expectations are the first important basis for norm compliance. In a social interaction, in order for a person to play her equilibrium strategy she has to predict what the others' strategies will be. The definition of "enough", according to Bicchieri, includes two aspects. First, different individuals may have different thresholds below which they consider the number of followers to be too small to count. Second, individuals may have different thresholds for different norms (Bicchieri 2010). A normative expectation is a belief

---

[1] Cristina Bicchieri, Ryan Muldoon, and Alessandro Sontuoso, "Social Norms," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (2018), 12, https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=social-norms

[2] Cristina Bicchieri, *Norms In The Wild: How to Diagnose, Measure, and Change Social Norms* (New York: Oxford University Press, 2017), 35.

[3] This term is applied here loosely to echo the intentional concepts in previous paragraph which are italicized.

about what behavior, in a situation relevant to a norm, people in the population approve or disapprove of. This is usually expressed as 'one ought to do X in situation Y'. Normative expectations are the second important basis for norm compliance (Sugden 1998, 2005; Bicchieri 2006). The ground of normative expectations is that we expect others' beliefs about what is praiseworthy or blameworthy that we have learned in the past to be continued in the future in the same situations.

Importantly, an expectation is a type of belief. It is a belief about "what is going to happen or what should happen."[4] In addition to distinguishing between empirical and normative social expectations, Bicchieri's model also distinguishes between *first-order* and *second-order* expectations/beliefs.[5] First-order beliefs are regular beliefs, and second-order beliefs are beliefs about beliefs.

First-order social beliefs are *empirical* in the sense that one might believe/expect people to act in a certain way in some situations. I call this type of beliefs first-order descriptive beliefs (hereafter 1D). *Second-order* social beliefs are *empirical* when one believes that others in her society hold some first-order empirical beliefs about what others and herself are going to do in some situations. These types of beliefs are called second-order descriptive beliefs (hereafter 2D).

First-order social beliefs are *normative*, when one believes that others *should* behave in a certain way in some situations. *Second-order* social beliefs are *normative* when one believes that others in her reference network hold the first-order normative beliefs about how others and herself *should* behave in some situations. These two types of beliefs are called first-order normative beliefs (hereafter 1N) and second-order normative beliefs (hereafter 2N), respectively.

Therefore, the complete range of types of *social beliefs* that count as the social expectations which matter to Bicchieri's analysis of social norms are as follows (paraphrased from Table 2.1, Bicchieri 2017):[6]

First-order Descriptive belief (1D): what *one believes* about what *others* do.

Second-order Descriptive belief (2D): what *one believes* about what *others believe I/others* do.

---

[4] Ibid., 11.
[5] Ibid., 70, Table 2.1.
[6] Ibid., 70.

First-order Normative belief (1N): what *one believes* about what others should do.

Second-order Normative belief (2N): what *one believes* about what *others believe I/others should* do.

According to Bicchieri, to confirm the existence of a norm, empirical expectations and normative expectations must have joint effects on people's behavior. In order for a person to play her equilibrium strategy in an interaction, she has to predict others' strategies. If she is confident in her expectations about the other players' beliefs, and these beliefs are consistent, she can expect the other players' strategies will follow from their beliefs or validate those very beliefs and this simultaneously extends to the other players whose behavior supports the norm. Thus, social norms are functions of networks of expectations. Meanwhile, the account bridges ascriptions at the social level and individual level and allows social norms to be identified in choice data about interacting individuals.

1.2 Conditional Preferences

A second key feature in Bicchieri's definition of social norms is the concept of conditional preference. The precise formal statement of the conditionality of social norms given by Bicchieri can be checked at various places (e.g., Bicchieri 2006, p.11) but the important point is that conditionality gives rise to *context* in explaining behaviors regarding social norms. In other words, our *preferences* concerning norm compliance are *conditional* on the decision context.[7]

Bicchieri (2010) argues that our propensity to follow social norms is not an unconditional one: it is conditional on having certain empirical and normative expectations, namely the empirical expectation that a sufficiently large subset of the reference group conforms to the norm, and the normative expectation that a sufficiently large subset of the reference group expects the agent to conform to the norm, prefer the agent to conform to the norm, and may sanction the norm-following behavior in such a kind of situation. Operationally, the conditionality for an individual to prefer conforming to a social norm, following Bicchieri, is the co-existence of the set of social beliefs as reviewed in section 1.1.[8]

---

[7] Cristina Bicchieri, *The Grammar of the Society: The Nature and Dynamics of Social Norms* (New York: Cambridge University Press, 2006), 11.

[8] The question of whether a full range of social beliefs are required for norm identification, as a general approach, is of relevance here. I have argued in Wang (2022, 19-39) that this is an empirical question, not a conceptual one. However, given the focus of this chapter, I set this issue aside.

Furthermore, Bicchieri (2008) argues that social norms *map* specific contexts into behavioral rules, and it is groups of such rules, taken together, that give meaning to concepts such as fairness or trustworthiness (Bicchieri 1999). This is a key feature of Bicchieri's model which distinguishes from other operational models of social norms in the literature. For example, the Fehr-Schmidt (1999)'s model captures the idea that people dislike unequal outcomes, and they lose utility propositional to a distance measure of inequality. This model is referred as the famous social preference account (see also Fehr and Gächter 2000, Fehr et al. 2002) in the literature, as one alternative approach to explaining norm compliant behavior. Social preference account models norms as strictly emergent from individuals' preferences over distributional social properties such as equality, equity, or reciprocity, and it reduces social norms to individuals' conceptions of the social good. Bicchieri and Zhang (2012) argue that the social preference account fails to recognize that individuals' preference for norm compliance are context-dependent, and so it ignores "a mapping from contexts to preferences that indicates in predictable ways how and why a given context or situation changes one's preference."[9]

A norm-based model should be able to directly test how contexts which matter for the norm can be studied in the lab and in the field without losing generality over different types of social norms (Binmore 2010). Evidence suggests that even simple changes to choice contexts are often sufficient to cause people to switch their choice-making (Engelmann and Strobel 2004). A norm-based model should also be able to show how a non-reducible account of social facts can be studied empirically through observable choices of individuals. Bicchieri's theory of social norms, given its emphasis on expectations and conditionality, suggests a methodological possibility for the identification of social norms as social facts from observations of the choices of individuals.

[9] Cristina Bicchieri, and Jiji Zhang, "An Embarrassment of Riches: Modeling Social Preferences in Ultimatum Games," in *Philosophy of Economics*, ed. Uskali Mäki (San Diego: North Holland, 2012), 581.

2. *The Intentional Stance and the Revealed Preference Theory*

2.1 Belief distributions and the Intentional Stance

As reviewed in section 1, intentional concepts such as beliefs, expectations, and preferences are the building blocks for Bicchieri's conceptualization of social norms. However, it is not clear in her conceptual analysis what these concepts refer to.

In everyday life, we typically appeal to others' propositional attitudes (e.g., their beliefs and preferences) to explain their actions and/or predict their further behavior. This is what philosopher Daniel Dennett calls taking the *intentional stance* (IS). The designation of *propositional attitudes* characterizes the family of relationships between agents and intentional (i.e., mental) content. According to Dennett (1991), propositional attitudes such as beliefs and desires are abstract posits that help us to track real patterns in observable behavior, including linguistic behavior. The IS is one of several predictive and explanatory strategies[10] one can adopt for explanation of systems (Dennett 1987). To adopt the IS is to treat the system as if it had motivating states characterized by propositional attitudes like belief, desire and preference which are about some object, situation, or event (Dennett 1987, 1988, 1991). According to Dennett, the intentional stance identifies "true believers". A true believer is a system whose behavior can only adequately be explained using the intentional stance. To treat someone as an intentional system, and therefore, as a true believer, is to treat them as rational, in a sense that Zawidzki (2007) formulates as follows:

> "(1) Intentional systems have the beliefs they ought to have, that is, true and relevant beliefs, given their perceptual capacities and informational needs; (2) Intentional systems have the desires they ought to have, given their biological needs; (3) Intentional systems behave in ways that count as rational given these beliefs and desires."[11]

Dennet (1991) argued that ignoring other agents' intentionality would lead to the missing of real patterns. Ladyman and Ross (2007) complete Dennett's theory of real patterns and provide a technical analysis of real patternhood in terms of computational efficiency and non-redundancy of projectability. This grounds a precise basis for rejecting mere instrumentalism about propositional attitudes by offering a rigorous test of pattern independence. The intentional stance therefore is not simply a strategy for taking an epistemological shortcut to apprehending minds. It is the stronger suggestion that all there is to having a belief that *p* is being a system that is predictable under the assumption that it

[10] The other two stances which act as predictive and explanatory strategies, are the physical stance and the design stance. (Dennett 1987, 1988, 1991)

[11] Tadeusz Zawidzki, *Dennett* (Oxford: Oneworld Publications, 2007), 48.

believes *p*. Ontologically, to say that *x* believes *p* is to assert that *x*'s behavior (verbal and otherwise) demonstrates a particular kind of regularity; namely, just that kind of regularity which justifies the 'projection' (the attribution by the observer) of *x*'s attitudes to the proposition that *p*. Methodologically, the propositional attitudes such as beliefs and desires are virtual properties of agents in environmental (particularly social) contexts which can be attributed from the IS.

The emphasis on expectations and conditionality in Bicchieri's theory of social norms is compatible with the intentional stance in a sense we can render as follows: Bicchieri's account of social norms suggests that the relevant networks of expectations are social facts the agent encounters, and that the best methodology for measuring behavioral responses to these social facts is to construct utility functions from choice behavior that are conditional on recognition of any norm, but do not vary with outcomes distinguished by reference to their normative assessments.

2.2 Revealed Preference Theory

In economics, Samuelson's Revealed Preference Theory (RPT) showed the necessary mathematical conditions which enable ascription of a consistent set of preference to an agent based on her behavior without taking a position about psychology. As Hands (2013) argues Contemporary Revealed Preference Theory (CRPT), as an extension of original RPT, not only allows for the construction of utility functions empirically from finite choice data, but also embeds philosophical and methodological claims. CRPT, like other versions of RPT, defines preference solely in terms of choice behavior and denies that preferences are causes of choices. It advocates interpreting RPT as a general methodological template for all of microeconomics. My arguments about RPT refers to CRPT.

Binmore (2009) points out that RPT "makes a virtue of assuming nothing whatever about the psychological causes of our choice behavior."[12] This means a RPT analysis is based only on information about what subjects choose. Given that observed behavior tells us people choose in some situations, these data can be used to deduce what they would choose in other situations. However, in making such inferences there are two assumptions about agents' choice behavior: stability and consistency.[13]

---

[12] Ken Binmore, *Rational Decisions* (Princeton: Princeton University Press, 2009), 8.
[13] Ibid., 9-12.

Stability and consistency allow the use of RPT to model agents as seeking to optimize the value of something. Whatever this abstract something may be in a particular context, economists call it "utility". The utility function constructed based on assumptions of stability and consistency therefore can be synthesized as the "real pattern" tracked by our intentional states from observable behavior.

Following the intentional stance, belief attribution requires a similarly thin assumption of rationality, according to which there is a systematic relationship between agents' choices and their goals - a real pattern. This sense of rationality is called "ecological rationality" in economics by Smith (2007). In RPT, the criterion that is relevant when picking one of the infinity of utility functions that represent a given preference relation is that of mathematical convenience. Von Neumann and Morgenstern's (VNM) approach to utility gives a primary role to risk. VNM utility measures how much an agent wants something by the relative magnitude of the risk she is willing to take to get it. Expected Utility Theory (EUT) requires that the preferences over lotteries that an agent reveals in risky situations be consistent. EUT was used by Savage (1972) to develop the standard model of subjective probability.

According to EUT, an agent acts as though she believes that each possible state of the world in a decision problem has a subjective probability. Savage then models the agent's utility by assuming that there exists a subjective probability distribution *and* a utility function such that observed choices can be characterized as maximizing her Subjective Expected Utility (SEU). The existence proof for identification of SEU requires satisfaction of a number of axioms which can be easily checked in the literature (Savage 1972, 18-31) and which, for our purposes, require that preferences are separable across events (the Sure-Thing principle), the preference between two acts depend solely on the consequences in states in which the payoffs of the two acts being compared are distinct (Monotonicity), and, betting preferences be independent of the specific consequences that define the bets (Weak comparative probability).

Savage argues that the applicability of SEU is restricted to what he calls a "small world". A small world is one within which there is always a possibility for the agent to "look before you leap."[14] In this world, it is arguably possible for an agent to take account in advance of the implications of all conceivable future information for her subjective utility. Updating information therefore can correct mistakes in her model.

[14] Ken Binmore, *Rational Decisions* (Princeton: Princeton University Press, 2009), 117.

RPT is frequently criticized in economics as a form of instrumentalism, echoing the most common objections to the intentional stance among philosophers.[15] This criticism states that modeling choices using RPT is just a way of applying math: utility or preference identified by RPT does not really exist. According to this critique, RPT is an instrument instead of a model of reality.[16] The critique is mainly based on two arguments: first, preferences cannot be revealed by choice alone as RPT suggests; second, RPT fails to support the assumption that behavior is caused by mental states (Hausman (1992, 2000, 2008); Sen (1973, 1980, 1993, 1997)). The basis for the realist interpretation of the intentional stance can also provide a clue as to why these criticisms are misguided.

Dennett and Ross maintain that what is real about mind is accessed through observable behavior by taking the intentional stance to identify real patterns, whereas Hausman et al. maintain that internal mental states, which are unobservable, are constitutive of real mental life.

Given the defense of the realist interpretation of the IS as opposed to the instrumentalist interpretation of the IS, we can apply the same realist interpretation to RPT. If IS realism stands, then the utility function revealed from an agent's choices, according to RPT, can be regarded as real, although not as the cause of choices. A utility function constructed on the basis of RPT is, rather, a model of the choices themselves (Binmore 2009), in a statistical sense.

For our purposes, the key point is that, for modeling social norms, there is no appeal to private preferences about how society should be. Instead, it is the networks of expectations as social facts which, from taking the intentional stance, we can see as conditioning agents' behavioral responses.

3. Quadratic Scoring Rule to Model Subjective Belief Distributions

Various experimental studies by Bicchieri and co-authors on norms have aimed to provide data to evaluate her theory of social norms: see Bicchieri and Xiao 2009; Bicchieri and Chavez 2010, 2013; Bicchieri, Xiao and Muldoon 2011; Xiao and Bicchieri 2010; Bicchieri and Zhang 2012. These experimental designs suffer from four main weaknesses.

---

[15] Don Ross, *Philosophy of Economics (*New York: Palgrave Philosophy Today, 2014), 69.

[16] For critiques about RPT along these lines, see Hausman (1992, 2000, 2008) and Sen (1973, 1980, 1993, 1997). For a detailed survey of the debates, see Hands (2013) and Ross (2014).

First, both empirical and normative expectations are measured by means of *modal* responses, rather than *distributions*.[17] For example, in Bicchieri and Xiao (2009) subjects were asked "How many dividers[18] in this room do you think answered 'Yes' to question (d)? (If your answer is the same as the actual number, you will receive an additional $1);"[19] in Bicchieri, Xiao and Muldoon (2011) subjects were asked "What is the option that you think most participants chose today? (If your answer is right, you will get one point)."[20] An answer to this type of question represents the *modal value* in statistics. Questions asked in such format do not elicit *distributions* of beliefs, and so are insensitive to degrees of confidence in beliefs. And points two, three and four, their experiments fail to apply sufficient incentive control, assume risk-neutrality despite overwhelming evidence that most people are risk averse, and employ a simplistic understanding of belief consistency.

I have identified and explained these limitations elsewhere (Wang 2022).[21]The purpose here will be to point out why some limitations that arise in Bicchieri and co-authors' empirical work can be solved by applying an alternative methodology.

In what follows, I will suggest the experimental methodology of applying a Quadratic Scoring Rule (QSR) for eliciting people's subjective beliefs in operationalizing Bicchieri's analysis of social norms. There are three reasons to promote the QSR. First, it considers the degrees of confidence in subjects' beliefs, which allows for rich econometric modeling. Second, it is compatible with the principle of risk-neutralization for elicitation of belief distributions. Third, it is compatible with salient and dominant incentivization. The QSR hence tackles some of the limitations that emerge in the experimental studies by Bicchieri and co-authors.

3.1 Subjective Probability for Subjective Belief

The concept of subjective belief captures the fact that individuals may have different

[17] The concepts of modal beliefs and distributions here reflect the key difference between point estimate versus distribution estimate in statistics. This difference is illustrated by Merkle and Weber (2011). Wang (2022) applies the concepts when analyzing the methodology of belief elicitations applied in some experimental work by Bicchieri and co-authors' empirical work.

[18] "Divider" is the term applied in their experimental instructions, which refers to "dictator" or "proposer" in the dictator game.

[19] Cristina Bicchieri, and Etre Xiao, "Do the Right Thing: But Only If Others Do So," *Journal of Behavioral Decision Making* 22, (2009): 201.

[20] Cristina Bicchieri, Erte Xiao, and Ryan Muldoon, "Trustworthiness is a Social Norm, but Trusting is Not," *Politics, Philosophy & Economics* 10, no. 2 (2011): 183-186.

[21] Cuizhu Wang, "Normative Expectations and Subjective Beliefs: An Incentivised Experimental Study" (PhD thesis, University College Cork, 2022), 49-71.

degrees of belief or judgments due to different backgrounds in their knowledge or experiences. It was formalized by Savage (1972) as an extension of Revealed Preference Theory. In statistics, subjective belief refers to the degree of belief as a quantified judgment[22] which is represented mathematically by subjective probability. Subjective probability enables one to explain the distribution of the belief about different events.

Operationally, subjective probabilities are defined as those probabilities that causally lead an agent to choose some prospects over others when the outcomes of those prospects depend on events that are not yet actualized. A simple prospect here might be a bet on whether one specific event will occur. Since the events are not yet actualized, and there are several possible events and many possible risky bets about these events, belief elicitation naturally involves risk. In general, we need to elicit or control for risk preference in order to infer belief.

Scoring rules, betting tasks and simulated auctions are the most popular methods[23] for eliciting subjective probabilities for individuals, and for the purpose of studying inferences about subjective beliefs. In economic terms, scoring rules are functions mapping a subject's reported beliefs about a random variable and the ex-post realization of that random variable into a payoff for the subject. The most popular scoring rule is the Quadratic Scoring Rule (QSR), first introduced by Brier (1950).

The formula for the QSR is as follows:

The score P is defined by

$$P = \frac{1}{n}\sum_{j=1}^{r}\sum_{i=1}^{n}\left(f_{ij} - E_{ij}\right)^2$$

where E_{ij} takes the value 1 or 0 according to whether the event occurred in class j or not. In the formula, n denotes the occasions an event can occur; j refers to the possible classes or categories of each occasion. For one of the occasions i, the forecast probabilities are f_{i1}, $f_{i2}, \ldots f_{ir}$. The r classes are chosen to be mutually exclusive and exhaustive so that

$$\sum_{i=1}^{\gamma} f_{ij} = 1, i = 1,2,3, \ldots n$$

An illustrative example discussing the score P is given further by Brier (1950).[24]

[22] Robert L. Winkler, and Allan H. Murphy, "Nonlinear Utility and the Probability Score," *Journal of Applied Meteorology and Climatology* 9, no. 1 (1970): 143-148.

[23] See Matheson and Winkler (1976), Savage (1972), and Anderson et al. (2014).

[24] Glenn W. Brier, "Verification of Forecasts Expressed in terms of Probability," *Monthly Weather Review* 75, (1950): 2.

The QSR relaxes the imposition of binary event structure on beliefs that might be about continuous event spaces. Borrowing an example given by Harrison and Ross (2016), a question "A typical American male will live 70 or more years. True or false?" concerns a binary event. By contrast, "How long will the typical American male live?" is its corresponding continuous event. The theory of belief elicitation for continuous events also applies to non-binary, discrete events. For example, the question, "A typical American male will live to be between 0 and 50, 51 and 65, or more than 65. How likely are each of these age categories?" inquiries about a 3-valued event space.[25] Beliefs over these 3 events define a probability *mass* function. Beliefs over the continuous event define a probability *density* function.

The framework of scoring rule for binary events can be formally viewed as a trading game between the agent and the reporter, following Savage (1972). The payment amount from the agent's report in the binary event is determined by two positive parameters α and β which define the QSR: α determines a reward and β determines a penalty. Assume an event has two possible outcomes A and B which are complements (so if A does not occur, B must occur), and let be θ the reported probability for A, and let Θ be the true binary valued outcome for outcome A. This means if A occurs $\Theta=1$, and if A doesn't occur $\Theta=0$. Andersen et al. (2014) show that the QSR score or payment penalizes the subject by the squared deviation of the report from the true binary valued outcome Θ.[26]

For a risk neutral agent, under the assumption of SEU,[27] her report is her belief. This is because under SEU, the Subjective Expected Value of a risk prospect is equivalent to the certainty equivalent of that prospect, therefore the risk premium is zero. This means the reported belief distribution for a risk neutral agent is her latent belief distribution. However, it is well known that the payment to the belief reports over binary events in the QSR can be affected by the risk attitude of an agent.[28] For a risk averse agent, the report will be different from the report of a risk neutral agent. Risk averse agents tend to report their beliefs close to ½ for binary events,[29] with varying degrees of distortions of reports from beliefs caused by

[25] Glenn W. Harrison, and Don Ross, "The psychology of human risk preferences and vulnerability to scare-mongers: experimental economic tools for hypothesis formulation and testing," *Journal of Cognition and Culture* 16 (2016): 403.

[26] Steffen Andersen, John Fountain, Glenn W. Harrison, and E. Elisabet Rutström, "Estimating Subjective Probabilities," *Journal of Risk and Uncertainty* 48, no. 3 (2014): 212.

[27] SEU refers to the theory of Subjective Expected Utility developed by Savage (1972).

[28] See Winkler and Murphy (1970), Savage (1972), and Kadane and Winkler (1988).

[29] Steffen Andersen, John Fountain, Glenn W. Harrison, and E. Elisabet Rutström, "Estimating Subjective Probabilities," *Journal of Risk Uncertainty* 48, no. 3 (2014): 212.

different degrees of risk aversion.

In the context of social norms, beliefs are better modeled as probability distributions over continuous events instead of binary events. For example, if we ask a Chinese person to predict whether a young Chinese couple in his social circle will spend all their savings and even bear debts to host a luxurious wedding for the sake of social image, this agent may not be able to report his belief simply as yes/no, but he may refer to his knowledge of the couple and evaluate the probability of or say "to what degree" the couple may follow the norm or not. Scoring Rules for continuous events are then a better model for eliciting such types of beliefs.

Matheson and Winkler (1976) developed the QSR for eliciting people's subjective beliefs for continuous distributions. The scoring rule can induce truthful reports if the agent is risk neutral. Each report in the scoring rule induces a lottery, and the risk attitudes affect the inference of the subjective belief reports. The subjects are rewarded for the accuracy of their beliefs, and the definition of the penalty can be severe if the accuracy misses the true answer.

Assume K events, which means the possible response to the question being asked is partitioned into K intervals, and that each agent's report over the K intervals adds up to 100 tokens. Harrison et al. (2017) ask each subject to conduct a task of allocating 100 tokens by sliding a bar across each of 10 bins that range across the possible realization values in order to discover their beliefs. As summarized in Harrison et al. (2017) that the QSR payment score works as follows: if k is the interval the actual value lies in, the reward in the score is a doubling of the report allocated to the true interval k, and the penalty in the score depends in a "quadratic manner" on the distribution of the reports over K intervals. For all SEU agents, in order to gain the highest reward, she would never allocate any tokens to a bin if her belief to that interval is zero; she will only allocate different levels of tokens to 2 bins if her belief differs for the 2 intervals (hence, if her belief is the same for the 2 intervals, she would allocate the same number of tokens to each bin of the 2 bins). With risk aversion, the agent would behave as if "flattening" the report for those bins she has positive belief of occurring, so as to avoid risk by reducing the variability of utility over different possible outcomes.

In Wang (2022), a QSR is applied to elicit subjects' subjective beliefs. Wang (2022) uses 4-bin belief questions to elicit discrete distributions and 10-bin belief questions to elicit discretized versions of continuous distributions.[30] Each bin is defined by an event if discrete

[30] Cuizhu Wang, "Normative Expectations and Subjective Beliefs: An Incentivised Experimental Study" (PhD thesis,

(e.g. 3 out of 4 people) in the 4-bin questions, or by an interval within which the subjective belief might lie if real-valued (e.g. 0% to 9% of all tokens) in the 10-bin questions. The first-order belief questions are 4-bin questions which ask for beliefs about how many out of 3 randomly selected people would or should choose to accept or reject a proposed split by the Proposer. The 4 bin options are "nobody", "1 out of 3 people", "2 out of 3 people" and "all people". The second-order belief questions elicit beliefs about the percentage of tokens allocated by everyone else in the task when reporting their first-order beliefs with respect to the 4 bin options. These second-order beliefs are 10-bin questions, where the intervals of the 10 bins are "0% to 9% of all tokens", "10% to 19% of all tokens", "20%-29% of all tokens",… to "90%-100% of all tokens". The application of the QSR with the interface developed by Harrison et al. (2017) allow experimenters to study individuals' belief distributions without burdening the subjects with descriptive formulae that most would not understand.

3.2 Confidence in Subjective Belief Distribution

The first reason for adopting the QSR is the capability to consider degrees of confidence when beliefs are elicited as distributions. Harrison and Swarthout (2020) note three distinct definitions of overconfidence in the literature. The first definition is **overestimation** of one's absolute ability to perform a task. For example, one may report a belief of a higher score than one can actually get in a math task. The second definition is **overplacement** of one's performance relative to others. Overplacement is usually known as the "better than average effect", where everyone is relatively better than average. The third definition considers **overprecision** about the certainty in the accuracy of one's belief. Overconfidence in the sense of overprecision refers to the belief expressing unwarranted certainty in its precision.

It is overprecision as one of the hypotheses about overconfidence that is particularly important to the study of belief elicitation regarding social norms. The need to evaluate the hypothesis about overprecision is explained by Moore and Healy (2008), Merkle and Weber (2011) and Benoît and Dubra (2011).[31]

University College Cork, 2022), 99-110.

[31] Harrison and Swarthout (2020) argue that it is not clear how inferences about statistically significant overestimation or overplacement can be made without knowing how much precision individuals have about certain beliefs. They point out that because beliefs about overestimation and overplacement are elicited in the current literature as point

The main issue here is that statistical tests about bias of estimates of distributions involve standard *errors*. However, confidence about belief distribution involves standard *deviations* of that distribution. For each statistic of the belief distribution the econometric model generates a point estimate and a standard error of that point estimate. One statistic for the normal distribution, assumed here to characterize beliefs, is the mean. So the mean has a point estimate and a standard error. The other statistic for the normal distribution is the standard deviation, so it also has a point estimate and a standard error.

This problem of confusing *standard deviation* with *standard error* causes confusion between tests of the significance of bias of a subjective belief distribution with tests of the bias of an estimate of the mean of the belief distribution.

Confidence is a concept of "imprecision" from statistical estimation. It could also be regarded as reflecting the variability of beliefs when the beliefs aren't normal distributions. This problem is assessed by Harrison and Swarthout (2020) which is sharply illustrated by the problem of "precise estimates of imprecise beliefs."[32]

3.3 Induce Risk Neutrality

A second reason for applying the QSR is to do with risk control for belief elicitation in experimental settings.

The scoring rule is predicted to directly induce truthful belief reports only if the agent is risk neutral. However, Andersen et al. (2014) show that risk attitudes will affect the incentive to report one's subjective probability "truthfully" in the QSR. The results from their study show that one has to be sensitive to the risk attitudes of subjects before drawing inferences about subjective probabilities from responses to scoring rules. Harrison et al. (2017) show that a subjective probability distribution can deviate from the reported distribution under a QSR applied to non-binary events, and that risk aversion theoretically causes the individual to report a "flattened" version of their true distribution.

An alternative way to account for risk preferences is to apply the Binary Lottery Procedure to induce risk neutral behavior. It allows for eliciting subjective probabilities without correcting for risk attitudes. Smith (1961) was the first to suggest the solution of a

---

estimates, standard practice downplays the importance of sufficient imprecision in beliefs.

[32] Glenn W. Harrison, and Todd J. Swarthout, "Belief distributions, Bayes rule and Bayesian overconfidence," *CEAR Working Paper WP 2020-11*, The Center for the Economic Analysis of Risk, (2020): 18. https://cear.gsu.edu/working-papers/

"probability currency", which allows inferences of the subjective probabilities over some binary event from subjects' choices over the bets, without necessarily knowing about their utility functions. Roth and Malouf (1979) employed this procedure in many experiments and assumed that we can evaluate the expected utility of two bargaining agents with this device.

Harrison et al. (2013) propose a simple payment procedure in the lab to implement the Binary Lottery Procedure. In their setting, subjects are asked to make a single choice over a pair of lotteries defined over money and objective probabilities. However, instead of earning money when one of their choices is played out, subjects earn points. These points convert into increased objective probability of winning another binary lottery. Through observed choice patterns, they found that the lottery procedure induces risk neutrality robustly when subjects are given one task. Harrison et al. (2013, 2014, 2015) evaluated the Binary Lottery Procedure in a laboratory experiment for inducing risk neutral behavior in a belief elicitation task.[33]

In summary, I argue that risk attitudes have to be elicited and jointly estimated with belief reports in order to recover beliefs, or that risk attitudes have to be neutralized by experimental design so that the subjects behave as if they are risk neutral for that belief elicitation task. A key point is, when belief elicitation is involved for the empirical study of social norms, one cannot just assume risk preferences away: one has to either jointly estimate *ex post* the data collection for risk preferences and belief reports, or design the experiment *ex ante* to risk-neutralize subjects.

3.4 Compatibility with Incentive Control

A third reason for applying the QSR for eliciting subjective beliefs is that the QSR can guarantee incentive controls for subjective probabilities over events.[34] In experiments on belief elicitation through survey questions, the first fundamental distinction is whether the surveys are *hypothetical* or financially[35] *incentivized*. Hypothetical choices mean that

[33] These studies confirmed the big advantage of applying the Binary Lottery Procedure in belief elicitation task is that it allows the data from the reported belief distributions to be directly applied as subjective beliefs without further need to recover beliefs from observed reports.

[34] See Florian Artinger, Filippos Exadaktylos, Hannes Koppel, and Lauri Saaksvuori, "Applying Quadratic Scoring Rule transparently in multiple choice settings: A note," *Jena Economic Research Papers*, no. 021 (2010), Friedrich Schiller University Jena and Max Planck Institute of Economics; Steffen Andersen, John Fountain, Glenn W. Harrison, and E. Elisabet Rutström, "Estimating Subjective Probabilities," *Journal of Risk and Uncertainty* 48, no. 3 (2014): 207-229.

[35] Incentives could also be in the form of gift card, for example. However, in this thesis, the financial incentives solely refer to monetary incentives.

preferences or beliefs revealed by the choices have *no* consequences for the decision maker, whereas incentivized choices refer to preferences and beliefs revealed when there are *real* consequences for the decision maker of alternative choices. Merely hypothetical choices give rise to so-called hypothetical bias. In some choice modeling contexts, however, it is not possible to avoid hypothetical bias entirely. Therefore, it must be mitigated if possible.

Harrison (2014) suggests two ways to mitigate hypothetical bias: by means of instrument calibration *before* the survey or by means of statistical calibration *after* the survey. Instrument calibration focuses on how the hypothetical questions are posed. For example, Blumenschein et al. (2001) suggested that survey questions be asked in a way that encourages subjects to understand what a "hypothetical yes" means. And Cummings et al. (1997), Cummings and Taylor (1998) and List (2001) suggested that the survey questions themselves could encourage subjects to avoid hypothetical bias by telling them about the issue of hypothetical bias. The evidence supporting these procedures is mixed. The approach of statistical calibration developed by Blackburn et al. (1994) focuses on applying some bias functions which attempt to predict the response the subjects would have given if facing incentives after observing how they actually responded when not facing incentives. The usefulness of statistical estimation rests upon the assumption that hypothetical choices are informative. The evidence for statistical calibration is positive so far when applied to valuation of common private goods.

The definition of incentivization includes three aspects.

The first aspect of incentivization is **salience**. Salience means that the rewards must vary with the performance of the subject in the choice task. That is, the rewards should connect to the *message* subjects send. In the case of belief elicitation, for example, salience means that the token allocations subjects make affect their payoffs. In one scoring rule task for belief elicitation, subjects allocate tokens by moving the slider between different possible outcomes. The fact that the behavior of a subject of moving the slider changes her payoff in some way means that the reward in the task is salient. Of course, incentives can be non-salient. Non-salient rewards refer to rewards that do not vary with the performance of the subject in the task. For example, if a subject gets paid $20 to complete a task, no matter what answer the subject chooses in the task, then this is a non-salient reward. Smith (1982) provided a precise rationale for the motivational role of salient financial incentives in experimental economics.

The second aspect of incentivization is **dominance.** Payoff dominance basically means that the rewards, apart from being salient, have to be large enough so that the subject cares about which response she is giving. Dominance means that the size of payoffs must be large enough to be *noticeable*. This is called "just-noticeable-difference" in psychology. Smith and Walker (1993) show that subjects respond to increased incentives by expending more cognitive effort, and that this effort can reduce the variance of responses. Of course, reduced variation is different from reduced bias.

One aspect of dominance, which is subtle, is that the payoffs must be sufficiently rewarding to avoid the "flat payoff problem" stressed by Harrison (2006). The difference in payoffs has to be enough to motivate the subject to engage in some cognitive effort. Imagine, a subject is uncertain about choosing between outcome A and outcome B, however, the payoff difference between changing from choosing A to B is too *unnoticeable*, such as $0.01. Such a small difference is arguably not *dominant* enough for a subject to be motivated to take more cognitive effort to decide on whether to choose one outcome or the other. A flat payoff function means no matter how the allocation of tokens is changed, subjects effectively get the same amount of payoff. A flat payoff function implies that the incentive is not dominant.

The third aspect of incentivization is **incentive compatibility**, assuming salience and dominance. Incentive compatibility means that the incentives are designed in the way that the rewards give the subjects incentives to *truthfully* reveal their preferences or their beliefs. Incentive compatibility of stated choice experiments is often confused with the issue of hypothetical bias. The concept of incentive compatibility is more than just providing real consequences for the choices respondents make. It requires that incentives are designed so that different consequences "make it in the best interest of the respondent to respond truthfully," and "this connection has to be behaviorally transparent and credible."[36] The QSR uses salient and dominant incentives, and it is incentive compatible if the objective is to elicit belief distributions assuming the subjects are risk neutral.

4. Conclusion

This chapter argues that beliefs, as defined in the theory of social norms developed by Bicchieri, must be elicited as distributions. Section 1 summarized Bicchieri's theory of social

[36] Ibid., 236.

norms with a focus on her applications of the intentional concepts. Section 2 provided a philosophical frame which aligns with Bicchieri's conceptual analysis of social norms. I argued that the IS as a realist interpretation of the mind is the preferred philosophical background for the study of social norms, and that RPT is the methodology which implements this philosophy in the context of experimentation. Empirical studies of social norms should adopt the IS as the philosophical background, and apply RPT for the construction of utility functions. Section 3 based on the concerns about the limitations arise from the experimental work by Bicchieri and co-authors, and suggested an alternative methodology which tackles those limitations. It provided three reasons why the QSR as the best methodology for eliciting expectations on which Bicchieri's account of social norms relies, and is the best tool for improved operationalization of Bicchieri's theory of social norms through observed behavior in the lab. QSR therefore should be applied when eliciting belief distributions for the measurement of social norms following Bicchieri's conceptual framework.

References

Andersen, Steffen, John Fountain, Glenn W. Harrison, and E. Elisabet Rutström. "Estimating Subjective Probabilities." *Journal of Risk and Uncertainty* 48, no. 3 (2014): 207-229.

Artinger, Florian, Filippos Exadaktylos, Hannes Koppel, and Lauri Sääksvuori. "Applying Quadratic Scoring Rule Transparently in Multiple Choice Settings: A note." *Jena Economic Research Paper*, no. 21 (2010). Friedrich Schiller University Jena and Max Planck Institute of Economics.

Blackburn, Mckinley, Glenn W. Harrison, and E. Elisabet Rutström. "Statistical Bias Functions and Informative Hypothetical Surveys." *American Journal of Agricultural Economics* 76, no. 5 (1994): 1084-1088.

Bicchieri, Cristina. "Local Fairness." *Philosophy and Phenomenological Research* LIX, no. 1 (1999): 229-236.

---. *The Grammar of the Society: The Nature and Dynamics of Social Norms.* Cambridge: Cambridge University Press, 2006.

---. "The Fragility of Fairness: An Experimental Investigation on the Conditional Status of Pro-social Norms." *Philosophical Issues* 18 (2008): 229-248.

---. *Norms In The wild: How to Diagnose, Measure, and Change Social Norms*. New York: Oxford University Press, 2017.

---. "When Equality Trumps Reciprocity." *Journal of Economic Psychology* 31 (2010): 456-470.

Bicchieri, Cristina, and Alex Chavez. "Behaving as Expected: Public Information and Fairness Norms." *Journal of Behavioral Decision Making* 23 (2010): 161-178.

---. "Norm Manipulation, Norm Evasion: Experimental Evidence." *Economics & Philosophy* 29, no. 2 (2013): 175-198.

Bicchieri, Cristina, and Erte Xiao. "Do the Right Thing: But Only If Others Do So." *Journal of Behavioral Decision Making* 22 (2009): 191-208.

Bicchieri, Cristina, Erte Xiao, and Ryan Muldoon. "Trustworthiness is a Social Norm, but Trusting is Not." *Politics, Philosophy & Economics* 10, no. 2 (2011): 170-187.

Bicchieri, Cristina, and Jiji Zhang. "An Embarrassment of Riches: Modeling Social Preferences in Ultimatum Games." In *Philosophy of Economics*, 577-595. Edited by Uskali Mäki. San Diego: North Holland, 2012.

Bicchieri, Cristina, Ryan Muldoon, and Alessandro Sontuoso. "Social Norms." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta (Winter Edition 2018). https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=social-norms

Binmore, Ken. *Rational Decisions*. Princeton: Princeton University Press, 2009.

---. "Social Norms or Social Preferences?" *Mind & Society* 9 (2010): 139-157. doi: 10.1007/s11299-010-0073-2

Benoît, Jean-Pierre, and Juan Dubra. "Apparent Overconfidence." *Econometrica* 79, no. 5 (2011): 1591-1625.

Brier, Glenn. W. "Verification of Forecasts Expressed in terms of Probability." *Monthly Weather Review* 75 (1950): 1-3.

Cummings, Ronald G., Steven Elliott, Glenn W. Harrison, and James Murphy. "Are Hypothetical Referenda Incentive Compatible?" *Journal of Political Economy* 105, no. 3 (1997): 609-621.

Cummings, Ronald G., and Laura Osborne Taylor. “Does Realism Matter in Contingent Valuation Surveys?” *Land Economics* 74, no. 2 (1998): 203-215.

Dennett, Daniel C. *The Intentional Stance*. Cambridge: MIT Press, 1987.

---. “Out of the Armchair and Into the Field.” *Poetics Today* 9, no. 1 (1988): 205-221.

---. “Real Patterns.” *The Journal of Philosophy* 88, no. 1 (1991): 27-51.

Engelmann, Dirk, and Martin Strobel. “Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments.” *American Economic Review* 94, no. 4 (2004): 857-869.

Fehr, Ernst, and Simon Gächter. “Fairness and Retaliation: The Economics of Reciprocity.” *Journal of Economic Perspectives* 14, no. 3 (2000): 159-181.

Fehr, Ernst, Urs Fischbacher, and Simon Gächter. “Strong Reciprocity, Human Cooperation, and the Enforcement of Social Norms.” *Human Nature* 13, no. 1 (2002): 1-25.

Fehr, Ernst, and Klaus M. Schmidt. “A Theory of Fairness, Competition, and Cooperation.” *The Quarterly Journal of Economics* 114, no. 3 (1999): 817-868.

Hands, Wade D. “Foundations of Contemporary Revealed Preference Theory.” *Erkenntnis* 78, no. 5 (2013): 1081-1108.

Harrison, Glenn W. “Making Choice Studies Incentive Compatible.” In *Valuing Environmental Amenities Using Stated Choice Studies*, 67-110. Edited by Barbara J. Kanninen. Dordrecht: Springer, 2006.

---. “Real Choices and Hypothetical Choices.” In *Handbook of Choice Modelling, 236-254*. Edited by Stephen Hess and Andrew Daly. Edward Elgar, 2014.

Harrison, Glenn W., Jimmy Martínez-Correa, and Todd J. Swarthout. “Inducing Risk Neutral Preferences with Binary Lotteries: A Reconsideration.” *Journal of Economic Behavior & Organization* 94 (2013): 145-159.

---. “Eliciting Subjective Probabilities with Binary Lotteries.” *Journal of Economic Behavior & Organization* 101 (2014): 128-140.

Harrison, Glenn W., Jimmy Martínez-Correa, Todd J. Swarthout, and Eric R. Ulm. "Eliciting Subjective Probability Distributions with Binary Lotteries." *Economics Letters* 127 (2015): 68-71.

Harrison, Glenn W., Jimmy Martínez-Correa, Todd J. Swarthout, and Eric R. Ulm. "Scoring Rules for Subjective Probability Distributions." *Journal of Economic Behavior & Organization* 134 (2017): 430-448.

Harrison, Glenn W., and E. Elisabet Rutström. "Risk Aversion in the Laboratory." In *Risk Aversion in Experiments, 41-196*. Edited by Cox, James C. and Glenn W. Harrison. Bingley: Emerald, 2008.

Harrison, Glenn W., and Don Ross. "The Psychology of Human Risk Preferences and Vulnerability to Scare-mongers: Experimental Economic Tools for Hypothesis Formulation and Testing." *Journal of Cognition and Culture* 16 (2016): 383-414.

Harrison, Glenn W., and Todd J. Swarthout. "Belief Distributions, Bayes Rule and Bayesian Overconfidence." *CEAR Working Paper WP 2020-11*, The Center for the Economic Analysis of Risk, (2020). https://cear.gsu.edu/working-papers/

Hausman, Daniel M. *The Inexact and Separate Science of Economics*. Cambridge: Cambridge University Press, 1992.

---. "Revealed Preference, Belief, and Game Theory." *Economics & Philosophy* 16, no. 1 (2000): 99-115.

---. "Mindless or Mindful Economics: A Methodological Evaluation." In *Foundations of Positive and Normative Economics: A Handbook,* 125-55. Edited by Andrew Caplin and Andrew Schotter. Oxford: Oxford University Press, 2008.

Kadane, Joseph, and Robert L. Winkler. Separating Probability Elicitation from Utilities. *Journal of the American Statistical Association* 83, no. 402 (1988): 357-363.

Ladyman, James, and Don Ross. *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press, 2007.

List, John A. "Do Explicit Warnings Eliminate the Hypothetical Bias in Elicitation Procedures? Evidence from Field Auctions for Sports Cards." *American Economic Review* 91, no. 5 (2001): 1498-1507.

Matheson, James E., and Robert L. Winkler. “Scoring Rules for Continuous Probability Distributions.” *Management Science* 22, no. 10 (1976): 1087-1096.

Merkle, Christopher, and Martin Weber. “True Overconfidence: The Inability of Rational Information Processing to Account for Apparent Overconfidence.” *Organizational Behavior and Human Decision Processes* 116, no. 2 (2011): 262-271.

Moore, Don. A., and Paul J. Healy. “The Trouble with Overconfidence.” *Psychological Review* 115, no. 2 (2008): 502.

Ross, Don. *Philosophy of Economics*. New York: Palgrave Philosophy Today, 2014.

Roth, Alvin E., and Michael W. Malouf. “Game-theoretic Models and the Role of Information in Bargaining.” *Psychological Review* 86, no. 6 (1979): 574.

Savage, Leonard J. *The Foundations of Statistics*. New York: Dover Publications, Second Edition. 1972.

Sen, Amartya. “Behaviour and the Concept of Preference.” *Economica* 40, no. 159 (1973): 241-259.

---. “Description as Choice.” *Oxford Economic Papers* 32, no. 3 (1980): 353-369.

---. “Internal Consistency of Choice.” *Econometrica* 61, no. 3 (1993): 495-521.

---. “Maximization and the Act of Choice.” *Econometrica* 65, no. 4 (1997): 745-779.

Smith, Cedric. A. B. “Consistency in Statistical Inference and Decision.” *Journal of the Royal Statistical Society: Series B (Methodological)* 23, no. 1 (1961): 1-25.

Smith, Vernon L. “Microeconomic Systems as an Experimental Science.” *The American Economic Review* 72, no. 5 (1982): 923-955.

---. *Rationality in Economics: Constructivist and Ecological Forms*. Cambridge: Cambridge University Press, 2007.

Smith, Vernon L., and James M. Walker. “Monetary Rewards and Decision Cost in Experimental Economics.” *Economic Inquiry* 31 no. 2 (1993): 245-261.

Sugden, Robert. “Normative Expectations: The Simultaneous Evolution of Institutions and Norms.” In *Economics, Values and Organization*, 73-100. Edited by Avner Ben-ner and Louise Putterman. New York: Cambridge University Press, 1998.

---. *The Economics of Rights, Co-operation, and Welfare*. Hampshire: Palgrave Macmillan, 2005.

Wang, Cuizhu. “Normative Expectations and Subjective Beliefs: An Incentivised Experimental Study.” PhD thesis, University College Cork, 2022.

Winkler, Robert L., and Allan H. Murphy. “Nonlinear Utility and the Probability Score.” *Journal of Applied Meteorology and Climatology* 9, no. 1 (1970): 143-148.

Von-Neumann, John, and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press, 1944.

Xiao, Erte, and Cristina Bicchieri. “When Equality Trumps Reciprocity.” *Journal of Economic Psychology* 31, no. 3 (2010): 456-470.

Zawidzki, Tadeusz. *Dennett*. Oxford: Oneworld Publications, 2007.