Three papers on belief updating and its applications

Chao Hung Chan

Dissertation submitted to the Faculty of the

Virginia Polytechnic Institute and State University

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Economics

Sudipta Sarangi, Chair

Matthew Kovach

Hans H. Haller

Hector Tzavellas

April 11, 2024

Blacksburg, Virginia

Keywords: Non Bayesian Updating, Belief Updating, Mechanism Design

Three papers on belief updating and its applications

Chao Hung Chan

(ABSTRACT)

The normative foundation (axioms) of Bayesian belief updating has long been established in the literature of decision science. However, psychology and experiments suggest that while rational decision making is ideal, it is rarely achievable for ordinary people. Therefore, it is important to explore the foundations and consequences of rational decision making within the field of economics.

This thesis involves three papers on this. In the first paper, I explore the consequences of wishful thinking on mechanism design. It suggests that wishful thinking bias could be profit-generating for mechanism designers.

In the second paper, I investigate conservative updating and provide a foundation for it. The main behavioral axiom, "conservative consistency," suggests that decision-makers may partially incorporate information, particularly when it requires them to revise their previous preferences (the preferences order according to their prior belief).

In the third paper, I reframe the model selection problem as a rational decision-making problem. The decision-maker is restricted to choosing an advisor to delegate their choices. I explore the conditions under which a rational decision-maker selects models (or advisors) according to Bayes factor criteria.

Three papers on belief updating and its applications

Chao Hung Chan

(GENERAL AUDIENCE ABSTRACT)

Most of us do not always make decisions completely rational. This thesis digs into how irrational decision-making fits into economics, with three papers to break it down.
The first paper looks at wishful thinking and how it affects our decisions. It suggests that if we understand our biases, we can design better mechanism to generate profit.
The second paper talks about conservative updating, which is all about how we pick and choose what information matters, especially when it clashes with our existing belief.
Lastly, the third paper explores how we choose advisors to help us make decisions. It looks at when it is smart to pick based on Bayes factor criteria.
Through these papers, this thesis helps us understand how rational decision-making plays out in real-life economics.

Dedication

This thesis is lovingly dedicated to my mother, Lan Chi Wong, whose unwavering love and encouragement have been my steadfast support throughout my graduate studies.

I also dedicate this work to my partner and my family, whose enduring support and belief in me have been invaluable.

Acknowledgments

The completion of this thesis would not have been possible without the support of many individuals.

I express my deepest gratitude to my major professor, Matthew Kovach, for his invaluable support, guidance, and encouragement. His patience and considerate nature have been instrumental in my academic journey.

I extend heartfelt thanks to Hans Haller, Sudipta Sarangi, and Hector Tzavellas for their wonderful suggestions and contributions as members of my committee. Their expertise as distinguished theorists has been invaluable to my work.

My gratitude also goes to Byron Tsang, whose open-mindedness has enriched many creative discussions that have significantly contributed to my research.

Special thanks to Joe Ng and Yuan Gu, who have been not only supportive colleagues but also great friends throughout this journey. I am grateful to all the graduate students who provided companionship and support when I needed someone to talk to.

Last but not least, I want to thank all the staff in the Economics Department, especially Will Bebout. Whenever I encountered administrative challenges, he was always there to assist me.

Contents

List of Figures

List of Tables

Chapter 1

Selling to Wishful Thinkers

1.1 Introduction

Beliefs about economic variables and the behavior of others play critical roles in many economic settings. For example, an investor's decision whether or not to purchase shares in a firm depends on her beliefs about the firm's future profits. Likewise, individuals decide on their insurance levels based on their beliefs about the likelihood of accidents or health complications. It is common to assume in economic models that beliefs are "correctly calibrated," or decision makers exhibit rational expectations, in contrast to the myriad of biases observed in experimental economics.

This paper explores the implications of a particular type of belief bias, wishful thinking (WT), in auctions with independent, private values. WT is a form of optimism where the desirability of an outcome influences its perceived probability. In an auction setting, a wishful thinker misperceives the distribution of the other's values and hence she also misperceives the distribution of payoffs from the mechanism. To model

WT, we adapt [33] to an auction setting. A bidder's beliefs are "twisted" via a distortion function that transforms the true distribution over values into a subjective distribution that "shifts" probability mass from bad outcomes to good outcomes. In short, this means that a wishful bidder puts excessive probability on other bidders realizing low valuations; she believes she is more likely to win and pay a lower price conditional on winning than she should.

We start our analysis by studying the impact of WT in common auction formats. First, we show that WT induces underbidding in the standard first-price auction because bidders overestimate their chances of winning given their bid. Notably, this does not arise in the second-price auction, since the (weakly) dominant strategy is independent of a bidder's beliefs about others. Therefore, revenue equivalence does not hold for wishful bidders.

The mechanism design literature primarily assumes (subjective) expected utility preferences, which allows for additive separation of agents' expected utility into expected allocation and expected payment. Using the envelope theorem, the expected revenue of an incentive-compatible mechanism becomes a function of the expected allocation for each type (e.g., Myerson-type characterization). When beliefs are distorted due to wishful thinking expectations are not consistently additive, complicating the application of Myerson-type characterization.

Consider a scenario with two risk-neutral bidders. One is susceptible to WT, while the other is a standard (subjective) expected utility agent. For the bidder not influenced by WT, all mechanisms that yield the same expected payment and expected allocation are considered equivalent. In this case, the seller can disregard the actual distribution of payment and allocation, focusing solely on their expected values.

However, this principle does not apply to the bidder affected by WT. Let's consider a situation where both bidders are presented with two different mechanisms, each offering the same expected allocation and expected payment from the perspective of the seller. One mechanism provides winning probability and payment unconditionally, irrespective of others' reports, while the other is conditional on those reports. The wishful thinker will typically prefer the mechanism with outcomes contingent on others' reports. This preference arises because wishful thinkers tend to adjust their beliefs to favor scenarios that yield more favorable results. When outcomes hinge on factors with subjective uncertainty, like others' reports, there is room for bidders to engage in wishful thinking and imagine more favorable outcomes. This leads them to overestimate the actual utility value of a particular bid (i.e., because they overestimate their chances of winning).

From the seller's perspective, both mechanisms are equally feasible. However, the one with outcomes conditioned on others' reports has the potential to generate more surplus that can be extracted, provided all bidders report truthfully. Consequently, the mechanism that conditions outcomes on others' types will consistently yield higher revenue.

Further complicating the problem is the seller's authority over the distribution of outcomes. This control plays a pivotal role in shaping the extent to which the bidder is influenced by WT, rendering the problem intricate. Presently, the seller bears the responsibility of crafting the intricate distribution of payment and allocation. This introduces non-additivity to the expectation operator, as beliefs evolve alongside the distribution of outcomes. Intuitively, WT induces a "positive feedback" loop between behavior and beliefs. Navigating this complexity is achieved through the application of the converse envelope theorem, as proposed by [59].

Our main results reveal that sellers can exploit WT-affected bidders by employing a variant of the sad-loser auction, as first introduced by [54]. Wishful thinkers typically underestimate their chances of being the "sad loser," making this auction strategy particularly effective in leveraging their behavior.

The remainder of the paper is structured as follows. Section 2 reviews the related literature. Section 3 provides a simple motivational example to illustrate how wishful thinking can influence bidders' behavior. In Section 4, we formalize the concept and introduce relevant techniques that we employ to address the problem. Section 5 and 6 delve into the specific examination of two types of wishful thinking distortions and their impact on the action design problem. Section 7 presents a discussion of the results and their implications

1.2 Related Literature

1.2.1 Wishful Thinking

Wishful Thinking (WT) occurs when an event's desirability influences its perceived probability. The psychological literature has long acknowledged this bias [7, 25, 27]. The psychology literature suggests that WT is context-specific, and is more frequently observed in competitive situations and those involving subjective probabilities. [35].

More recently, wishful thinking has been gaining traction in economics, with experimental evidence in [39] and [18] and an axiomatic model developed by [33]. In [39], subjects were assigned roles as either *farmers* or *bakers*, and were tasked with predicting wheat prices from data. If subjects form rational beliefs, then they should make similar predictions across roles because they observed the same information.

The results, however, show that subjects' beliefs are influenced by their role, in line with wishful thinking. [18] finds evidence for wishful thinking when subjects are faced with anxiety due to future discomfort or losses. They also find that Wishful thinking is more pronounced with ambiguous information.

1.2.2 Mechanisms with non-SEU bidders

Our paper contributes to the growing literature on mechanisms designed under bounded rationality or behavioral bidders.

Most mechanism design with non-SEU bidders has focused on two specific preference theories: Maxmin Expected Utility[24] and Cumulative Prospect Theory [61]. For example, a recent addition to the literature on loss aversion in auctions includes [55]. In this paper, he considers sequential auction design and shows that loss aversion induces history-dependence and a discouragement effect.

Other forms of bounded rationality have also been considered. For instance, [22] considered the impact of projection bias in auctions. With independent private values, such bidders overbid in first-price auctions, but not in second-price ones. The focus of the paper, however, is on common value auctions. They show that in such settings second-price auctions are less efficient than first-price auctions.

1.2.3 The Sad-Loser Auctions

Our main results show that the revenue-maximizing auction for WT bidders is a variant of the sad-loser auction, first introduced by [54]. This auction type is characterized by a system where only the losers pay, while the winners receive goods free of charge. The sad-loser auction has been shown to maximize revenue in a variety of settings, including Tullock contests [8, 38, 42] and when bidders are risk loving (e.g., bidders have exponential utility)[49].

1.3 Motivating Example: Wishful Thinking and Underbidding in FPA

Consider a private value auction with two bidders who exhibit Wishful Thinking (WT). In this setting, the bidders' valuations θ are independently drawn from a uniform distribution $U[0,1]$. WT is a bias in the bidders' beliefs in which they assign excessive probability to "better" outcomes. We start with a simple way to capture this idea and suppose that each bidder "shifts" their beliefs slightly toward the best-case scenario. In this example, the best case occurs when the other bidders do not value the good at all. We can therefore model this bias using the distorted belief $G^{\delta}(\theta) = (1-\delta)G(\theta) + \delta\mathbb{1}\{\theta = 0\}$, with δ capturing the degree of bias.

Our objective is to examine the implications of WT on bidder behavior in two common auction formats: the First-Price Auction (FPA) and the Second-Price Auction (SPA). FPA involves bidders submitting sealed bids, with the highest bidder winning the item and paying their bid. In SPA, the highest bidder wins but pays the price of the second-highest bid.

To illustrate the effects of WT, we analyze the symmetric Bayesian Nash equilibrium (BNE) in both FPA and SPA, assuming a limit case where the minimum bid approaches zero. Our focus is on the bidding strategies and resulting seller revenue in the presence of WT.

$\beta(\theta)$	SEU	WT
FPA	$\frac{\theta}{2}$	$\max\{\frac{\theta}{2} - \frac{\delta}{1-\delta}, 0\}$
SPA	θ	θ

Table 1.1: Bidding Strategies in FPA and SPA

Table 1.1 summarizes the bidding strategies in FPA and SPA for both WT bidders and bidders with Subjective Expected Utility (SEU). The strategies reflect the influence of WT on bidder behavior. We observe that WT bidders tend to underbid in FPA, adjusting their bids downward due to their belief that others' valuations are lower than standard agents would believe. Conversely, in SPA, WT bidders tend to reveal their true valuations, resulting in bidding strategies that align with SEU bidders.

These findings have important implications for auction outcomes and seller revenue. In FPA, the bias induced by WT leads to reduced seller revenue compared to auctions with SEU bidders.[1] However, in SPA, the bidding strategy of revealing the true type dominates, resulting in revenue equivalence between WT bidders and SEU bidders.

By examining the impact of WT on bidder behavior and auction outcomes, we shed light on the need for alternative mechanisms. One direction is to look at seller optimal mechanisms - can the seller exploit the bias to increase their profits? In the following sections, we delve into the manipulation of beliefs and the design of optimal mechanisms, aiming to maximize seller revenue and mitigate the impact of WT-induced underbidding.

[1] Previous experimental results found a tendency to overbid in first-price sealed-bid auctions, compared with the risk-neutral Nash Equilibrium [28]. The literature offered various explanations, such as risk aversion [9], regret aversion [17], and spiteful motives [44]. However, a recent study by [47] suggests the overbidding behavior may be influenced by the information feedback process in FPA's standard setting. Our research builds upon this observation and proposes that Wishful Thinking (WT) may provide an explanation for underbidding in FPA.

1.4 Framework and Notation

This section presents the general framework and notation used throughout the paper. We consider a single good auction with n bidders subject to wishful thinking bias. The set of bidders is denoted as $N = 1, 2, ..., n$, with bidder i being the typical bidder. Each bidder i has a quasi-linear utility function defined as $u(q, t, \theta_i) = \theta_i q - t$, where q represents the probability of winning the good, t denotes the monetary payment (transfer to the mechanism), and θ_i represents bidder i's private valuation or type. The payment is positive and bounded by the budget constraint $\bar{t} > 1$. The valuations θ_i are drawn independently from a distribution F with a continuous density f over the possible types $\Theta = [0, 1]$. The distribution of types for the bidders is identical and independent, denoted as $G = F^{N-1}$ over $\Theta_{-i} = \Theta^{N-1}$.

To describe uncertainty that depends on others' types, we introduce random variables defined on the probability space $(\Theta_{-i}, \mathcal{B}^{n-1}, G)$, where $\Theta_{-i} = [0, 1]^{n-1}$ and $\mathcal{B}^{n-1}$ is the Borel σ-algebra of $\mathbb{R}^{n-1}$ restricted to $[0, 1]^{n-1}$. Capital calligraphic letters, such as $\mathcal{Q}, \mathcal{P}, \mathcal{T}$, represent sets of real-valued random variables defined on Θ_{-i}, while boldface lowercase letters, such as $\mathbf{q}, \mathbf{p}, \mathbf{t}$, denote specific elements of these sets. Furthermore, boldface capital letters, such as $\mathbf{Q}, \mathbf{P}, \mathbf{T}$, represent functions that map types to random variables defined on Θ_{-i}. For example, $\mathbf{Q} : \Theta \to \mathcal{Q}$ implies $\mathbf{Q}(\theta)(\theta_{-i})$ can be denoted as $Q(\theta, \theta_{-i})$, where $Q : \Theta \times \Theta_{-i} \to [0, 1]$.

In the presence of wishful thinking bias, bidders' beliefs about the probability of winning and payment depend on others' types θ_{-i}. This introduces uncertainty for bidder i in the form of an uncertain state. For any state-dependent utility $\mathbf{u} : \Theta_{-i} \to \mathbb{R}$, the subjective expectation is denoted as $\mathbb{E}_{-i}\mathbf{u} = \int_{\theta_{-i} \in \Theta_{-i}} \mathbf{u}(\theta_{-i}) G(d\theta_{-i})$. With wishful thinking, beliefs are distorted towards favorable events. This is captured by a positive non-decreasing distortion function δ, and the distorted expectation becomes $\mathbb{E}^{\delta}_{-i}\mathbf{u} = \int_{\theta_{-i}\in\Theta_{-i}} \mathbf{u}(\theta_{-i})\delta(\mathbf{u}(\theta_{-i}))G(d\theta_{-i})$. For any random variable $\mathbf{x}$ defined on $(\Theta_{-i}, \mathcal{B}^{n-1}, G)$, the distorted expectation $\mathbb{E}^{\delta}\mathbf{x}$ represents the expectation of $\mathbf{x}$ under the distribution G^{δ}, where $G^{\delta}(d\theta_{-i}) = \delta(\mathbf{x}(\theta_{-i}))G(d\theta_{-i})$. Further discussion on the distortion will be provided in the following subsection.

The seller aims to maximize the expected payment by selecting a direct mechanism. Each bidder i submits a report $\hat{\theta}i$ to the seller, and based on the profile of reports $\hat{\theta} = (\hat{\theta}1, ..., \hat{\theta}N)$, the mechanism determines the probability of winning and the payment for each bidder, denoted as $Q_i(\hat{\theta}), T_i(\hat{\theta})i \in N$. The mechanism is anonymous with respect to bidders, meaning that for bidder i and j, when $(\theta_i, \theta{-}i) = (\theta_j, \theta{-}j)$, we have $Q_i(\theta_i, \theta_{-i}) = Q_j(\theta_j, \theta_{-j})$ and $T_i(\theta_i, \theta_{-i}) = T_j(\theta_j, \theta_{-j})$. To emphasize the uncertainty of other types of an agent, we treat the mechanism as a pair of random variables, $\mathbf{Q}$ and $\mathbf{T}$. Both take an agent's report as input and return a random variable defined on the probability space of others' reports. The schedule $\mathbf{Q}$ maps reports, θ_i, to a real random variable $\mathbf{Q}(\theta_i)$ that defines the winning probability conditional on others report θ_{-i}. Similarly $\mathbf{T}$ maps a report, θ_i, to a real random variable $\mathbf{T}(\theta)$ that the payment conditional on others report θ_{-i}. The set $\mathcal{Q}$ denotes the measurable functions with range $[0,1]$ defined on the probability of others' types, $\mathcal{Q} = \{\mathbf{q} : \Theta_{-i} \to [0,1] | \mathbf{q} \text{ is measurable}\}$; the set $\mathcal{T}$ denotes the measurable functions with range $[0,\bar{t}]$ defined on the probability of others' type, $\mathcal{T} = \{\mathbf{t} : \Theta_{-i} \to [0,\bar{t}] | \mathbf{t} \text{ is measurable}\}$. A mechanism is denoted as $(\mathbf{Q}, \mathbf{T})$, $\mathbf{Q} : \Theta_i \to \mathcal{Q}$ and $\mathbf{T} : \Theta_i \to \mathcal{T}$. With mechanism $(\mathbf{Q}, \mathbf{T})$, a type θ_i agent reporting $\hat{\theta}_i$ will have a random utility, $\mathbf{U}(\hat{\theta}_i, \theta_i | \mathbf{Q}, \mathbf{T}) = u(\mathbf{Q}(\hat{\theta}_i), \mathbf{T}(\hat{\theta}_i), \theta_i)$. When the mechanism is clear from the context, we will omit the notation $\mathbf{Q}, \mathbf{T}$.

The mechanism is feasible when the typical Bayesian Incentive Compatible (BIC), Individual Rationality (IR), and Plausibility (P) holds.

$$\forall\theta_i,\theta_i'\in\Theta,\theta_i'\neq\theta_i: \quad \mathbb{E}^{\delta}_{-i}\{\mathbf{U}(\theta_i,\theta_i)\}\geq\mathbb{E}^{\delta}_{-i}\{\mathbf{U}(\theta_i',\theta_i)\}, \tag{BIC}$$

$$\forall\theta_i\in\Theta: \quad \mathbb{E}^{\delta}_{-i}\{\mathbf{U}(\theta_i,\theta_i)\}\geq 0, \tag{IR}$$

$$\forall(\theta_1,...,\theta_n)\in\Theta^n \quad \sum_{j\in N}Q(\theta_j,\theta_{-j})\leq 1,\ \text{and}\ \forall i\in N: Q(\theta_i,\theta_{-i})\geq 0, \tag{P}$$

where $Q(\theta_i,\theta_{-i})=\mathbf{Q}(\theta_i)(\theta_{-i})$.

1.4.1 Incorporating Wishful Thinking

We first clarify how wishful thinking is modeled in this setting. In an auction, there are two sources of uncertainty. The randomness of others' reports $\hat{\theta}_{-i}$, and given a report the mechanism may further give a randomized allocation $Q(\hat{\theta}_i,\hat{\theta}_{-i})\in(0,1)$. We assume that WT only distorts a bidder's beliefs about others' reports, not the way the mechanism works.[2]

We study two forms of WT axiomatized in [33]: the best-case binary distortion and the consequential distortion. In both cases, real random variables (acts) (e.g., $\mathbf{x}$ defined on probability space (Ω,Σ,F)) are evaluated by the distorted expectation $\mathbb{E}^{\delta}\mathbf{x}=\mathbb{E}\mathbf{x}\delta(\mathbf{x})$. The distortion δ is an increasing function that captures wishful thinking. The best-case wishful thinker increases the probability of the best-case scenario by determining other cases' probability proportionally, while the consequential wishful thinker reweighs the probability density of each event by a distortion function and

[2] This is in line with experimental evidence suggesting that wishful thinking arises more frequently in subjective or ambiguous environments (such as competitions), rather than games of objective chance (such as roulette). [35].

renormalizing afterward. The distortion function is increasing in utility to capture wishful thinking:

1. Best-case binary distortion:

$$\delta(\mathbf{U}(\hat{\theta}_i, \theta_i)(\theta_{-i})) = \begin{cases} 1-\delta & \text{if } \theta_{-i} \notin \mathcal{B}(\hat{\theta}_i, \theta_i), \\ 1-\delta+\delta\frac{1}{G(\mathcal{B}(\hat{\theta}_i, \theta_i))} & \text{if } \theta_{-i} \in \mathcal{B}(\hat{\theta}_i, \theta_i), \end{cases}$$

   $\mathcal{B}(\hat{\theta}_i, \theta_i)$ is the θ_{-i}'s that give the maximum utility for type θ_i with report $\hat{\theta}_i$. $\delta \in (0,1)$ is some constant.[3] Alternatively, one may think of the distorted belief as $G^\delta = (1-\delta)G + \delta D$, a convex combination between correct belief G and some distortion belief D different from G. In this example, D is given by a uniform distribution over "best case types," which could be discrete. With this definition, we extend the best-case binary distortion to the cases of discrete best cases.

2. Consequential distortion:

$$\delta(\mathbf{U}(\hat{\theta}_i, \theta_i)) = \frac{v(\mathbf{U}(\hat{\theta}_i, \theta_i))}{\mathbb{E}_{-i}v(\mathbf{U}(\hat{\theta}_i, \theta_i))},$$

   where the distortion function v is a continuous and increasing function.

Note that bidders have correct beliefs when they have no stake in the game (or choose a constant act). This reflects the nature of wishful thinking, which distorts "subjective" events optimistically but does not distort objective lotteries.

We only require that the distorted preference is monotonic with respect to First

[3] In its most general form, δ is not a constant but depends on the act (report) and type. As we focus on truthful reporting cases, it should depend on the type. We assume all types have the same level of distortion for simplicity here.

Order Stochastic Dominance (FOSD). Note that the utility function depends on the mechanism so that we denote the distortion as $\delta(\theta_{-i}|\hat{\theta}_i, \theta_i, Q, T) := \delta(\mathbf{U}(\hat{\theta}_i, \theta_i)(\theta_{-i}))$, and the partial derivative of it with respect to the truth type as $\delta_\theta(\theta_{-i}|\hat{\theta}_i, \theta_i, \mathbf{Q}, \mathbf{T}) = \frac{\partial \delta(\theta_{-i}|\hat{\theta}_i, \theta_i, \mathbf{Q}, \mathbf{T})}{\partial \theta_i}$. As a random variable defined on Θ_{-i} denoted as $\delta_\theta(\mathbf{U}(\hat{\theta}_i, \theta_i))$.

1.4.2 The Envelope Theorem and Its Converse

The primary tools used in this paper are the generalized envelope theorem by [41] and its generalized version with converse by [59]. We also make reference to the work of [45], which characterizes feasible mechanisms in private value auctions with quasi-linear utility.

In a private value auction with quasi-linear utility, [45] characterizes the feasible mechanisms based on the monotonicity of allocation, envelope formulas, individual rationality for the lowest type, and plausibility conditions. With these characterizations and a regularity assumption to ensure the monotonicity of allocation, the envelope theorem and individual rationality are equivalent to the expression $\mathbb{E}\mathbf{T}(\theta) = \int_{\theta'=0}^{\theta} \mathbb{E}\mathbf{Q}(\theta') d\theta' - \mathbb{E}\mathbf{Q}(\theta)\theta$. Consequently, the expected payment in the revenue maximization problem can be replaced by a function of the expected allocation.

In mechanism design, the choice set can be arbitrary, and the traditional envelope theorem does not hold. [41] provide a general version of the envelope theorem that holds for an arbitrary choice set. It is well known that outside of the quasi-linear context, the converse envelope theorem is needed to characterize feasible mechanisms. [59] provide a generalized envelope theorem with a converse counterpart. They establish an implementability theorem using this theorem, which is a generalized version of the single crossing property with increasing allocation.

In our context, given a mechanism (Q, T), a bidder with type θ_i maximizes their distorted expected utility by choosing their report. We denote the maximum utility as $V(\theta_i) = \max_{\hat{\theta}_i \in \Theta} \mathbb{E}^\delta \mathbf{U}(\hat{\theta}_i, \theta_i) = f(\hat{\theta}_i, \theta_i)$. The optimal report for type θ_i is denoted as $\hat{\theta}^*(\theta_i)$.

1.4.3 Stochastic Ordering

In order to define an increasing allocation, it is necessary to properly order the allocation space. We utilize three stochastic orders that are particularly relevant for our analysis. These orders are integral stochastic orders, which are generated by specific sets of real-valued functions denoted as $\mathcal{F}$. For random vectors $\mathbf{X}$ and $\mathbf{Y}$, the order $\mathbf{X} \geq_{\mathcal{F}} \mathbf{Y}$ holds when $\mathbb{E}[f(\mathbf{X})] \geq \mathbb{E}[f(\mathbf{Y})]$ for every function f in the set $\mathcal{F}$ for which the expectation exists.

The increasing convex order ($\leq_{icx}$) is generated by the set of all increasing convex functions. This order measures variability, as convex functions emphasize the tail of the distribution. The convex order ($\leq_{cx}$) is generated by the set of convex functions, and it is the multivariate generalization of a mean-preserving spread. The standard order ($\leq_{st}$) is the integral stochastic order generated by the set of bounded increasing functions and serves as the multivariate generalization of first-order stochastic dominance.

We list a few properties of these stochastic orders, all of which can be found in [46] and [57]:

P0 If $\mathbf{X} \leq_{cx} \mathbf{Y}$, then $\mathbf{X} \leq_{icx} \mathbf{Y}$ and $\mathbf{X} \leq_{st} \mathbf{Y}$.

P1 If $\mathbf{X} \leq_{icx} \mathbf{Y}$, there exist random vectors $\mathbf{Z}$ and $\mathbf{Z}'$ such that $\mathbf{X} \leq_{st} \mathbf{Z} \leq_{cx} \mathbf{Y}$ and $\mathbf{X} \leq_{cx} \mathbf{Z}' \leq_{st} \mathbf{Y}$.

P2 For $\mathbf{X} \leq_{icx} \mathbf{Y}$, there exist random vectors $\hat{\mathbf{X}}$ and $\hat{\mathbf{Y}}$ (defined on the same probability space) such that $\hat{\mathbf{X}} = st\mathbf{X}$, $\hat{\mathbf{Y}} = st\mathbf{Y}$, and $E\hat{\mathbf{Y}}|\hat{\mathbf{X}} \geq E\hat{\mathbf{X}}$ almost surely. This implies that $(\hat{\mathbf{X}}, \hat{\mathbf{Y}})$ forms a submartingale.

P3 For random vectors $\mathbf{X}$ and $\mathbf{Y}$:

(a) If $\mathbf{X} \leq_{icx} \mathbf{Y}$, then $\mathbb{E}[\mathbf{X}] \leq \mathbb{E}[\mathbf{Y}]$.

(b) If $\mathbf{X} \leq_{st} \mathbf{Y}$, then $\mathbb{E}[\mathbf{X}] \leq \mathbb{E}[\mathbf{Y}]$.

(c) If $\mathbf{X} \leq_{cx} \mathbf{Y}$, then $\mathbb{E}[\mathbf{X}] = \mathbb{E}[\mathbf{Y}]$.

In the next section, we demonstrate that the expectation operator is additive for the best-case distortion. As a result, the Myerson-type characterization of optimal mechanisms can be applied. Although the expectation is not additive for the consequential distortion, the insights gained from the best-case distortion still hold. In both cases, the optimal auction resembles a variant of the Sad-loser auction initially introduced by [54].

1.5 Optimal Auction with Best-case Binary Distortion

Assuming a regular distribution[4], which is common in the mechanism design literature, we establish two Lemmas that enable a Myerson-type characterization. The first lemma demonstrates the stability of the best case under a slight change in type,

[4] $\psi(\theta) = \theta - \frac{1-F(\theta)}{\theta}$, where ψ is a monotone non-decreasing function

while the second lemma shows the uniqueness of the allocation and payment within each best case. Both lemmas rely on the continuity of the utility function.

Lemma 1.1. $\mathcal{B}(\hat{\theta},\theta) = \mathcal{B}(\hat{\theta},\theta+\epsilon)$ *and* $\mathbb{E}_{-i}\mathbf{U}(\theta)\delta_\theta(\mathbf{U}(\theta)) = 0$.

Proof. Fix a report $\hat{\theta}$. By definition, $\theta_{-i} \in \mathcal{B}(\hat{\theta},\theta)$ if for every $\theta'_{-i} \notin \mathcal{B}(\hat{\theta},\theta)$, $Q(\hat{\theta},\theta_{-i})\theta - T(\hat{\theta},\theta_{-i}) > Q(\hat{\theta},\theta'_{-i})\theta - T(\hat{\theta},\theta'_{-i})$. This implies that for a small enough ϵ, $Q(\hat{\theta},\theta_{-i})(\theta+\epsilon) - T(\hat{\theta},\theta_{-i}) > Q(\hat{\theta},\theta'_{-i})(\theta+\epsilon) - T(\hat{\theta},\theta'_{-i})$. Hence, $\theta_{-i} \in \mathcal{B}(\hat{\theta},\theta+\epsilon)$. Since the best case remains unchanged for small changes in type, we conclude that $\delta_\theta \equiv 0$. □

Lemma 1.2. *For every* $\theta_{-i}, \theta'_{-i} \in \mathcal{B}(\hat{\theta},\theta)$, $Q(\hat{\theta},\theta_{-i}) = Q(\hat{\theta},\theta'_{-i}) =: \bar{Q}(\hat{\theta},\theta)$ *and* $T(\hat{\theta},\theta_{-i}) = T(\hat{\theta},\theta'_{-i}) =: \bar{T}(\hat{\theta},\theta)$.

Proof. For $\theta_{-i}, \theta'_{-i} \in \mathcal{B}(\hat{\theta},\theta)$, we have $Q(\hat{\theta},\theta_{-i})\theta - T(\hat{\theta},\theta_{-i}) = Q(\hat{\theta},\theta'_{-i})\theta - T(\hat{\theta},\theta'_{-i})$. The previous Lemma shows $\mathcal{B}(\hat{\theta},\theta) = \mathcal{B}(\hat{\theta},\theta+\epsilon)$, which also implies that $Q(\hat{\theta},\theta_{-i})(\theta+\epsilon) - T(\hat{\theta},\theta_{-i}) = Q(\hat{\theta},\theta'_{-i})(\theta+\epsilon) - T(\hat{\theta},\theta'_{-i})$. Combining both conditions gives $Q(\hat{\theta},\theta_{-i}) = Q(\hat{\theta},\theta'_{-i})$ and $T(\hat{\theta},\theta_{-i}) = T(\hat{\theta},\theta'_{-i})$. We denote them as $\bar{Q}(\hat{\theta},\theta), \bar{T}(\hat{\theta},\theta)$. When $\hat{\theta} = \theta$, we denote them as $\bar{Q}(\theta), \bar{T}(\theta)$. □

With these two Lemmas, we can apply a Myerson-type characterization of optimal mechanisms to auctions with best-case binary distorted bidders. We define the distorted expected winning probability as $q(\theta_i) = (1-\delta)\mathbb{E}_{-i}\mathbf{Q}(\theta_i) + \bar{Q}(\theta_i)\delta$. Similarly, we define the distorted expected payment as $t(\theta_i) = (1-\delta)\mathbb{E}_{-i}\mathbf{T}(\theta_i) + \bar{T}(\theta_i)\delta$. Consequently, $V(\theta) = q(\theta)\theta - t(\theta)$. We can rewrite the BIC condition as $V(\theta) = V(0) + \int_0^\theta q(x)dx$. IR and profit maximization require $V(0) = 0$. Expanding $V(\theta)$

gives:

$$t(\theta_i) = q(\theta_i)\theta_i - \int_0^{\theta_i} q(x)dx. \tag{1.1}$$

Equation (1.1) summarizes BIC and IR.

The seller's objective is to maximize the expected revenue. Rearranging $t(\theta_i)$ and combining it with equation (1.1) gives the expected payment for type θ_i:

$$\mathbb{E}_{-i}\mathbf{T}(\theta_i) = \mathbb{E}_{-i}\mathbf{Q}(\theta_i)\theta_i - \int_{x=0}^{\theta_i} \mathbb{E}_{-i}\mathbf{Q}(x)dx + \frac{\delta}{1-\delta}\left(\bar{u}(\theta_i) - \int_{x=0}^{\theta_i} \bar{Q}(x)dx\right),$$

where $\bar{u}(\theta) = \bar{Q}(\theta)\theta - \bar{T}(\theta)$. The expected revenue from a bidder i is given by:

$$\begin{aligned}\int_{\theta_i=0}^{1} \mathbb{E}_{-i}\mathbf{T}(\theta_i)F(d\theta_i) = &\int_{\theta_i=0}^{1} \mathbb{E}_{-i}\mathbf{Q}(\theta_i)\psi(\theta_i)F(d\theta_i)\\ &+ \frac{\delta}{1-\delta}\int_{\theta_i=0}^{1}(\bar{Q}(\theta_i)\psi(\theta_i) - \bar{T}(\theta_i))F(d\theta_i),\end{aligned}$$

where $\psi(\theta_i) = \theta_i - \frac{1-F(\theta_i)}{f(\theta_i)}$ is the virtual valuation when the bidder is unbiased.

Observation 1. When the transfer to the mechanism is unbounded, the seller can achieve arbitrarily large revenue. In other words, as $\bar{T}(\theta_i) \to -\infty$, $\int_{\theta_i=0}^{1} \mathbb{E}_{-i}\mathbf{T}(\theta_i)dF(\theta_i) \to \infty$.

This is because the mechanism designer can directly offer the bidder a chance to place a bet on another player's type. Since bidders have biased beliefs, the mechanism designer can achieve unbounded profit by offering an unfair lottery to the biased bidders, which may have nothing to do with the good itself. To focus on the incentive for obtaining the good and exclude such cases, we restrict the transfer to be positive, i.e., $\forall\theta \in \Theta^N : T(\theta) \geq 0$.

The expected revenue is decreasing in $\bar{T}(\theta_i)$. The optimal mechanism has $\bar{T}(\theta_i) = 0$. This provides a Myerson-type characterization of payment. The expected revenue per bidder is:

$$\int_{\theta_i=0}^{1} \mathbb{E}_{-i}\mathbf{T}(\theta_i)F(d\theta_i) = \int_{\theta=0}^{1} \left(\mathbb{E}_{-i}\mathbf{Q}(\theta_i) + \frac{\delta}{1-\delta}\bar{Q}(\theta_i) \right) \psi(\theta_i)F(d\theta_i). \tag{1.2}$$

Observation 2. The first term, $\int_{\theta=0}^{1} \mathbb{E}_{-i}\mathbf{Q}(\theta)\psi(\theta)F(d\theta)$, represents the revenue when the bidders are unbiased. Therefore, bias increases the seller's revenue. The second term, $\int_{\theta=0}^{1} \frac{\delta}{1-\delta}\bar{Q}(\theta)\psi(\theta)F(d\theta)$, represents the surplus induced by the WT bias. It is increasing in δ. As a result, the more biased the bidders are, the more revenue the seller can generate.

Observation 3. There is no restriction on $\mathcal{B}(\theta_i)$.

As the type distribution is regular, a simple solution is as follows:

$$\begin{aligned} Q_i(\theta_i, \theta_{-i}) &= 1 && \Leftrightarrow && \forall j \neq i : \theta_i > \theta_j, \text{ and } \psi(\theta_i) \geq 0, \\ Q_i(\theta_i, \theta_{-i}) &= \frac{1}{\bar{N}} && \Leftrightarrow && \forall j \neq i : \theta_i \geq \theta_j, \text{ and } \psi(\theta_i) \geq 0, \\ Q_i(\theta_i, \theta_{-i}) &= 0 && \Leftrightarrow && \forall j \neq i : \theta_i < \theta_j, \text{ or } \psi(\theta_i) < 0, \end{aligned}$$

If $\psi(\theta_i) < 0$, the bidder never wins and must have zero transfer, so the best case is $\mathcal{B}(\theta_i, \theta_i) = \Theta_{-i}$. For other cases, since there is no restriction on $\mathcal{B}(\theta_i, \theta_i)$, one can freely choose some $\mathcal{B} \subset [0, \theta_i)^{N-1}$ to make it the best case. The payment is given by equation (1.1) with $\bar{T}(\theta) = 0$. Here, $\bar{N}$ is the number of bidders submitting the same highest report. When there are multiple bids with positive virtual valuations, the good is allocated to the bidder with the highest valuation.

To see that the optimal mechanism has a loser-pay feature, consider a bidder i. In the best case, bidder i obtains the good for free and only has to pay in other cases. Thus, this is a sad-loser auction in the best case, which is the case where the bidder has wishful thinking and distorted beliefs.

1.5.1 Implementation: SPAr with Sad-loser Lottery

A Second Price Auction (SPA) with a reservation price and an additional lottery can easily implement the above mechanism. For simplicity, let's consider the case of 2 bidders with uniformly distributed valuations $\theta \sim U[0,1]$, where types are distributed independently. In this case, the virtual valuation is $\psi(\theta_i) = 2\theta_i - 1$. We choose the best case as the other bidder reporting zero, $\mathcal{B}(\theta_i, \theta_i) = 0$.

Consider a Second-Price Auction with a reservation price of $\frac{1}{2}$, where the seller withholds the good if either bidder bids zero. In SPAr, revealing the truth is still weakly dominant. The revenue in auctions with Expected Utility Theory (EUT) bidders is given by:

$$\Pi^{SEU} = \sum_{i \in \{1,2\}, j \neq i} \left(\int_0^1 \int_0^1 (2\theta_i - 1) Q_i(\theta_i, \theta_j) d\theta_j d\theta_i \right).$$

It is worth noting that the optimal mechanism for auctions with EUT bidders can be implemented by an SPA with a reservation price of $\frac{1}{2}$, denoted as $\Pi^{\text{EUT}*} = \Pi^{\text{SPAr,EUT}}$. The SPAr, where the seller withholds the good when either bidder has a zero valuation, yields the same revenue, i.e., $\Pi^{\text{SPAr,WT}} = \Pi^{\text{SPAr,EUT}}$. Equation (1.2) shows that the revenue could be improved by:

$$\Pi^{WT} - \Pi^{SEU} = \sum_{i \in \{1,2\}} \left(\frac{\delta}{1-\delta} \int_0^1 (2\theta_i - 1) Q(\theta_i, 0) d\theta_i \right).$$

The optimal mechanism gives $Q(\theta_i, 0) = 1$ if $\theta_i > \frac{1}{2}$ and zero otherwise. Thus, the potential improvement is $\sum_{i\in 1,2}\frac{\delta}{4(1-\delta)}$. To extract this surplus, an extra lottery is offered to each bidder.

Recall that the seller withholds the good if either bidder has a zero valuation. The seller could offer a sad-loser lottery $L(x)$ to bidders conditional on this event. The lottery holder obtains the good for free if the other bidder bids zero (the best case) and pays x otherwise. Bidders with valuations greater than $\frac{(1-\delta)}{\delta}x$ accept the lottery. The lottery represents a bet on a null event. Whenever the lottery is accepted, the revenue is improved by x. For each bidder, the expected revenue from the lottery is:

$$\Pi^{L(x)} = \int_{\frac{(1-\delta)}{\delta}x}^{1} x d\theta = x - \frac{1-\delta}{\delta}x^2.$$

The revenue is maximized when $x = \frac{\delta}{2(1-\delta)} := x^*$. The expected revenue for each bidder is $\Pi^{L(x^*)} = \frac{\delta}{4(1-\delta)}$. All possible improvements have been explored. This proves that the SPA with a reservation price of $\frac{1}{2}$ and the additional sad-loser lottery $L(x^*)$ is revenue maximizing.

1.5.2 The Mechanism Behind the Sad-loser Lottery

This section explores how the sad-loser lottery improves the seller's revenue, particularly in the context of wishful thinking (WT) bidders. The lottery is profitable when bidders have biased beliefs, and the revenue is higher when the payment is conditional on loss, an underweighted event.

The figure below illustrates the lottery from the perspective of wishful thinking bidders, Subjective Expected Utility (SEU) maximizing bidders, and the seller. For wishful thinkers, the perceived gain of the lottery is $\delta\theta_i > 0$, and the perceived cost is $(1-\delta)x$. Bidders with valuations θ_i greater than $\frac{\delta x}{(1-\delta)}$ find the lottery profitable and accept it. EUT bidders, on the other hand, do not find the lottery profitable and reject the offer. Since the actual probability of winning is zero, it is free for the seller to provide the lottery. Anyone who accepts the lottery increases the seller's revenue by x.

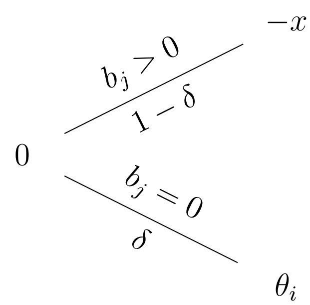

Figure 1.1: Perceived payoff for WT bidders

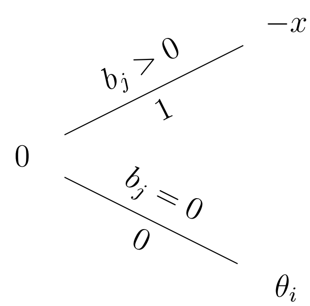

Figure 1.2: Perceived payoff for WT bidders

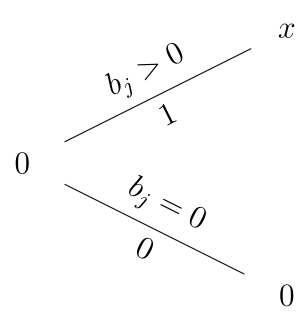

Figure 1.3: Payoff of the seller

Separating the case of receiving the good from paying the lottery fee always increases surplus, with the subjective expected payment remaining unchanged. SEU bidders are indifferent between paying x conditional on loss or unconditionally. However, wishful thinkers prefer paying x conditional on loss since they underweight the probability of loss. There is a perceived discount of δ on the payment when it is separated from the best case. If the lottery's payment is made unconditionally, the lottery only yields the seller an expected revenue of $\frac{\delta}{4}$ per bidder, which is suboptimal. A similar effect also arises in the winning case. Wishful thinkers have a biased perceived gain of δx from the sad-loser lottery as they overestimate the winning probability. Thus, the sad-loser feature improves revenue for mechanisms targeting wishful thinking bidders. The next section demonstrates that this intuition remains valid in the more general case of consequential distortion.

The mechanism is robust in markets with both best-case binary WT and SEU bidders. SEU bidders reject any lottery that is not in their favor, while wishful thinkers accept the lottery. We separate WT bidders from EUT bidders, even if the bidders themselves fail to recognize their bias. However, to offer the optimal sad-loser lottery, the seller must know the parameter δ.

1.6 Consequential Wishful Thinking

This section examines the auction problem with wishful thinking (WT) modeled by the more general consequential distortion. We demonstrate that under a mild condition on the shape of the distortion function v, the optimal mechanism is a variant of the loser-pay auction, and the bias increases the seller's revenue.

We require the distortion function to satisfy the following condition:

$$\forall u \leq 1 : v''(u) \leq 2v'(u) + uv''(u). \tag{1.3}$$

This condition ensures that the agent's preference is monotonic with respect to first-order stochastic dominance (FOSD).[5] It also guarantees a convex distorted expected utility.

Lemma 1.3. *The distorted preference satisfies FOSD when*

$$\forall u \leq 1 : v'(u) \leq v(u) + uv'(u). \tag{1.4}$$

The condition (1.4) is always satisfied when $u = 1$ since $v > 0$. By the Fundamental Theorem of Calculus, (1.3) can be seen as a smooth version of (1.4).

Before delving into the problem of the optimal auction with wishful bidders, let us review some basic properties of consequential distortion. Based on these properties, we will define the maximum spread auction, which intuitively should be the optimal auction. Finally, we will prove that it is indeed the case under certain assumptions.

First, we observe that wishful thinkers prefer a spread.

Lemma 1.4. *Given a set of acts (reports) with the same subjective (unbiased) expected payoff, wishful thinkers prefer the one with a greater spread.*

Proof. Define the function $\phi(u) = uv(u)$. Then, the distorted expectation for a random variable (act) $\mathbf{u}$ can be written as $\frac{\mathbb{E}\phi(\mathbf{u})}{\mathbb{E}v(\mathbf{u})}$. Approximating the value of the distorted expected value around the mean gives:

[5] For random variables X and Y, if X first-order stochastically dominates Y ($X >_{st} Y$), then the agent prefers X over Y ($X \succ Y$).

$$\mathbb{E}^{\delta}\mathbf{u} = \frac{\mathbb{E}\phi(\mathbf{u})}{\mathbb{E}v(\mathbf{u})} \approx \frac{\phi(\mathbb{E}\mathbf{u}) + \phi''(\mathbb{E}\mathbf{u})Var(\mathbf{u})}{v(\mathbb{E}\mathbf{u}) + v''(\mathbb{E}\mathbf{u})Var(\mathbf{u})}.$$

Since (1.3) ensures $\phi'' > v''$, for two random variables $\mathbf{u}$ and $\mathbf{u}'$ with $\mathbf{u}'$ being the mean-preserved spread of $\mathbf{u}$, $Var(\mathbf{u}') > Var(\mathbf{u})$. Thus, $\mathbb{E}^{\delta}\mathbf{u}' > \mathbb{E}^{\delta}\mathbf{u}$. □

The agent's beliefs are increasingly biased toward events with good outcomes. Whenever there is a spread between good and bad events, while keeping the expected outcome unchanged, the agent will believe that she is more likely to win and likely to pay less. Since the maximum utility for a type θ agent is θ, and the minimum payoff is $-\bar{t}$, a maximum exists for every distribution of states (others' types).

Figure 1 provides an intuitive explanation. The black curve represents the state-dependent random utility $u(\theta_{-i})$, and the red line depicts the mean-preserved deviation $u'(\theta_{-i})$ of $u(\theta_{-i})$. Suppose that for every $\theta^{+}_{-i} \in \Theta^{+}_{-i}$ and $\theta^{-}_{-i} \in \Theta^{-}_{-i}$, $u(\theta^{+}_{-i}) > u(\theta^{-}_{-i})$. By increasing $u(\theta_{-i})$ in Θ^{+}_{-i} by a fixed amount Δ while preserving the mean, we can reduce $u(\theta_{-i})$ in Θ^{-}_{-i} by $\Delta\frac{G(\Theta^{-}_{-i})}{G(\Theta^{+}_{-i})}$. Since $u'(\theta_{-i})$ in Θ^{+}_{-i} is increased, while $u'(\theta_{-i})$ in Θ^{-}_{-i} is decreased, the distorted agent further increases the distorted probability assessment on Θ^{+}_{-i} and decreases those on Θ^{-}_{-i}. As a result, the distorted utility for u' is greater than u. By continuing to increase the distorted utility through the selection of two regions and increasing the spread, two possible cases emerge. The support of the utility function is either $\theta, 0, -\bar{t}$ or $\theta, \theta - \bar{t}, -\bar{t}$. Figures 2 and 3 illustrate these two cases.

This provides intuition as to why the loser-pay auction is optimal in the previous section. Naturally, it offers the maximum spread to the agent. The seller has no preference for the distribution of the winning probability and the payment for any report. At the same time, the bidder prefers the one that offers the greatest spread to the generated random utility. We consider the case where some report does not provide the maximum spread to the bidder as a potential for Pareto improvement. The seller could extract the surplus from the improvement with an appropriate mechanism. Thus, we formally define the maximum spread report and the maximum spread auction and then show their optimality under certain additional assumptions.

Definition 1.5 (Maximum spread report)**.** Given a mechanism $\mathbf{Q}, \mathbf{T}$, a report θ is a maximum spread report when the random utility it generates for the corresponding type, $\mathbf{U}(\theta) = \theta\mathbf{Q}(\theta) - \mathbf{T}(\theta)$, is $\geq_{cx}$-maximum among the mechanisms $\mathbf{Q}', \mathbf{T}'$ that give the same expected winning probability and payment for θ, i.e., $\mathbb{E}_{-i}\mathbf{Q}'(\theta) = \mathbb{E}_{-i}\mathbf{Q}(\theta)$ and $\mathbb{E}_{-i}\mathbf{T}'(\theta) = \mathbb{E}_{-i}\mathbf{T}(\theta)$.

The maximum is well-defined as the support is bounded.

Definition 1.6 (Maximum spread auction)**.** A mechanism is called a maximum spread auction when all reports are maximum spread reports.

This means that there are functions $\alpha, \beta : \Theta \to [0, 1]$ such that

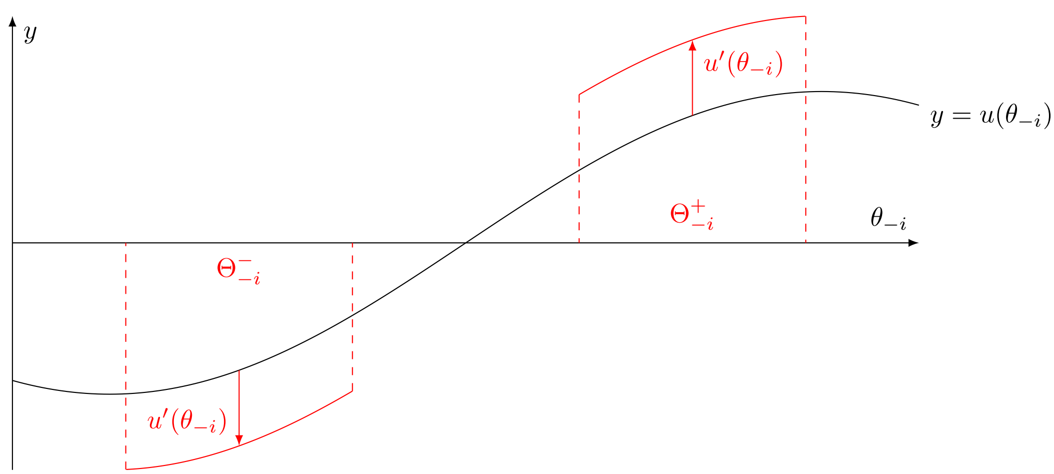

Figure 1.4: Mean-preserved Spread of u that increases the distorted expectation.

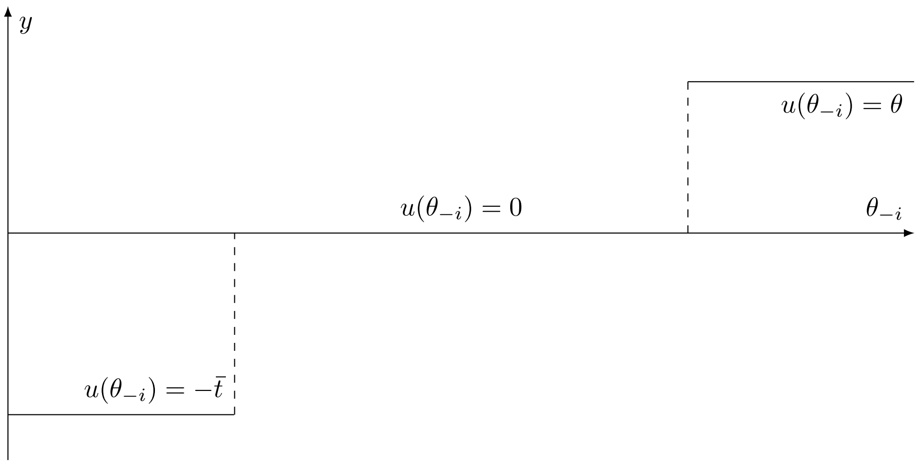

Figure 1.5: Case 1: Support of the utility function is $\theta, 0, -\bar{t}$.

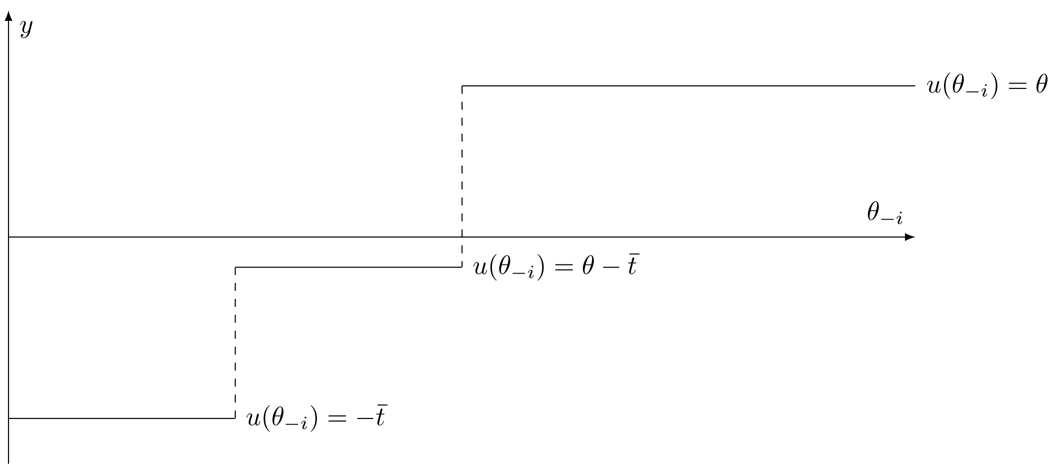

Figure 1.6: Case 2: Support of the utility function is $\theta, \theta - \bar{t}, -\bar{t}$.

$$G[\theta_{-i} : Q(\theta, \theta_{-i}) = 1 \text{ and } T(\theta, \theta_{-i}) = 0] = 1 - \beta(\theta),$$

$$G[\theta_{-i} : Q(\theta, \theta_{-i}) = 0 \text{ and } T(\theta, \theta_{-i}) = \bar{t}] = 1 - \alpha(\theta),$$

$$G[\theta_{-i} : Q(\theta, \theta_{-i}) = 1 \text{ and } T(\theta, \theta_{-i}) = \bar{t}] = (\alpha(\theta) + \beta(\theta) - 1)^+,$$

$$G[\theta_{-i} : Q(\theta, \theta_{-i}) = 0 \text{ and } T(\theta, \theta_{-i}) = 0] = (1 - \alpha(\theta) - \beta(\theta))^+,$$

where $(z)^+ = \max z, 0$. Here, $\alpha(\theta)$ is the probability of getting the goods, and $\beta(\theta)$

is the probability of paying the whole budget.

A necessary condition is that for every $\theta_i \in \Theta, \theta_{-i} \in \Theta_{-i}$:

$$Q(\theta_i, \theta_{-i}) \in \{0, 1\}$$
$$T(\theta_i, \theta_{-i}) \in [0, \bar{t}].$$

The maximum spread auction is a variation of the sad-loser auction, which, in turn, is a variation of an all-pay auction. Our main theorem shows that it is optimal when the budget is large, and the expected winning probability schedule is increasing.

Theorem 1.7 (Main Theorem)**.** *When the budget is large and the expected winning schedule,* $\mathbb{E}_{-i}\mathbf{Q}$*, is an increasing function, the optimal mechanism is of maximum spread.*

An immediate corollary of the theorem is the following.

Corollary 1.8 (Main Theorem (Alternative))**.** *When the budget is large, the optimal efficient mechanism is of maximum spread.*

To prove the main theorem, we will show that when the allocation schedule increases, individual rationality (IR) is met for the lowest type, and the envelope formula is satisfied and plausible, then the mechanism is feasible. We will mainly demonstrate that IR for the lowest type implies IR for all types, and we define the allocation space and order in such a way that the Spence-Mirrlees condition implies that when the envelope formula holds, the Bayesian incentive compatibility (BIC) condition also holds.

Proposition 1.9. *When condition (1.3) holds, for any mechanism that satisfies BIC,* $V(0) \geq 0 \Rightarrow \forall \theta \in \Theta : V(\theta) \geq 0$ *(IR is satisfied).*

Proof. Note that the same report from a higher type first-order stochastically dominates the same report from lower types. By Lemma 2, (1.3) implies that for every type θ_i, $\mathbb{E}^{\delta}_{-i}\mathbf{U}(0,\theta_i) \geq \mathbb{E}^{\delta}_{-i}\mathbf{U}(0,0)$. BIC requires $\mathbb{E}_{-i}\mathbf{U}(\theta_i,\theta_i) \geq \mathbb{E}_{-i}\mathbf{U}(0,\theta_i)$. As V is increasing, $V(0) \geq 0 \Rightarrow \forall\theta \in \Theta : V(\theta) \geq 0$. □

Proposition 1 replicates the result in [45] that the IR constraint only restricts the initial value of the envelope formula. Next, we construct the implementability theorem as in [59] in our context. With the implementability theorem, any increasing allocation is implementable: there is a payment schedule such that the mechanism is BIC. The payment schedule and the allocation must jointly satisfy the envelope formula.

Next, we give an implementability theorem. The implementability theorem isolates payment, which the seller cares about, from allocations, for which the bidder has a preference. In our context, the bidders have a preference over the distribution of payment. Thus, we extend the allocation space to include the distribution of both payment and winning probability. For a mechanism $\mathbf{Q}, \mathbf{T}$, we separate the expected payment from its distribution by defining functions $\tau, \mathbf{P}$ from the payment schedule $\mathbf{T}$. $\tau : \Theta \to [0,k]$ is the expected payment schedule, $\tau = E\mathbf{T}$. $\mathbf{P} : \Theta \to \mathbb{R}^+$ is the proportion of the expected payment schedule, $\mathbf{P} = \frac{\mathbf{T}}{\tau}$. When $\tau = 0$, which implies $\mathbf{T} = 0$, we define $\mathbf{P} = 0$. The allocation schedule of a mechanism is denoted as $\mathbf{Y} = (\mathbf{Q}, \mathbf{P})$. Thus, a mechanism can be written as $\mathbf{M} = (\mathbf{Q}, \mathbf{T}) = (\mathbf{Q}, \mathbf{P}, \tau) = (\mathbf{Y}, \tau)$. The allocation space and payment space are defined as $\mathcal{Y} = \mathcal{Q} \times \mathcal{P}$ and $\mathcal{P}$, respectively.

We order the allocation space $\mathcal{Y}$ by the following order $\geq_{\mathcal{Y}}$. For $\mathbf{y}, \mathbf{y}' \in \mathcal{Y}$, with $\mathbf{y} = (\mathbf{q}, \mathbf{p})$ and $\mathbf{y}' = (\mathbf{q}', \mathbf{p}')$, we have $\mathbf{y}' \geq_{\mathcal{Y}} \mathbf{y}$ if $(\mathbf{p}', -\mathbf{t}') \geq_{icx} (\mathbf{p}, -\mathbf{t})$.

The implementability theorem of [59] suggests that if $\mathcal{Y}$ is regular and the payoff function f satisfies the outer Spence-Mirrlees condition, then any increasing allocation is implementable. [6] [7] [8] [9]

The Spence-Mirrlees condition is a single crossing property on the payoff function that interprets as the higher type is more willing to pay for an increase in allocation. In the appendix, we show that if we focus on a restricted set of mechanisms, the outer Spence-Mirrlees condition holds.

The restricted set of the mechanism is such that for every $\theta' > \theta$, $\mathbf{Y}(\theta') = (\mathbf{q}', \mathbf{p}'), \tau(\theta') = t', \mathbf{Y}(\theta) = (\mathbf{q}, \mathbf{p}), \tau(\theta) = t$, $\hat{\mathbf{U}}(\theta) = \mathbf{q}\theta - t\mathbf{p}$:

$$\mathbb{E}_{-i}\{\phi'(\hat{\mathbf{U}}(\theta))(\mathbf{q}' - \mathbf{q})\} \geq \mathbb{E}_{-i}\{[\phi'' - v''](\hat{\mathbf{U}}(\theta))\mathbf{q}\mathbf{p}\}(t' - t) \text{ for every } \theta \in [0, 1]. \quad (1.5)$$

To see why this restriction is needed, we could focus on the right-hand side of the inequality. When $\mathbf{q}$ and $\mathbf{p}$ have separated support, the right-hand side is equal to zero. (1.5) holds by the property O2. When the support for payment and winning is mixed, there are two effects of increasing type on the willingness to pay. First, the higher types have a greater valuation for winning, thus giving a more distorted probability assessment and expected value. On the other hand, when the support for

[6] The outcome space $\mathcal{Y}$ is *regular* iff it is order-dense-in-itself, countably chain-complete and chain-separable.

[7] For payoff function f that maps allocation, y, payment, p, and type, t, to payoff, $f(y, p, t)$. f is *regular* iff (i) the type derivative f_3 exists and is bounded, and $f_3(y, \cdot, t)$ is continuous for each $y \in \mathcal{Y}$ and $t \in [0, 1]$, and (ii) for every chain $\mathcal{C} \subseteq \mathcal{Y}$, f is jointly continuous on $\mathcal{C} \times \mathbb{R} \times [0, 1]$ when $\mathcal{C}$ has the relative topology inherited from the order topology on $\mathcal{Y}$.

[8] f satisfies the (strict) *outer Spence-Mirrlees condition* iff for any increasing $\mathbf{Y} : [0, 1] \to \mathcal{Y}$, any $\tau : [0, 1] \to \mathbb{R}$ and any $r < t$ in $(0, 1)$,

$$n \mapsto \frac{\bar{d}}{\bar{d}m} \int_r^t f(\mathbf{Y}(s + m), \tau(s + m), s + n)ds \bigg|_{m=0}$$

is (strictly) single-crossing, where $\bar{d}/\bar{d}m$ denotes the upper derivative.

[9] An allocation $\mathbf{Y} : [0, 1] \to \mathcal{Y}$ is *implementable* iff there is a payment schedule $\tau : [0, 1] \to \mathbb{R}$ such that $(\mathbf{Y}, \tau)$ is incentive-compatible. An increasing allocation is one that provides higher types with larger outcomes (in the partial order on $\mathcal{Y}$).

payment and winning are overlapped, higher types have greater utility for the region of payment, thus increasing the distorted expected payment. As a result, the aggregate effect is ambiguous. The second effect reduces when the payment and winning are less aligned, $\mathbb{E}_{-i}\mathbf{qp}$ reduces. We say the mechanism has sufficiently separated regions of winning and payment if (1.5) holds. This gives the following implementability theorem.

Proposition 1.10. *Any mechanism with sufficiently separated regions of winning and payment with increasing allocation is implementable.*

With the implementability theorem in hand, we could fix any increasing allocation and look for the expected payment schedule using the envelope formula. When 1.5 holds, the result mechanism must be BIC.

Next, we address the problem of determining which mechanism with an increasing allocation is optimal. We aim to identify profitable deviations for feasible mechanisms that are not of maximum spread. Since agents prefer spread, our target deviation is to increase the spread for the mechanism, which we consider a Pareto improvement. We then seek a mechanism to extract the surplus generated by this improvement.

We begin by demonstrating that for any increasing allocation with some reports not of maximum spread, we can replace those reports with reports of greater spread without altering the monotonicity of the allocation. Thus, the revised allocation remains implementable.

Lemma 1.11. *For an increasing allocation with some reports not of maximum spread, there exists another increasing allocation with increased spread for those reports.*

Proof. Suppose the increasing mechanism is $\mathbf{y_1} \leq_{icx} \mathbf{y_2} \leq_{icx} \mathbf{y_3}$ with $\mathbf{y_2}$ not of maximum spread. By property O1 (define if necessary), there exists $\mathbf{z} \neq_{st} \mathbf{y_2}$ such that $\mathbf{y_1} \leq_{icx} \mathbf{y_2} \leq_{cx} \mathbf{z} \leq_{st} \mathbf{y_3}$. Since $\mathbf{y_1} \leq_{icx} \mathbf{z} \leq_{st} \mathbf{y_3}$, there exists $\mathbf{z}'$ such that $\mathbf{y_1} \leq_{cx} \mathbf{z}' \leq_{st} \mathbf{z} \leq_{st} \mathbf{y_3}$. Note that there is an interval between $\mathbf{z}$ and $\mathbf{z}'$ in the *st* order where we can replace $\mathbf{y_2}$ with some $\mathbf{y_2'}$ between $\mathbf{z}$ and $\mathbf{z}'$ in the *st* order. We choose $\mathbf{y_2'}$ to have a higher spread than $\mathbf{y_2}$, i.e., $\mathbf{y_2'} \geq_{cx} \mathbf{y_2}$. □

Next, we show that when we disregard the monotonicity constraints, the seller would always want to offer a mechanism with maximum spread.

Lemma 1.12. *If $\mathbf{U}(\theta|Q',T') \geq_{cx} \mathbf{U}(\theta|Q,T)$, then $f_\theta(\theta|Q',T') > f_\theta(\theta|Q,T)$, and $f_\theta(\theta)$ is an increasing function.*

The lemma shows that when we replace a non-maximum spread report with one that has more spread, both the indirect utility and its slope increase.

Lemma 1.13. *$V(\theta|\mathbf{Q},\mathbf{T})$ and $f_\theta(\theta|\mathbf{Q},\mathbf{T})$ are increasing in the probability of winning schedule $\mathbf{Q}(\theta)$ and decreasing in the payment schedule $\mathbf{T}(\theta)$, both ordered by first-order stochastic dominance (the st order). For changes in $\mathbf{Q}$ or $\mathbf{T}$ that have the same impact on the resulting random utility, $\mathbf{U}(\theta|\mathbf{Q}',\mathbf{T}) = \mathbf{U}(\theta|\mathbf{Q},\mathbf{T}')$, the relative impact on V and f_θ is different.*

For a fixed θ, let $\mathbf{q} = \mathbf{Q}(\theta)$ and $\mathbf{t} = \mathbf{T}(\theta)$. $V(\mathbf{q},\mathbf{t}) = V(\theta)$ and $f_\theta(\mathbf{q},\mathbf{t}) = f_\theta(\theta)$. For a functional F that takes functions a and b as inputs, let η be the derivative of F with respect to a, i.e., $F_a^\eta(a,b) = \lim_{\epsilon\to 0} \frac{F(a+\epsilon\eta,b)-F(a,b)}{\epsilon}$. For every $\eta^q \leq 0$ and $\eta^t \geq 0$:

$$
\begin{aligned}
V_q^{\eta^q} &\leq 0, & f_{\theta,q}^{\eta^q}(\mathbf{q},\mathbf{t}) &\leq 0,\\
V_t^{\eta^t} &\leq 0, & f_{\theta,t}^{\eta^t}(\mathbf{q},\mathbf{t}) &\leq 0.\\
\text{For every } \eta : \Theta_{-i} \to \mathbb{R}^+, & & \frac{V_q^{\theta\eta}}{f_{\theta,q}^{\theta\eta}}(\mathbf{q},\mathbf{t}) &\neq \frac{V_t^{-\eta}}{f_{\theta,t}^{-\eta}}(\mathbf{q},\mathbf{t})
\end{aligned}
$$

Lemma 7 suggests that when we replace a non-maximum spread report with one that has more spread, the value function increases. Lemma 8 shows that we can reduce it back to the original value function by decreasing the probability of winning or increasing the payment in the *st* order. In the first case, we have some free winning probability to assign. We can assign it to the report that has the greatest payment, giving more revenue to the seller as the report with the greatest payment becomes more attractive. In the second case, it is obvious that the deviation is profitable. In sum, Lemmas 7 and 8 show that when both mechanisms are feasible, the optimal mechanism is the one with maximum spread.

We return to the implementability condition that requires an increasing allocation and (1.5) to hold. Given that the maximum spread auction is feasible when the budget is large and the expected winning probability schedule is increasing, (1.5) naturally holds. Also, note that the maximum spread auction with an increasing expected winning schedule implies an increasing allocation, which proves our main theorem.

In conclusion, the main theorem states that when the budget is large and the expected winning schedule is an increasing function, the optimal mechanism is of maximum spread. The proof is feasible because if a mechanism satisfies the envelope formula, and we have another mechanism that gives the same values of $V(\theta), f_\theta(\theta)_{\theta\in\Theta}$, the alternative mechanism also satisfies the same envelope formula. We find a process in which the revenue increases while keeping the envelope formula unchanged, and we show that the optimal mechanisms from the previous process are implementable.

Figure 4 provides a conceptual understanding of the proof. The figure plots the expected utility derived from a fixed report $\hat{\theta} = \theta$ as a function of the true type Θ. Given a mechanism $(\mathbf{Q}, \mathbf{T})$ which yields a non-maximum spread report θ, the distorted expected utility (and similarly, the slope) is augmented by altering the report θ to a maximum spread one in $(\mathbf{Q}', \mathbf{T}')$. The seller can then seize the surplus by suitably decreasing allocation and increasing payment such that the new mechanism $(\mathbf{Q}'', \mathbf{T}'')$ delivers the same expected utility function for report θ as found in the original mechanism $(\mathbf{Q}, \mathbf{T})$. Consequently, the initial envelope formula applies, and revenue surges, demonstrating that when the maximum spread auction is viable, it is the optimal choice. The proof is finalized by illustrating that if the budget is ample and the expected winning probability schedule is ascending, the maximum spread auction is feasible.

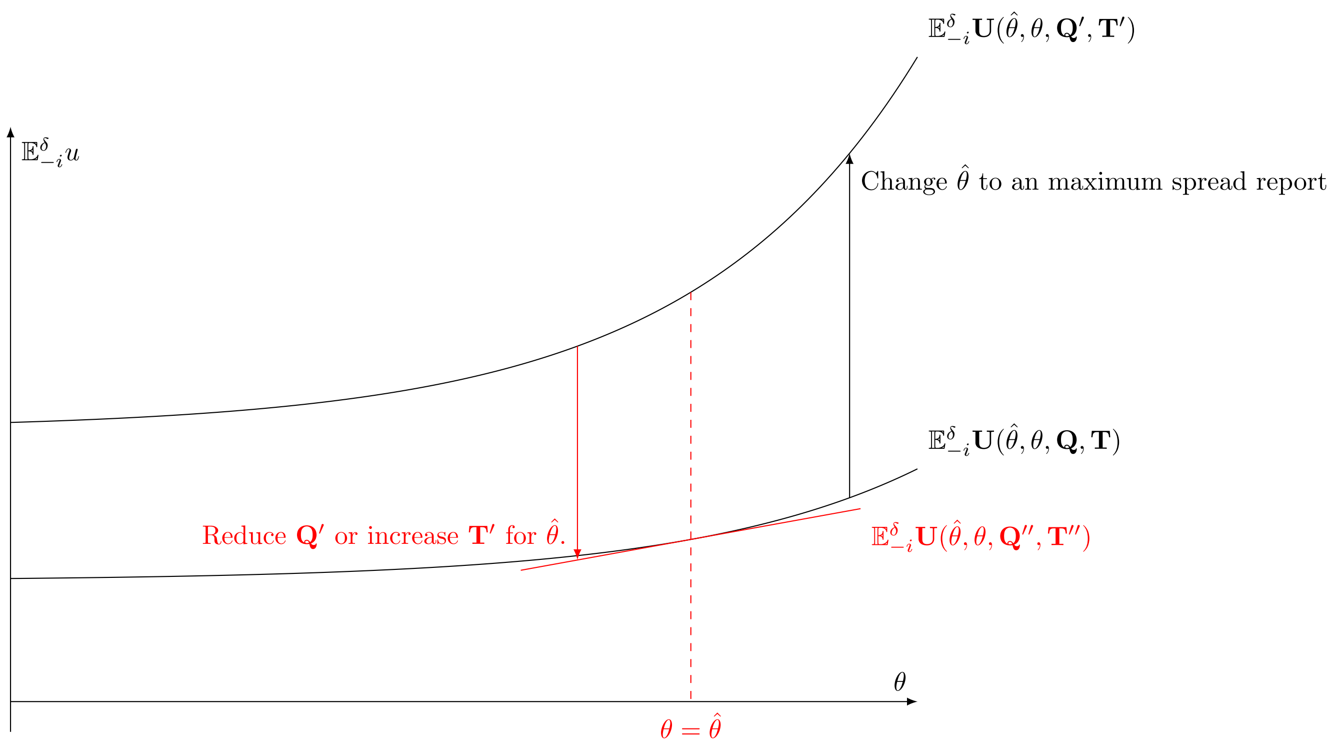

Figure 1.7: The non-maximum spread report is not optimal.

1.6.1 Comparative Statics with Respect to the Level of Distortion

In this section, we examine how the revenue is affected by the level of distortion. First, we compare the revenue between auctions with unbiased bidders and auctions with wishful thinking (WT) bidders as a baseline. Recall that WT bidders make accurate assessments of the probability of purely random events generated by the mechanism but have distorted beliefs about others' types. The optimal mechanism for selling to unbiased bidders can be implemented as a constant mechanism for wishful thinkers by absorbing the randomness from profiles and randomizing the allocation and payment itself. As a result, the wishful thinker becomes unbiased. Therefore, the revenue from auctions with WT bidders is weakly greater than the revenue from auctions with unbiased bidders. This result is summarized in the following proposition.

Proposition 1.14. *The revenue of the seller in auctions with WT bidders is weakly greater than that with unbiased bidders.*

Next, we investigate how the revenue changes with the level of distortion. To do this, we define the notion of "more distorted."

Definition 1.15 (The distortion function v_2 is more wishful than v_1)**.** v_2 is more wishful than v_1 if, for any two random variables $\mathbf{X}$ and $\mathbf{Y}$ such that $\mathbf{X}$ first-order stochastically dominates $\mathbf{Y}$,

$$\mathbb{E}^{\delta(v_2)}\mathbf{X} > \mathbb{E}^{\delta(v_1)}\mathbf{X} \tag{1.6}$$

$$\mathbb{E}^{\delta(v_2)}\mathbf{X} - \mathbb{E}^{\delta(v_2)}\mathbf{Y} > \mathbb{E}^{\delta(v_1)}\mathbf{X} - \mathbb{E}^{\delta(v_1)}\mathbf{Y} \tag{1.7}$$

These two conditions ensure that when v_2 is more distorted than v_1, the increase in expected utility under distortion of v_2 is greater than that under v_1 (condition (1.6)). Moreover, the increase in the spread of expected utility is also greater under v_2 than under v_1 (condition (1.7)).

To understand how the revenue is affected by the level of distortion, consider that if v_2 is more distorted than v_1, the distorted expected utility f_θ under v_2 is greater than that under v_1 due to condition (1.7). As a result, the seller can mimic any mechanism for v_1 under the context of v_2 by reducing allocation and increasing payments. This leads to the following proposition.

Proposition 1.16. *The revenue is increasing in the level of wishfulness.*

The proof is as follows: Suppose (Q, T) is the optimal mechanism for less distorted bidders with distortion function v_1. Given the utility distribution $V(\theta), f_\theta(\theta)\theta \in \Theta$ that it generates, the same (Q, T) will yield a greater utility distribution $V(\theta)', f\theta(\theta)'\theta \in \Theta$ with more distorted bidders using v_2. By Lemma 6, we can reduce it back to the original utility distribution $V(\theta), f\theta(\theta)_{\theta\in\Theta}$ by increasing T and reducing Q. With this change, the new mechanism generates more profit and remains feasible.

As a corollary of the previous propositions, we can simplify the seller's optimization problem under the assumption of a large budget and an increasing expected winning probability schedule, denoted as $\mathbb{E}_{-i}\mathbf{Q}$:

Corollary 1.17. *Under the assumptions of a large budget and an increasing expected winning probability schedule, the seller's optimization problem can be simplified as* *follows:*

$$\max_{a,b,c,d\in\Theta^{[0,1]}} \int_{\Theta} b(\theta)+c(\theta)F(d\theta)\bar{t}$$

$$\begin{aligned}
s.t. \int_{\theta'=0}^{\theta} f_\theta(\theta')d\theta' =& \frac{\sum_{p\in x} p(\theta_i)v_p(\theta_i)u_p(\theta_i)}{\sum_{p\in x} p(\theta_i)v_p(\theta_i)} \text{ for every } \theta\in\Theta, \\
f_\theta(\theta) =& \frac{\left(\sum_{p\in w} p(\theta)(v_p(\theta)+u_p(\theta)v'p(\theta))\right)\left(\sum p\in xp(\theta)v_p(\theta)\right)}{(\sum_{p\in x} p(\theta)v_p(\theta))^2} \\
&- \frac{\left(\sum_{p\in w} p(\theta)v'p(\theta)\right)\left(\sum p\in xp(\theta)u_p(\theta)v_p(\theta)\right)}{(\sum_{p\in x} p(\theta)v_p(\theta))^2}, \\
a+b+c+d =& 1, \\
\int_{\Theta} a(\theta)+c(\theta)F(d\theta) =& \frac{1}{N}, \text{ and } d(0)=1,
\end{aligned}$$

where a,b,c,d represent the probabilities of all possible cases, and $x = a,b,c,d$ and $w = a,c$ represent the winning cases. The utility and distortion functions for each case are denoted as $u_p(\theta)$ and $v_p(\theta)$, respectively.

With consequential distortion, mechanisms that yield the same expected allocation and payment for every type can result in different revenues. Therefore, the seller needs to consider the profile-dependent utility distribution for each report. With the above propositions, the seller can focus on maximum spread auctions, and the problem becomes assigning winning and losing probabilities for each type. Instead of choosing (Q,T) functions that map Θ^N to $[0,1]$ and $[0,\bar{t}]$, the seller can focus on choosing (a,b,c,d) functions that map types in Θ to probabilities in $[0,1]$. This simplifies the dimensionality of the problem. Another possibility is to use the (α,β) functions as defined above, which further reduces the dimensionality of the problem but introduces some non-smoothness due to the use of the $(\cdot)^+$ operator.

1.7 Discussion

Wishful thinking (WT) leads bidders in standard auctions to overestimate their chances of winning, prompting underbidding as they underestimate competitors' valuations. This miscalculation of payment expectancy reduces the potential surplus, lowering the seller's revenue. However, the implementation of sad-loser auctions can reverse this trend, enabling the seller to profit from bidder bias. These auctions further allow the mechanism designer to manipulate the level of distortion by adjusting the outcome distribution. Thus, the more biased the bidders, the greater the seller's profit.

Sad-loser auctions, while not common in goods sales, appear frequently in contracts and market structures like the 'loser pays attorney fees' clauses and R&D processes. This type of system, where 'the winner gets all,' can incentivize WT individuals to participate. Whether this is beneficial depends on the designer's objectives. Given the correlation between optimism and traits like creativity [53], risk-taking [1], and procrastination [58], the implications vary. 'Loser-pays' contracts may be unsuitable for hiring or loans if WT individuals are prone to risky decisions, but may foster creativity in R&D contests.

Our findings indicate that WT individuals underbid in private value auctions. However, the literature suggests that bidders in common value auctions may suffer from the winner's curse, bidding above true value due to underestimated correlations between others' actions and information [13]. Future research could investigate how WT individuals behave in common value auctions when subject to similar biases.

Chapter 2

Conservative Updating

Abstract: This paper provides a behavioral analysis of conservatism in beliefs. We introduce a new axiom, Dynamic Conservatism, that relaxes Dynamic Consistency when information and prior beliefs "conflict." When the DM is a subjective expected utility maximizer, Dynamic Conservatism implies that conditional beliefs are a convex combination of the prior and the Bayesian posterior. Conservatism may result in belief dynamics consistent with confirmation bias, motivated reasoning, and the good news-bad news effect, suggesting a deeper behavioral connection between these biases. An index of conservatism and a notion of comparative conservatism are characterized. Finally, we extend conservatism to the case of an DM with incomplete preferences that admit a multiple priors representation.

Keywords: Conservative updating, prior-bias, non-Bayesian updating, confirmation bias, motivated reasoning, good news-bad news effect, multiple priors.

JEL: D01, D81, D9.

2.1 Introduction

Many papers in both economics and psychology have identified biases in belief updating[1]. A frequent finding is that people often exhibit conservatism ([52], [14], [4], and [43]), meaning that they only partially incorporate new information into their beliefs and give too much weight to their prior beliefs. Despite its prevalence, conservatism has yet to be behaviorally founded.

To address this gap, we introduce a novel behavioral postulate called Dynamic Conservatism, which captures conservatism within the framework of conditional preferences over acts (see [56] and [2]). This axiom weakens Dynamic Consistency to accommodate "preference stickiness," allowing for violations of Dynamic Consistency only when their is a conflict between initial preferences and new information. Furthermore, we show that conservative preferences are consistent with several well-known biases, including confirmation bias, motivated reasoning, and the good news-bad news effect.

For a deeper understanding of Dynamic Conservatism, consider the following hypothetical scenario of a decision maker's (DM) reaction to information regarding climate change. Suppose this DM initially favors using coal for power, but concedes that using alternative energy is better if climate change is occurring. If she receives information indicating the scientific consensus that climate change is occurring, would she now support using alternative energy over coal? If she is dynamically consistent (i.e., Bayesian), then the answer is yes because she has already conceded that alternative energy is better in that contingency. If she is conservative, she may still prefer coal. When there is a conflict between information (climate change news) and a DM's initial preference (use coal), conservatism may result in a violation of Dynamic Consistency.

[1] See for instance, [5], [29], [16])

Dynamic Conservatism allows for such violations.

Dynamic Conservatism restricts the DM so that she may violate Dynamic Consistency only in situations in which there is a conflict between the information and her initial preference. To illustrate, consider how our DM would feel if she had initially preferred alternative energy to coal. Now that the evidence for climate change is consistent with her initial preference, she must continue to prefer alternative energy. That is, no matter how conservative she is, it would be absurd for her to (i) initially support alternative energy, (ii) acknowledge that alternative energy is better if climate change is occurring, and then (iii) state that she supports coal upon receipt of information regarding the consensus that climate change is occurring. Dynamic Conservatism precludes reversals of this form.

Suppose the DM is a subjective expected utility maximizer (SEU) and has prior beliefs over payoff relevant states and signals. In this case, Dynamic Conservatism characterizes a subjective expected utility DM whose conditional beliefs after any signal are given by a convex combination of the prior and the Bayesian posterior. The weight on her prior is signal dependent and is always between zero and one. Therefore, our model nests the standard Bayesian agent — when the weight on her prior is zero, she is Bayesian. When the weight on the prior is positive, the DM is reluctant to move away from her initial beliefs, and therefore she is conservative. In the extreme case where the weight on her prior is one, she is unresponsive to the information. Consequently, the weight on her prior can be interpreted as her signal-dependent degree of conservatism.[2] In general, we refer to such a representation as a **conservative subjective expected utility** representation (henceforth, Conservative SEU).

[2]Relatedly, the weight on the Bayesian posterior might be thought of as a measure of her confidence in the new information. This interpretation will be useful for the discussion of source dependence in section B.3.

An important feature of the representation in is the dependence of the degree of conservatism on the realized signal: bias may be source-dependent. Because of source dependence, the conservative SEU representation is able to capture belief dynamics consistent with confirmation bias ([48]), the motivated reasoning ([36]), and the good news-bad news effect ([15]). Although these biases have been thought of as distinct types of behavior, this finding suggests that these biases can also be thought of as different forms of conservatism.

We then provide an analysis of various aspects related to an DM's degree of conservatism. The weight placed on the prior after signal θ is denoted by $\delta(\theta)$ and captures how conservative a DM is (at θ). When $\delta(\theta)$ is event independent, $\delta(\theta) = \delta$ for some $\delta \in [0, 1]$, then the DM is described by a single behavioral parameter and δ is an index of conservatism. We establish a behavioral characterization of this case through Weak Consequentialism, a novel weakening of Dynamic Consistency. We then provide a simple definition of comparative conservatism that rests upon the comparison of a binary act with a constant outcome. This provides an efficient method for ordering people by each one's degree of conservatism.

We close by exploring the implications of conservatism when DMs do not have precise beliefs. To do so, we consider a collection of possibly incomplete conditional preferences and impose an adapted version of Dynamic Conservatism on these preferences, which we call Unambiguous Dynamic Conservatism. In this setting, the DM has a set of possible beliefs and prefers an act, f, to another, g, when f provides (weakly) greater expected utility than g for every possible belief. Unambiguous Dynamic Conservatism characterizes a conceptually similar representation to the SEU case: the DM's set of posterior beliefs is derived by mixing the initial set of priors and the set of Bayesian posteriors. A crucial distinction from the subjective expected utility case is that now information may induce a genuine expansion of the set of beliefs. Thus subjectively ambiguous information may result in a structural change in behavior; an initially subjective expected utility DM may become ambiguity averse after information.

2.1.1 Outline and Related Literature

The setup and model are presented in section B.1. Behavioral foundations are presented in section B.2. The connection between conservatism and other belief biases, and the analysis of degrees of conservatism is discussed in section B.3. We present the extension to multiple priors and incomplete preferences in section B.4. We close with a conclusion remarks in section 2.5. The remainder of this section discusses related literature.

While numerous experimental papers suggest that people do not update their beliefs in a Bayesian manner, relatively few provide an axiomatic analysis of non-Bayesian updating. [20] provided one of the first axiomatic analysis of non-Bayesian updating. He utilized a setup of preferences over menus and modeled non-Bayesian updating as a temptation (á la [26]) to modify prior beliefs in response to an interim signal. This was extended to an infinite-horizon model by [21]. Because of the menu-preference setting, beliefs are typically dependent on both the information and the option set. More recently, [50] utilized the conditional preference approach to introduce the Hypothesis Testing Representation. In his model, a decision maker may be surprised by low probability events and, in response, adopt a new prior before applying Bayes' rule. This representation captures "over-reaction" to information and is conceptually distinct from conservatism.[3]

[3]Additionally, such a decision maker always satisfies Consequentialism, which is typically violated

The notion of conservative updating we characterize results in violations of Dynamic Consistency.[4] While violations of the property is well documented, the causes of these violations are still not fully understood. Violations of Dynamic Consistency and Consequentialism were documented by [12] in an Ellsberg-style experiment. While these violations were more frequent among ambiguity averse subjects, they also occur among ambiguity neutral subjects.[5]

Conservative updating of the form we characterize has recently been applied to cheap-talk games in [37]. They found that a conservative receiver may induce the sender to provide more accurate signals, and thus conservatism may be welfare improving. Hence, while our results show that conservatism leads to "mistakes" in individual choice (the DM violates Dynamic Consistency) there may be benefits from conservatism in other situations. [10] extended the model of Bayesian persuasion [31] to general non-Bayesian receivers (e.g., a patient who exhibits conservatism bias).

2.2 Model and Setup

Our setting distinguishes between payoff-relevant states, $S = \{s_1, .., s_n\}$, and signals, $\Theta = \{\theta_1, ..., \theta_m\}$. Further, we suppose there are at least two states and at least two signals. Our DM considers various acts, which are functions that map payoff-relevant states to a (nonempty) set X of consequences; $f : S \to X$. The set of acts is denoted $\mathcal{F}$. Our DM understands that she may receive information, and so has in mind the

by conservative updating.

[4]Consequentialism is also violated in a more general information structure.

[5]Further, all subjects report a "loss of confidence" in their choices after information, with the greatest loss of confidence occurring among subjects who violated both Dynamic Consistency and Consequentialism. This is consistent with the interpretation of $1 - \delta(\theta)$ as the degree of confidence in the information, with lower confidence leading to more violations.

extended state space $\Omega := S \times \Theta$.[6] Let $\mathcal{P}$ denote the set of functions that map extended states to consequences $a : \Omega \to X$. We call the "extended acts" of $\mathcal{P}$ *plans*, and a, b are typical elements of $\mathcal{P}$.

We can identify each act with the constant plan that implements act f regardless of the signal; for each $f \in \mathcal{F}$ by $f \in \mathcal{P}$ we mean the constant plan $a(s_i, \theta) = f(s_i)$ for all $\theta \in \Theta$. Accordingly, we can embed $\mathcal{F}$ into $\mathcal{P}$.

Similarly, for any $x \in X$, by $x \in \mathcal{P}$ (or $x \in \mathcal{F}$) we mean the constant plan (act) that returns x in every state (payoff state). Lastly, for any $f, g \in \mathcal{F}$ and for any signal $\theta \in \Theta$, let $f\theta g \in \mathcal{P}$ denote the plan that implements act f after θ and implements act g for all other signals.

Following the literature, we assume that X is a convex subset of a vector space.[7] Thus, mixed acts can be defined point-wise, so that for every $f, g \in \mathcal{F}$ and $\lambda \in [0, 1]$, by $\lambda f + (1 - \lambda)g$ we mean the act that returns $\lambda f(s) + (1 - \lambda)g(s)$ for each $s \in S$. Mixed plans are defined similarly. For every $a, b \in \mathcal{P}$ and $\lambda \in [0, 1]$, $\lambda a + (1 - \lambda)b$ means the plan that returns $\lambda a(\omega) + (1 - \lambda)b(\omega)$ for each $\omega \in \Omega$.

We assume that the DM has an ex-ante preference over plans and a collection of preferences over $\mathcal{F}$ conditional on a signal. Note that $\{S \times \{\theta\}\}_{\theta \in \Theta}$ forms a partition of Ω, and so we simply identify θ with the corresponding cell of this partition. Formally, the DM has an ex-ante preference $\succsim$ over $\mathcal{P}$, and a collection of preference relations, $\{\succsim_\theta\}_{\theta \in \Theta}$ over $\mathcal{F}$, where $\succsim_\theta$ denotes her preferences after observing θ. For any preference relation $\succsim$, let $\succ$ and $\sim$ represent the asymmetric and symmetric parts of it. We let $\{\succsim, \{\succsim_\theta\}_{\theta \in \Theta}\}$ denote the collection of ex-ante preferences and preferences after

[6]The product structure supposed here is merely for convenience; we can characterize the main result in a general state space with some algebra of events. See [34] for details.

[7]For instance, X may be an interval of monetary prizes or, as in the classic [2] setting, a set of lotteries over some prize space.

signals.

In some cases it may be important to distinguish between the beliefs over the extended state space $\Omega = S \times \Theta$, which we denote by $\hat{\mu} \in \Delta(\Omega)$, and beliefs over the payoff-relevant states which we denote by $\mu \in \Delta(S)$. When this distinction is important we will refer to μ as the payoff-relevant belief. Note that each $\hat{\mu} \in \Delta(\Omega)$ induces a payoff-relevant belief $\mu \in \Delta(S)$, where for every $s \in S$, $\mu(s) = \sum_{\theta\in\Theta} \hat{\mu}(s, \theta)$. For any prior $\hat{\mu} \in \Delta(\Omega)$, and any $\theta \in \Theta$, define the Bayesian update of $\hat{\mu}$ given θ as $\mathcal{B}(\hat{\mu}, \theta)(s) = \frac{\hat{\mu}(s,\theta)}{\sum_{s'\in S} \hat{\mu}(s',\theta)}$. We will suppose that $\hat{\mu}$ is a full support prior, and each signal is unique. That is, Bayes' rule maps each signal to a distinct posterior probability.[8] Finally, for any two utility functions $u, v : X \to \mathbb{R}$, say $u \approx v$ if u is a positive affine transformation of v.

The ex-ante preference over plans is represented by a utility function u and a probability measures $\hat{\mu} \in \Delta(\Omega)$. For $a, b \in \mathcal{P}$, $a \succsim b$ if and only if $\sum_{\omega\in\Omega} u(a(\omega))\hat{\mu}(\omega) \geq \sum_{\omega\in\Omega} u(b(\omega)\hat{\mu}(\omega)$. As we embed $\mathcal{F}$ into $\mathcal{P}$. As a result, restriction of $\succsim$ top $\mathcal{F}$ is represented by the same utility function u and the prior belief $\mu \in \Delta(S)$ induced by $\hat{\mu}$. $f \succsim g$ if and only if $\sum_{s\in S} u(f(s))\mu(s) \geq \sum_{s\in S} u(g(s))\mu(s)$.

Definition 2.1. Say that a collection of preferences $\{\succsim_\theta\}_{\theta\in\hat{\Theta}}$ admits a **conservative subjective expected utility** (conservative SEU) representation if there are a non-constant utility function $u : X \to \mathbb{R}$, a prior belief $\hat{\mu} \in \Delta(\Omega)$, and a function $\delta : \Theta \to [0, 1]$ such that for all $\theta \in \Theta$,

$$a \succsim b \Longleftrightarrow \sum_{\omega\in\Omega} u(a(\omega))\hat{\mu}(\omega) \geq \sum_{\omega\in\Omega} u(b(\omega))\hat{\mu}(\omega) \tag{2.1}$$

[8]The two assumptions can be guaranteed by behavioral axioms: (i) there is no savage null state in Ω; and (ii) for every $\theta \neq \theta'$, there exists act f and consequences x, y such that $f\theta x \sim x \neq y \sim f\theta' y$.

$$f \succsim_\theta g \Longleftrightarrow \sum_{s\in S} u(f(s))\mu_\theta(s) \geq \sum_{s\in S} u(f(s))\mu_\theta(s) \tag{2.2}$$

and

$$\mu_\theta = \delta(\theta)\mu + (1-\delta(\theta))\mathcal{B}(\hat{\mu},\theta). \tag{2.3}$$

Example 2.2. An investor is deciding between making an investment into a risky company or a safe asset and plans to decide after reading a market report. There are two payoff states, $S = \{H, L\}$ and two signals $\Theta = \{h, l\}$. Let $\Omega = S \times \Theta$. The prior belief $\hat{\mu} \in \Delta(\Omega)$ is given in Table 2.1 below.

$\hat{\mu}$	h	l
H	4/8	1/8
L	1/8	2/8

Table 2.1: Prior over $S \times \Theta$.

This example has a natural interpretation as a standard signaling experiment with payoff-relevant belief on S given by $\mu = (\frac{5}{8}, \frac{3}{8})$, and (asymmetric) signal accuracy $\sigma(h|H) = \frac{4}{5}$, $\sigma(l|L) = \frac{2}{3}$.

If the DM admits a conservative SEU representation, the DM's posterior belief in H after h and l, are

$$\mu_h(H) = \delta(h)\frac{5}{8} + (1-\delta(h))\frac{4}{5}$$

$$\mu_l(H) = \delta(l)\frac{5}{8} + (1-\delta(l))\frac{1}{3}.$$

Since δ is information dependent, our setting allows for a rich variety of other biases to emerge. For instance, when $\delta(h) < \delta(l)$, the DM's conservatism mimics confirmation bias. That is, because her conservatism bias is greater when she receives a signal that is counter to her "prior hypothesis" (i.e., $\mu(H) > \frac{1}{2} > \mu(L)$) she is more reluctant to incorporate "disconfirming information." In the extreme case of $\delta(h) = 0$ and

$\delta(l) = 1$, the DM only reacts to the positive signal (h) and completely ignores the negative one (l). This and other biases are further discussed in section B.3.

As a second example, asymmetric conservatism may induce discrimination. A recent experiment by [40] shows that employers exhibit conservative updating and are more conservative toward positive signals from disadvantaged groups.

Example 2.3. There are two groups of workers $i \in \{g, r\} = G$, each endowed with an ability level $t \in \{h, m, l\} = S$. The disadvantaged group r, is believed to have lower ability. Workers decide whether to pursue education with a completion rate that depends on their ability, $t \in \{e, ne\} = T$. The employer would like to hire the a worker with higher ability conditional on education level. The group membership and (endogenous) choice of education level act as signals for ability (i.e. $\Theta = G \times T$). The employer is conservative, and more so against the disadvantage group, whenever $0 < \delta(g, e) = \delta(g, ne) < \delta(r, e) = \delta(r, ne) < 1$.

2.3 Behavioral Foundations

The first axiom imposes a subjective expected utility (SEU) representation at each information set. The conditions for this are well-established in the literature.

Axiom 1 (Conditional SEU)**.** Each preference in the collection $\{\succsim, \{\succsim_\theta\}_{\theta\in\Theta}\}$ admits a non-degenerate subjective expected utility representations.

Before introducing the conservatism axiom, for comparison, we first state the classic axioms of Dynamic Consistency and Consequentialism. An excellent discussion of both axioms is provided in [23], and so we will only briefly discuss them.

Axiom 2 (Dynamic Consistency)**.** For any $\theta \in \Theta$, and $f, g \in \mathcal{F}$,

$$f\theta g \succsim g \Leftrightarrow f \succsim_\theta g.$$

Axiom 3 (Consequentialism)**.** For all $\theta \in \Theta$, and $f, g \in \mathcal{F}$,

$$f \sim_\theta f\theta g.$$

In essence, $f\theta g \succsim g$ reveals that the decision maker believes she would abandon g for f if θ occurred. Dynamic Consistency states that if this is the case, then f must be preferred to g after being told θ has occurred. Consequentialism states that whenever two acts are identical within θ, the DM must be indifferent between them after θ.[9] It is well known that Dynamic Consistency and Consequentialism, when combined with Conditional SEU, are necessary and sufficient for Bayesian updating [23]. Hence, these must be relaxed to allow for conservatism.

To illustrate how these axioms are relaxed, recall the introductory example on the use of alternative energy versus coal. The act f is "use alternative energy," g is "use coal," and the event θ is "climate change is occurring." The DM concedes that alternative energy is better if climate change is occurring. Depending on the DM's initial preference between f and g, there are two relevant cases:

Case 1: $f \succsim g$ and $f\theta g \succsim g$,

Case 2: $g \succ f$ and $f\theta g \succsim g$.

[9] In our setting, $f \in \mathcal{F}$ and $g \in \mathcal{F}$ implies $f = f\theta g$. Thus, Consequentialism is trivially met. But in a more general setting, our model do have implications on violations of Consequentialism, see section 2.5 for details.

In case 1, the preference for the fixed action f (over g) is consistent with the preference for the contingent action $f\theta g$ (over g). In case 2, the preferences are in conflict. Dynamic Consistency requires the DM to conclude that f is better than g in the conditional preference in both cases. However, intuition suggests that a conservative DM may violate Dynamic Consistency in case 2. Indeed, she does so because she does not fully believe θ has occurred. Similarly, a conservative DM would never violate Dynamic Consistency in case 1 since she prefers f to g irrespective of θ's occurrence. The following axiom takes this intuition as the defining behavior of conservatism.

Axiom 4 (Dynamic Conservatism)**.** For any $\theta \in \Theta$ and for all $f, g \in \mathcal{F}$,

$$\left.\begin{array}{ll} \text{(i)} & f \succsim g \\ \text{(ii)} & f\theta g \succsim g \end{array}\right\} \implies f \succsim_\theta g.$$

Further, if both (i) and (ii) are strict, then $f \succ_\theta g$.

Dynamic Conservatism requires that if the DM (i) prefers f to g ex-ante and (ii) forecasts that she would abandon g for f if θ occurred, $f\theta g \succsim g$, then she must prefer f to g conditional on θ. Therefore, while Dynamic Conservatism allows for some violations of Dynamic Consistency (e.g., the climate change example), it restricts violations to cases where the DM cannot be completely sure that she is making the right decision. Hence, we can think of Dynamic Conservatism as allowing for a form of "stickiness" or skepticism about new information. Seen through this lens, Dynamic Conservatism can be viewed as a cautious response when the reliability of information is (subjectively) uncertain.[10] Consequently, Dynamic Conservatism permits a DM to

[10]Another way of viewing Dynamic Conservatism is through a two-self interpretation. Condition (i) captures a self that does not learn, while condition (ii) captures a Bayesian self. Dynamic Conservatism states that whenever the selves agree, then the DM's behavior necessarily reflects this agreement. When they disagree, then Dynamic Conservatism says nothing. The representation result, however, shows that the DM's behavior must be governed by a "compromise" belief.

state that they "will prefer f after θ," $f\theta g \succsim g$ but once θ occurs they still maintain a preference for g, $g \succ_\theta f$.

Theorem 2.4. *The following are equivalent:*

(i) The collection $\{\succsim_\theta\}_{\theta\in\hat{\Theta}}$ satisfies Conditional SEU and Dynamic Conservatism;

(ii) The collection $\{\succsim_\theta\}_{\theta\in\hat{\Theta}}$ admits a conservative subjective expected utility representation.

Further, if $(u, \hat{\mu}, \delta)$ and $(u', \hat{\mu}', \delta')$ both represent $\{\succsim_\theta\}_{\theta\in\hat{\Theta}}$, then (i) u' is a positive affine transformation of u, (ii) $\hat{\mu}' = \hat{\mu}$, and (iii) $\delta'(\theta) = \delta(\theta)$ for all $\theta \in \Theta$.

Theorem 1 shows that when conditional preferences admit a SEU representation, Dynamic Conservatism is the precise behavioral content of conservative updating. Further, the degree of conservatism, $\delta(\theta)$, is uniquely pinned down at essentially every θ.

Example 2.5 (continues=ConfBias)**.** To further illustrate this result, recall our example with two payoff states $S = \{H, L\}$ and two signals $\Theta = \{h, l\}$. As illustrated in Figure 2.1, Dynamic Conservatism ensures that the lower contour set of the conditional indifference curve passing through f must contain the intersection of the lower contour sets determined by (i) the initial preference (e.g., μ) and (ii) the preference that corresponds to performing Bayesian updating $\mathcal{B}(\hat{\mu}, h)$. This intersection is shaded in blue. In essence, conservatism pulls the conditional indifference curve to to be more aligned with the initial preference.

Remark 2.6. It is easy to see from Figure 2.1 that over-inference or "prior-neglect" may be captured by rotating the (orange) indifference curve away from μ and beyond

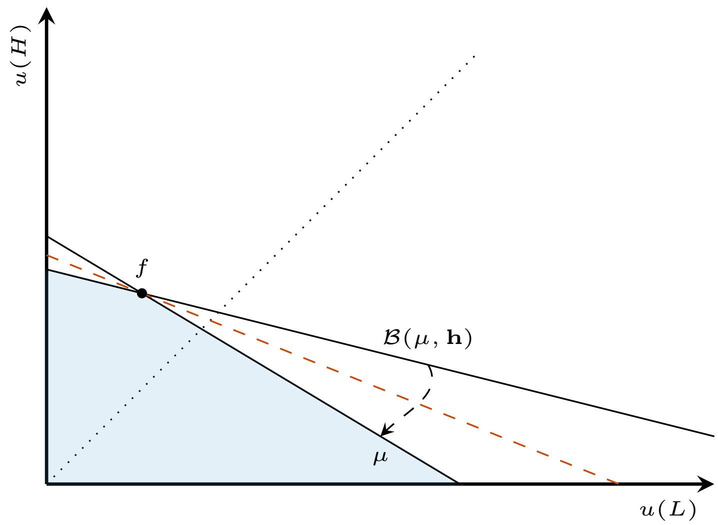

Figure 2.1: Indifference curves after the DM receives an h signal. The (orange) dashed line corresponds to an indifference curve passing through f for some $\delta(\mathbf{h}) \in (0, 1)$.

$\mathcal{B}(\mu, \mathbf{h})$ (e.g., intuitively, this is captured by $\delta(\mathbf{h}) < 0$). This would imply the existence of an act g, such that $f \succsim g$, $f\mathbf{h}g \succsim g$ but $g \succ_{\mathbf{h}} f$, violating Dynamic Conservatism. However, over-inference is difficult to capture in the conditional preference framework with a general state space. This is because $\delta(\mathbf{h}) < 0$ results in negative values for the probabilities of certain states. Thus, the behavioral foundations of such behavior remain an open question.

2.4 Degrees of Conservatism and Context Dependence

2.4.1 A Conservatism Index

B.2 allows for the degree of conservatism to depend on the information received. In many settings, this is useful as it allows for source-dependent reactions to news.

However, it is often convenient to fully describe behavior with a single parameter, or an index of conservatism. Further, a constant degree of conservatism simplifies the task of identifying behavioral parameters from data and increases the predictive power of the model. We show that constant conservatism is characterized by a consistency condition linking conditional preferences across information sets that may be viewed as a weak form of Dynamic Consistency.

Axiom 5 (Signal Neutrality)**.** For any $f, f' \in \mathcal{F}$, $\theta, \theta' \in \Theta$, and $x \in X$ such that $f \sim f'$ and $f\theta x \sim x \sim f'\theta' x$, there exists $y \in X$ such that $f \sim_\theta y \sim_{\theta'} f'$.

To see why Weak Consequentialism is a weak form of Dynamic Consistency, suppose $f\theta x \sim x \sim f'\theta' x$. Then Conditional SEU and Dynamic Consistency require that x is indifferent to f if θ occurs, and indifferent to f' if θ' occurs. That is, Dynamic Consistency requires not just that they have the same utility value, but that the value is exactly $u(x)$. Instead, we only impose that the value of f after θ is equal to f' after θ', without specifying the value.

For a Conservative DM, conditional preferences are affected by the ex-ante preference for fixed actions (like f) and contingent actions (like $f\theta x$). The statement at least requires an ex-ante indifference between the fixed actions ($f \sim f'$). However, this is not sufficient. The conservative attitudes toward signals also affect conditional preferences. Weak Consequentialism rules out the possibility of varying conservative attitudes toward signals. It requires the DM to be equally conservative towards all signals and ensures that $f \sim f'$ and $f\theta x \sim x \sim f'\theta' x$ sufficient for the statement that "f conditional on θ is indifferent to f' conditional on θ'."

Proposition 2.7. *Suppose the collection* $\{\succsim_\theta\}_{\theta \in \hat{\Theta}}$ *admits a Conservative SEU representation* (u, μ, δ)*. Then the following are equivalent:*

(i) The collection $\{\succsim_\theta\}_{\theta\in\hat{\Theta}}$ *satisfies Weak Consequentialism.*

(ii) There is a unique $\delta \in [0,1]$ *such that* $\delta(\theta) = \delta$ *for all* θ*.*

A key strength of this result is that now δ may serve as a simple index of conservatism bias and may be elicited with only a few questions.

2.4.2 Contextual Ordering of Signals

The DM's response to information may depend on various characteristics of the signal. For example, many experiments document an asymmetric willingness to accept good vs bad news ([15], [43], and [6]). This effect is exhibited whenever θ *is better news than* θ' implies $\delta(\theta) \leq \delta(\theta')$. Alternatively, an DM may react more strongly toward anticipated signals, thereby exhibiting confirmation bias (see [48] for an excellent review). This effect is exhibited whenever θ *is more likely than* θ' implies $\delta(\theta) \leq \delta(\theta')$.

To characterize these examples and other biases, it is convenient to abstract away from the specific context. Accordingly, we introduce an abstract **contextual ordering over signals**, denoted by $\trianglerighteq^\Theta$, to capture the DM's reaction to signals. Intuitively, $\theta \trianglerighteq^\Theta \theta'$ means that the DM has a stronger reaction to θ than to θ'. We let $\approx^\Theta$ denote the symmetric part of $\trianglerighteq^\Theta$, and so $\theta \approx^\Theta \theta'$ means that the DM is equally reactive to each signal. For example, the contextual ordering $\theta \trianglerighteq^\Theta \theta'$ may capture θ *is better news than* θ', θ *is more likely than* θ', or other, potentially more complex, features.

Definition 2.8. $\trianglerighteq^\Theta$ is a **contextual ordering over signals** if it is a complete and transitive ordering over Θ.

We will generally impose a behavioral interpretation upon $\trianglerighteq^\Theta$, and, as stated earlier, suppose that it captures the DM's relative responsiveness to signals. However, it also permits an interpretation as the perceived veracity or reliability of the source. A conservative DM exhibits a sort of "epistemic cold feet" and behaves as if she doubts the information she just received. Consequently, she discounts the information received and her beliefs move less than they should.[11]

Essentially all papers on updating make extensive use of the information revealed from preferences over contingent acts of the form $f\theta z$. This is because the statement $f\theta z \succsim g\theta z$ admits the interpretation that "f is better than g conditional on signal (or event) θ." We will build on this idea to impose some weak form of consistency on the DM's response to different signals. To do so, we derive the general contingent preference "f after θ is better than f' after θ'" from the ex-ante preference and denote this implied contingent preference $\succsim^c$. Formally, for all acts $f, f' \in \mathcal{F}$, and signals $\theta, \theta' \in \Theta$, and consequences $x, y \in X$:

$$(f, \theta) \succsim^c (f', \theta') \text{ when } f\theta x \sim x, f'\theta' y \sim y \text{ and } x \succsim y. \tag{2.4}$$

Under expected utility, $f\theta x \sim x$ implies that x is the (anticipated) certainty equivalent for f when the signal θ occurs. Thus, the standard contingent statement $f\theta z \succsim g\theta z$ is equivalent to $(f, \theta) \succsim^c (g, \theta)$. The ex-ante preference over constant plans f is already embedded into $\succsim^c$. Letting $\emptyset$ denote "no information," then $(f, \emptyset) \succsim^c (f', \theta')$ means $f \succsim y \sim f'\theta' y$. Finally, note that $(f, \theta) \succsim^c (f, \emptyset)$ means the signal is "good news" for act f.

We now introduce a general consistency axiom that ensures that the behavioral response to a signal is consistent with $\trianglerighteq^\Theta$. We will later show how to formally connect $\trianglerighteq^\Theta$ to specific biases.

[11] Under this interpretation, we refer to $\trianglerighteq^\Theta$ as the "reliability ordering" and say that "θ is more reliable than θ'" if $\theta \trianglerighteq^\Theta \theta'$.

Axiom 6 (Contextual Consistency)**.** For all $f, f' \in \mathcal{F}$, $\theta, \theta' \in \Theta$, such that $f' \succsim f, (f', \theta') \succsim^c (f, \theta)$, and $(f, \emptyset) \succsim^c (f, \theta)$, then $\theta \trianglerighteq^\Theta \theta'$ implies

$$\forall x \in X : [f \succsim_\theta x \Rightarrow f' \succsim_{\theta'} x] \text{ and } [x \succsim_{\theta'} f' \Rightarrow x \succsim_\theta f].$$

To understand Contextual Consistency, let's consider two indifferent acts f and f', and two signals θ, and θ'. The signals θ and θ' are "bad news" for f and f' respectively. (i.e. $f \sim f'$, $(f, \emptyset) \succsim^c (f, \theta)$ and $(f', \emptyset) \succsim^c (f', \theta')$). When the anticipated impact of θ on f is greater than that of θ' on f', the DM should prefer acting on f' and receiving θ' than to acting on f and receiving θ. Unless the DM is more conservative when receiving θ than when receiving θ'. Given $\{\succsim_\theta\}_{\theta \in \hat{\Theta}}$, Contextual Consistency requires the ordering $\trianglerighteq^\Theta$ to be consistent with the degree of reliability of signals.

When the collection of preferences $\{\succsim_\theta\}_{\theta \in \hat{\Theta}}$ is known. Instead of saying that the collection $\{\succsim_\theta\}_{\theta \in \hat{\Theta}}$ and the ordering $\trianglerighteq^\Theta$ jointly satisfy Contextual Consistency, we simply say the ordering $\trianglerighteq^\Theta$ satisfies Contextual Consistency.

Proposition 2.9. *Suppose the collection $\{\succsim_\theta\}_{\theta \in \hat{\Theta}}$ admits a Conservative SEU representation (u, μ, δ), and the ordering $\trianglerighteq^\Theta$ satisfies Contextual Consistency, then*

$$\theta \trianglerighteq^\Theta \theta' \Leftrightarrow \delta(\theta) \leq \delta(\theta').$$

This proposition suggests that the constant δ in B.3 reflects an equal treatment for signals; reflects an indifferent contextual ordering, $\theta \approx^\Theta \theta'$ for any signals. Instead of this uniform case, usually, the response differs between signals, depending on factors such as prior belief, preference over states, and reference point. The remainder of this section shows that those factor-dependent contextual orderings correspond to some well-known biases in the literature.

Generalized Confirmation Bias

Confirmation bias (see [48] for an excellent review) refers to a tendency to accept information that supports already believed hypotheses and to downplay conflicting information. Confirmation bias emerges from conservatism bias when the weight attached to the prior is greater for "disconfirming" news than for "confirming" news.

Example 2.10 (continues=ConfBias)**.** There are two payoff states, $S = \{H, L\}$, and two signals, $\Theta = \{h, l\}$, and the ex-ante belief $\hat{\mu}$ is given in Table 2.1 below.

$\hat{\mu}$	h	l
H	4/8	1/8
L	1/8	2/8

If the DM admits a conservative SEU representation, posterior beliefs after **h** and **l**, written in terms of the marginal belief on R, are

$$\mu_{\mathbf{h}}(H) = \delta(\mathbf{h})\frac{5}{8} + (1 - \delta(\mathbf{h}))\frac{4}{5}$$

$$\mu_{\mathbf{l}}(L) = \delta(\mathbf{l})\frac{5}{8} + (1 - \delta(\mathbf{l}))\frac{1}{3}.$$

When $\delta(\mathbf{h}) < \delta(\mathbf{l})$, the DM's beliefs exhibit confirmation bias in addition to conservatism; when she receives a signal that is counter to her "prior hypothesis" (i.e., $\mu(H) > \frac{1}{2} > \mu(L)$), she is more reluctant to incorporate "disconfirming information."

By weakening Weak Consequentialism to depend on the ex-ante likelihood of events, a behavioral characterization of conservative preferences that are consistent with confirmation bias may be obtained. To do so, we first define a qualitative likelihood ordering over events.

Definition 2.11. For any $\theta, \theta' \in \Theta$, say that θ is **more likely than** θ', denoted $\theta \geq_l \theta'$ if for all $x, y \in X$, $x \succsim y$ implies $x\theta y \succsim x\theta' y$.

Recall that Weak Consequentialism ensures a constant δ. This is independent of the ex-ante relative likelihood of θ or θ'. Confirmation bias, on the other hand, suggests that subjectively more likely events are incorporated more accurately. Less doubt arises after receiving a well-anticipated signal.

Axiom 7 (Generalized Confirmation Bias)**.** For any $\theta, \theta' \in \Theta$, if $\theta \geq_l \theta'$, then $\theta \trianglerighteq^{\Theta} \theta'$.

Proposition 2.12. *Suppose the collection $\{\succsim_\theta\}_{\theta \in \hat{\Theta}}$ admits a Conservative SEU representation (u, μ, δ). Then the following are equivalent:*

(i) The ordering $\trianglerighteq^{\Theta}$ satisfies Contextual Consistency and Generalized Confirmation Bias.

(ii) $\mu(\theta) \geq \mu(\theta')$ if and only if $\delta(\theta) \leq \delta(\theta')$.

Motivated Reasoning

Motivated reasoning refers to the tendency of the cognitive process to be biased by motivations to arrive at "desirable" conclusions, given some seemingly reasonable justification [36]. In general, desirability is generously defined. A conclusion can be desirable for goals and reasons that are not necessarily related to payoff [19]. [43] conducted a large-scale experiment to test the motivated reasoning bias arising from updating beliefs about "ego-relevant" information. The experiment shows that overall the subjects are conservative. Also, they systematically react more to a positive feedback than a negative one.

In the cognitive process of belief updating, the conclusion to be arrived at is a posterior belief, a subjective probability over states. Whether it is desirable depends on the desirability of the state it weighted heavily. The DM is biased towards signals that weight their preferred state heavily, according to a preference $\trianglerighteq^S$ over states S.

Axiom 8 (Motivated Reasoning)**.** For any $\trianglerighteq^S$ with a maximal element s^*, $\theta, \theta' \in \Theta$, if $\hat{\mu}(s^*, \theta) \geq \hat{\mu}(s^*, \theta')$, then $\theta \trianglerighteq^\Theta \theta'$.

The axiom states that a signal will be perceived as more reliable when it provide more support for the desirable state s^*.

Proposition 2.13. *Suppose the collection $\{\succsim_\theta\}_{\theta \in \hat{\Theta}}$ admits a Conservative SEU representation (u, μ, δ). Then the following are equivalent:*

(i) The orderings $\trianglerighteq^\Theta$ and $\trianglerighteq^S$ satisfy Contextual Consistency and Motivated Reasoning.

(ii) $\hat{\mu}(s^*, \theta) \geq \hat{\mu}(s^*, \theta')$ if and only if $\delta(\theta) \leq \delta(\theta')$,

where s^* is a maximal element of $\trianglerighteq^S$.

Example 2.14. Suppose the DM takes an ability test. Ability is either high or low, denoted by H and L, and possible results from the test, h or l, are signals of ability. The signal is ego-relevant, and having a higher ability is more desirable: $H \trianglerighteq^S L$. When a report indicating low ability is received, the DM will start to doubt the reliability of the result. In contrast, a report indicating high ability will be accepted directly. Thus, $\delta(l) > \delta(h) = 0$: the DM is conservative, and positive feedback is perceived as more reliable than negative feedback.

Reference Dependence and the Good-news Bad-news Effect

In many contexts preferences appear to be dependent on some reference point. For example, the influential model of prospect theory [30] introduced the idea of loss aversion. Whether an outcome is a loss or a gain is defined relative to some reference point (typically the status quo or endowment). [60] considers the status quo as reference point, and [32] considers the most recent expectations. For a similar reason, the good news-bad news effect [15] can also be considered to fall into this category. Our setting connects *ex-ante* plans with θ-conditional preferences. Suppose the DM has a reference act, and when θ is more in favor of the reference act, it is perceived as more reliable.

Axiom 9 (Reference Dependence)**.** Let $f^* \in \mathcal{F}$ denote a fixed reference act. Then for every $\theta, \theta' \in \Theta$,

$$\text{if } (f^*, \theta) \succsim^c (f^*, \theta'), \text{ then } \theta \trianglerighteq^\Theta \theta'.$$

The axiom states that the reliability of a signal depends on how favorable it is to the reference act.

Proposition 2.15. *Suppose the collection $\{\succsim_\theta\}_{\theta\in\hat{\Theta}}$ admits a Conservative SEU representation (u, μ, δ). Then the following are equivalent:*

(i) There exist f^* such that the ordering $\trianglerighteq^\Theta$ satisfies Contextual Consistency and Reference Dependence.

(ii) $(f^*, \theta) \succsim^c (f^*, \theta')$ if and only if $\delta(\theta) \leq \delta(\theta')$.

How the reference act is chosen is determined is outside the scope of this setting. For instance, it could be an exogenous status quo. The following example consider the case when an act chosen in a previous choice problem serves as the reference act.

Example 2.16. Consider a variation on example 1. An investor is deciding between investing in the risky company or in the safe asset, but in two rounds. In the first round, the DM chooses before reading a market report, and in the second round, after reading the report. The consequences of the two decisions are realized after the two decisions are made. The payoffs are given in the following table.

Suppose the market report turns out to be l. Before reading the market report, according to the prior belief in Table 2.1, the DM decides to invest in the risky company. A Bayesian DM will invest in the safe asset after reading the report l. However, l is bad news for those who are investing in risky asset, which is the choice of the DM in the first round. When the DM is really conservative to bad news, they perceive l as an unreliable signal and keep investing in the risky company regardless of the report.

2.4.3 Comparative Conservatism

Intuitively, the more conservative a DM is, the less responsive to signals she is. This can be captured in the representation by a larger weight on the prior belief. Analogously, one person is more conservative than another if he places a larger weight on his prior belief than she does on her prior belief. This can be formally defined in terms of preferences.

Definition 2.17. Say that $\{\succsim^1_\theta\}_{\theta\in\hat{\Theta}}$ is **more conservative** than $\{\succsim^2_\theta\}_{\theta\in\hat{\Theta}}$ if for all θ, all $f_1, f_2 \in \mathcal{F}$, satisfying (i) $f_1 \succsim^1 x \iff f_2 \succsim^2 x$ for all $x \in X$, (ii) $(f_1,\theta) \succsim^{c1} x \iff (f_2,\theta) \succsim^{c2} x$ for all $x \in X$, and (iii) $(f_i,\theta) \succsim^{ci} (f_i,\emptyset)$,

$$\text{for all } x \in X : f_1 \succsim^1_\theta x \implies f_2 \succsim^2_\theta x.$$

Consider two DMs with similar preferences. Both of them prefer f to g, but planned to choose g over f contingent on receiving signal θ. That is $f \succ g$, and $g\theta k \succ f\theta k$ for any act k. Upon receiving the signal, a Bayesian DM will always implement the plan. A conservative DM, however, may revise her plan. She doubts her prior judgment about the reliability of the signal. Consequently, the signal is discounted, and the plan is less likely to be implemented. Thus, when the more conservative DM implements the plan, the less conservative one will also implement it.

Proposition 2.18. *Suppose* $(\{\succsim^i_\theta\}_{\theta\in\hat{\Theta}})_{i=1,2}$ *admit Conservative SEU representations* $(u_i, \mu_i, \delta_i)_{i=1,2}$ *where* $u_1 \approx u_2$*. Then the following are equivalent:*

(i) $\{\succsim^1_\theta\}_{\theta\in\hat{\Theta}}$ *is more conservative than* $\{\succsim^2_\theta\}_{\theta\in\hat{\Theta}}$.

(ii) $\delta_1(\theta) \geq \delta_2(\theta)$ *for every* θ.

B.7 may be useful when attempting to classify subjects based on their degree of conservatism because it shows that subjects may be compared with relatively few questions. Given an signal of interest, the experimenter need only elicit (conditional) certainty equivalents for particular binary acts (on θ). Under the assumption of constant conservatism (i.e., Weak Consequentialism holds), a single elicitation of a (conditional) certainty equivalent suffices to order all subjects.

2.5 Concluding Remarks

This paper characterizes a form of conservative belief updating in which the posterior belief is a convex combination of the prior and the Bayesian posterior. While our setting separates payoff-relevant states from signals by directly imposing a product structure on the state space, so that $\Omega = S \times \Theta$, all major results go through with a significantly more general state space. That is, for any finite state space Ω and algebra of events $\Sigma \subset 2^{\Omega}$, we may characterize $\mu_A = \delta_A \mu + (1 - \delta_A)\mathcal{B}(\mu, A)$.[12]

[12] See [34] for details.

Chapter 3

When Does a Rational Decision-Maker Use Bayes Factor?

3.1 Introduction

Model selection is a fundamental aspect of sophisticated decision-making. However, traditional model selection methods often lead to dynamically inconsistent updates, which conflicts with the principles of the subjective Bayesian framework. To address this issue, one approach is to establish a new behavioral foundation, as proposed by [51], or to extend the existing subjective Bayesian framework to accommodate model selection behavior. This paper adopts the latter approach.

Instead of directly choosing an action, decision-makers (DMs) often face the task of selecting a model to make decisions on their behalf. When a DM must select a model before knowing the specifics of the future decision problem, a rational approach involves forming beliefs about potential decision problems and selecting the model that maximizes subjective expected utility (minimizes expected losses). We refer to this approach as Rational Model Selection (RMS). However, common model selection criteria often lack a precise loss function (i.e., expected utility).

An intuitive method for model selection is to choose the model with the highest belief. This approach, termed Bayes Factor Model Selection (BFMS), bypasses the need to consider beliefs about decision problems (choice beliefs). While RMS and BFMS may coincide in some cases, they often differ. This paper characterizes the set of choice beliefs for which BFMS is equivalent to RMS, regardless of the underlying beliefs about uncertain states.

This suggests a way to assume that DMs select models according to BFMS without violating traditional axioms of rational decision-making and belief updating. The intuition is that when choice beliefs exhibit sufficient randomness, biases toward particular models cancel each other out. We formalize this idea and demonstrate that this randomness must satisfy specific symmetric conditions.

This is not a trivial task. A surprising counterexample arises in problems where actions draw state-dependent utilities from uniform distributions, as they do not satisfy the required condition. Instead, choice beliefs must involve an action representing the status quo (i.e., the origin), with alternative actions being perturbations of the status quo.

We apply this model to explain various phenomena that cannot be elucidated within standard frameworks without abandoning assumptions of utility maximization and belief updating via Bayes' rule. Examples include sudden halts in foreign direct investment (FDI), the significant incentive for firms to advertise in new markets, and the emergence of monopolies in industries with low entry costs.

3.2 Setup and Main Theorem

There is a relevant state space S representing payoffs. Models are probability distributions over S, denoted as $\Delta(S)$. The DM's subjective belief over states (state belief henceforth) is generated by averaging two models. For some $\beta \in [0,1]$ and models $M_1, M_2 \in \Delta(S)$, the state belief is given by $\beta M_1 + (1-\beta)M_2$. When M_1 and M_2 are clear from context, we identify this state belief as β. $\mathbb{E}^\beta$ is the expectation operator with β as the true distribution.

Acts map states to utility. For act f and state s, $f(s) \in [0,1]$ represents the utility of f in state s. The act that yields zero payoff in every state, denoted as $\mathbf{0}$, is called the origin or the status quo act, depending on the intended interpretation. A choice problem P is a pair of acts $\{f, g\}$. Let $\mathscr{P}$ be the set of possible choice problems. A choice belief γ is a probability distribution over possible choice problems, denoted as $\Delta(\mathscr{P})$.

Instead of choosing an act from the realized choice problem, the DM delegates her choice to a model. A model chooses acts that maximize its expected utility. Denote P^M as the choices induced by model M in problem P, where $P^M = \arg\max_{f \in P} \mathbb{E}^M f$.

Given models M_1, M_2, and the subjective belief β, the utility loss from selecting M_i to decide for choice problem P is

$$L(\beta, M_i, P|\{M_1, M_2\}) = \mathbb{E}^\beta\{P^{M_i} - P^\beta\}.$$

When M_1, M_2 are clear from context, we write $L(\beta, M_i, P)$ for $L(\beta, M_i, P|\{M_1, M_2\})$.

Given choice belief γ, the expected loss of choosing M_i is

$$L(\beta, M_i, \gamma) = \mathbb{E}^{\gamma}_{P \in \text{supp}(\gamma)} L(\beta, M_i, P).$$

Definition 3.1 (Agreement over a choice problem P)**.** Two models agree over a choice problem P when they both pick the same act; otherwise, they disagree.

When two models disagree over P, the one that agrees with β wins, while the other loses. The loss function treats the model β as the true model. Thus, the model agreeing with β always wins and has zero loss. Winners always exist as the true model is the average of the two. When both models agree, both win. Thus, we focus on the set of problems on which they disagree.

Some problems generate the same losses for all possible models and beliefs.

Definition 3.2 (Equivalent Choice Problems)**.** For $P, P' \in \mathscr{P}$, $P \simeq P'$ if

$$\forall \beta, M_1, M_2 \in \Delta(S), L(\beta, M_i, P|\{M_1, M_2\}) = L(\beta, M_i, P'|\{M_1, M_2\}).$$

Lemma 3.3. *For problem $P = \{f, g\}$ and constant c, all problems $P' = \{g, f\}$, $P'' = \{0, f - g\}$, $-P = \{-f, -g\}$, and $P + c = \{f + c, g + c\}$ give the same losses. We call $P', P'', -P, P + c$ equivalent problems of P.*

Definition 3.4 (Equivalent Beliefs over Choices)**.** For $\gamma, \gamma' \in \Delta(\mathscr{P})$, $\gamma \simeq \gamma'$ if

$$\forall \beta, M_1, M_2 \in \Delta(S), L(\beta, M_i, \gamma|\{M_1, M_2\}) = L(\beta, M_i, \gamma'|\{M_1, M_2\}).$$

Two beliefs γ, γ' over choice problems are equivalent when they place the same weighting on choice problems. Equivalent choice problems are treated as the same problem.

Thus, two beliefs are equivalent when they give the same loss over all possible models and confidence levels.

Definition 3.5 (Choice Belief γ with Origin)**.** For all $P \in \text{supp}(\gamma)$, $\mathbf{0} \in P$.

By 3.3, any problem $P = \{f, g\}$ is equivalent to a problem with origin $P^0 = \{\mathbf{0}, g - f\}$. Thus, any choice belief γ has an equivalent choice belief with origin γ^0. A common example of such a choice belief is a problem with a status quo act and a random alternative with the status quo state-dependent utilities normalized to 0.

Definition 3.6 (Balanced Choice Belief with Origin)**.** The belief γ such that for all $P \in \text{supp}(\gamma)$, $\mathbf{0} \in P$, and for all $P \in \text{supp}(\gamma)$, $\gamma(P) = \gamma(-P)$.

Each choice belief is equivalent to one of the balanced choice beliefs with an origin. Note that models always agree when there is a dominant act. Thus, we can ignore all those problems.

Definition 3.7 (Rational Model Selection (RMS))**.** $M_1 \succsim^R M_2$ if

$$L(\beta, M_2, \gamma | \{M_1, M_2\}) \geq L(\beta, M_1, \gamma | \{M_1, M_2\}).$$

Definition 3.8 (Bayes Factor Model Selection (BFMS))**.** $M_1 \succsim^B M_2$ if

$$\beta \geq 1 - \beta.$$

A rational DM will select the model that minimizes her expected loss given a choice belief; meanwhile, the BFMS prefers the most believed model, regardless of the choice belief. Each of them induces a preference on acts $\trianglerighteq^R$ and $\trianglerighteq^B$. In the remainder of this section, we inspect the conditions on the choice belief γ under which $\succsim^R = \succsim^B$ for any possible models and beliefs, making BFMS rational and expected utility maximizing.

The β hyperplane is generated by the normal vector β passing through the origin in the state utility space. The DM is indifferent between any act in the β hyperplane and the status quo. For any act f and true belief β, there is an act f^r that is the reflection of f to the β hyperplane. When the true belief is β, the loss of a choice problem $\{0, f\}$ incurred by the losing model is the distance of f to the β hyperplane. This leads to the following lemma.

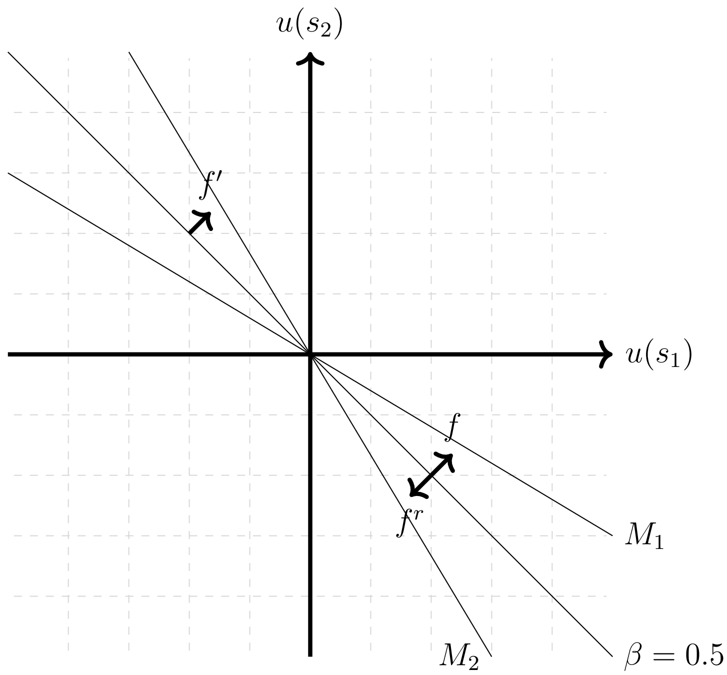

Figure 3.1: The reflection problem P^r and the equivalent problem P' of P

All vectors have equal lengths and are orthogonal to the β hyperplane.

Lemma 3.9. *If model M_1 loses to M_2 in the problem $P = \{0, f\}$ and incurs losses of L, then, in problem $P^r = \{0, f^r\}$, M_2 will either agree with M_1, or incur the same losses of L.*

Theorem 3.10. *For a choice belief γ, RMS is equivalent to BFMS if and only if there exists an equivalent choice belief such that its graph is symmetric with respect to any possible β hyperplane.*

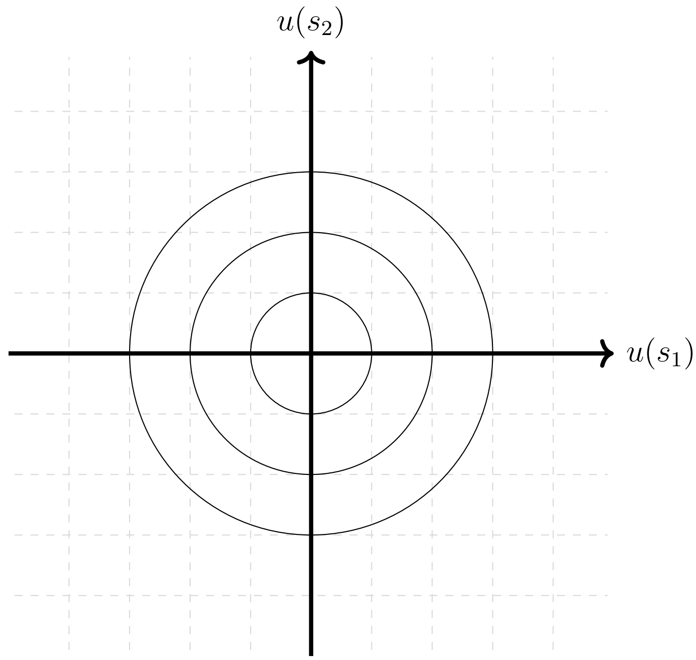

Figure 3.2: Level curves of γ that are symmetric to all possible β hyperplanes.

Proof. ($\Longleftarrow$) A direct consequence of 3.9.

($\Longrightarrow$) Suppose the condition does not hold, and the asymmetric guarantee that a problem will be weighted differently from its reflection with respect to a β-hyperplane. Thus, we can find two models such that $\beta = 0.5M_1 + 0.5M_2$, but RMS has a decisive choice over one model, say $M_1 \succ^R M_2$. Now if we change β slightly to place more weighting on the model M_2, set $\beta' = (0.5 - \epsilon)M_1 + (0.5 + \epsilon)M_2$. The suggestion of RMS maintains $M_1 \succ^R M_2$, but the BFMS will suggest otherwise, $M_1 \prec^B M_2$. □

Thus, the level curves of the distribution of γ have to be circles of radius r in the state utility space. This is not trivial. A counterexample is the choice belief generated by randomly drawing the state-dependent utility of each act from a uniform distribution. γ such that $P = \{f, g\}, \forall s \in S, f(s), g(s) \in U[0, 1]$. This choice belief yields box-like level curves. The selections of RMS are not always equivalent to the selections of BFMS. There is only one way to construct beliefs that allow us to assume the DM is selecting a model according to BFMS while being rational.

3.3 Updating and Dynamic Consistency

To see which model a rational DM would select, we need to add more structure to the framework. This section specifies how a rational DM changes their model selection according to data. For any $D \subset S$ and prior β, a DM who makes dynamically consistent choices will update her belief according to Bayes' rule. For a prior state belief $\beta = bM_1 + (1-b)M_2$, and event $E \subset D \subset S$,

$$\begin{aligned}\beta(E|D) &= \frac{\beta(E)}{\beta(D)} \\ &= \frac{bM_1(E) + (1-b)M_2(E)}{bM_1(D) + (1-b)M_2(D)} \\ &= \frac{bM_1(D)}{bM_1(D) + (1-b)M_2(D)} \frac{M_1(E)}{M_1(D)} + \frac{(1-b)M_2(D)}{bM_1(D) + (1-b)M_2(D)} \frac{M_2(E)}{M_2(D)} \\ &= b(D)M_1(E|D) + (1-b(D))M_2(E|D),\end{aligned}$$

where $b(D) = \frac{bM_1(D)}{bM_1(D)+(1-b)M_2(D)}$ is the updated belief on model M_1.

Thus, given data D, the rational DM will have true belief $\beta(D)$. Rational model selection is defined with $\beta(D)$ as the true state belief, and Bayes factor model selection is defined with the updated weighting $b(D)$ on M_1.

Suppose we observe conditional preferences on acts $\{\trianglerighteq_D\}_{D \subset S}$ that do not satisfy the dynamic consistency axiom. The conditional preferences will induce a set of state beliefs $\{\beta_D \in \Delta(S)\}_{D \subseteq S}$. Suppose we observe that the induced state beliefs are generated by Bayes factor model selection with models M_1 and M_2. By 3.10, the induced state beliefs can also be generated by RMS with the choice belief satisfying the conditions in the theorem. In this case, we can say that the DM using BFMS is rational, as she will make dynamically consistent choices if she is able to choose acts directly but not select a model to delegate her choice to.

3.4 Criterion for MS and General Scheme for Applications

The scheme for application is to compare the case of model selection with the case of choosing acts. We assume that there is a large number of agents with identical beliefs. The choice belief is consistent in that each agent will face a choice problem drawn from a distribution that coincides with the belief.

The main application is that when the agent is choosing an act from the choice problem, the mapping from belief to optimal act is continuous, then the observed behavior will not change drastically without a large change in belief. While the agent is selecting a model, the behavior will change drastically when the agent switches from one model to another according to Bayes factor.

3.4.1 Application: Sudden Stopping FDI

Assume that a continuum of agents chooses whether to invest in a foreign country (I) or not (NI). Two states indicate the foreign country's investment environment, either favorable (F) or bleak (B). The utility for not investing is zero in both states, which is the origin of the state utility space. The choice problem is random in that different agents will have different investment opportunities, which gives different payoff contingent on the investment environment of the foreign country. In state B, the payoff of investing is lower than not investing; vice versa in state F.

The state-dependent utility of investments is uniformly drawn from a quarter circle of radius r centered at the origin (not investing), with probability density function $f(\theta) = \frac{2}{\pi}$, $\theta \in [0, \frac{\pi}{2}]$, and the payoff of investing is $(I(F), I(B)) = (r\sin\theta, -r\cos\theta)$. The payoff of not investing is the origin $(NI(F), NI(B)) = (0, 0)$. Normalize r to 1.

Act\State	Favorable	Bleak
Investing	$\sin\theta$	$-\cos\theta$
Not Investing	0	0

There are two possible models, one being an optimist and the other pessimist, M_O and M_P, respectively. It could be that the optimist model anticipates normal political circumstances in a foreign country, and the pessimist model anticipates a crisis. The state belief of the agents is identical and generated by averaging the two models, $\beta = bM_O + (1-b)M_P$. We consider the following two cases: (i) the agents can make the foreign investment by themselves, and (ii) they must delegate their investment choice to an investment advisor. There are two advisors. Each dedicated advisor corresponds to one of the two models, A_O and A_P.

In case 1, the proportion of agents investing in the foreign country is continuous in the agents' belief in the optimist model b. While in case 2, the proportion of agents investing in the foreign country will drop drastically when b passes the 0.5 threshold, which corresponds to whether or not the Bayes factor is greater than 1.

3.4.2 Willingness to Pay for Advertising

A similar story can be told for the willingness to pay for advertising brands in the newborn industry. Products of different brands provide different experiences to different

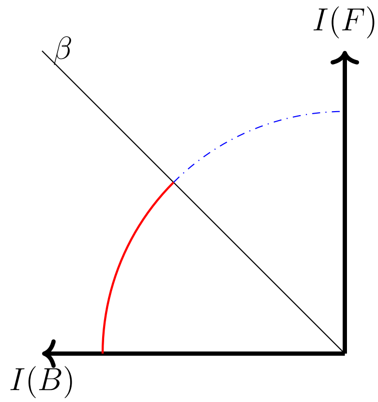

Figure 3.3: Case 1. Invest when the choice problem lies on the red thick arc, and not invest when on the blue dotted arc.

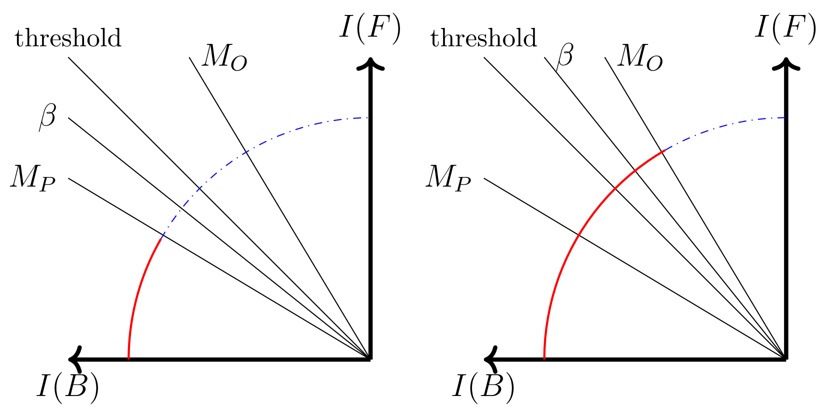

(a) $b < 0.5$, select A_O (b) $b > 0,5$, select A_P

Figure 3.4: Case 2: β passes the threshold.

users. Suppose two brands are producing a similar product. There are two models, each indicating that one brand will bring better experiences than another. Suppose the belief over two models is close to the 0.5 threshold and the brands can affect consumers' beliefs by advertising. The brands' incentive to advertise is huge, as they are competing against each other to push the consumers' beliefs across the threshold. Following an argument similar to the Bertrand competition, the two brands will keep increasing their investment in advertisement.

3.5 Discussions and Conclusion

3.5.1 Discussion on the Traditional Justification

The textbook example of justification of Bayes factor model selection criteria considers the case of a DM selecting a model to believe in, and the state space is which model is the true model that generates the data with the following payoff table.

Act \State	M_1	M_2
Selecting M_1	1	0
Selecting M_2	0	1

There are a few things that make this difficult to be compatible with the subjective Bayesian framework. First, selecting a model directly and setting the belief in another to zero is not entirely justified. Second, the payoff is arbitrary, but this can be solved by carefully constructing the loss function. Third, the subjective Bayesian framework does not define the notion of a true model. Most importantly, model selection will lead to dynamically inconsistent choices when the model selection is not a one-time choice before any decision is made. Generally, the DM making model selection decision will not just do once. Upon receiving new evidence, they update their belief and the selected model. Thus, dynamic inconsistency always results. At last, under the Bayesian framework, once the belief in a model becomes zero, it cannot turn positive again. Thus, the DM will be stuck at the selected model regardless of any new data collected.

3.5.2 Discussion on Counterexample

We formalize the idea that when the choice belief is sufficiently random, the bias induced from using the BFMS toward either model will cancel out each other. This formalization is non-trivial. The following counterexample showcases this.

There are two states $S = \{s_1, s_2\}$ and the choice problem $P = \{f_1, f_2\}$. Both have state-dependent utilities generated by the uniform distribution, for any $f \in P, s \in S$: $f(s) \sim U[0, 1]$. The level curves of this choice belief will be squares in the state utility space. In such a context, we cannot assume the DM selecting a model according to BFMS while keeping the assumption of rationality.

3.5.3 Conclusion Remarks

This paper proposes a modeling method for drastic changes in behavior without dropping the assumption of rationality. Such changes violate the dynamic consistency axiom, a foundation for Bayesian rationality. We re-frame the choice problem as a delegation problem such that the DM has to select a model to delegate their choice to. In fact, whenever the DM has to select a model, there must be a threshold belief such that she switches from one model to another when the belief passes the threshold. When the rationality of DM is assumed, one has to solve the agent's utility maximization problem for the threshold whenever new data becomes available. The theory in this paper states a condition on choice belief such that the DM will always select models according to the Bayes factor, and the threshold value is 1 (the 0.5 threshold of b) regardless of information. This simplifies the solution a lot. The model switching threshold varies upon the availability of new information if the condition does not hold. When the condition is satisfied, the rational agents select the model by whether or not the Bayes factor is greater than 1. Another advantage is that the Bayes factor is already available in the estimating process when estimating $bM_1 + (1 - b)M_2$.

Appendices

Appendix A

Appendix for the paper "Selling to Wishful Thinkers"

A.0.1 Road Map

In this proof, I will demonstrate the application of [59] to establish an implementability condition for an increasing allocation. Subsequently, I will prove that for any increasing allocations, there exists an alternative increasing allocation with greater spread. Consequently, the property of being increasing remains unchanged, and the revised allocation remains implementable. Additionally, I will establish that such a revision is advantageous to the seller. Consequently, the optimal mechanism must adhere to the icx-max criterion, which I define as the maximum spread mechanism.

Theorem A.1. *[59] If $\mathcal{Y}$ and f are regular, f satisfies the strict outer Spence-Mirrlees condition, and $\mathcal{Y}$ is a chain, then all and only increasing allocations are implementable.*

A.0.2 Proof: The Outcome Space $\mathcal{Y}$ is Regular

The outcome space $\mathcal{Y}$ is regular when it satisfies the following conditions:

1. Order-dense-in-itself: $\forall a < a' \in \mathcal{Y}$, $\exists b \in \mathcal{Y}$ such that $a < b < a'$.

2. Countably chain-complete: For every countable chain in $\mathcal{Y}$ with a lower (upper) bound in $\mathcal{Y}$, it has an infimum (a supremum) in $\mathcal{Y}$.

3. Chain-separable: For every chain $C \subset \mathcal{Y}$, there exists a countable set $B \subset \mathcal{Y}$ that is order-dense in C.

Remark: $B \subset \mathcal{Y}$ is order-dense in $C \subset \mathcal{Y}$ if for every $c < c'$ in C, there exists a $b \in B$ such that $c \leq b \leq c'$.

To align with the content, let the underlying state be $\theta \sim G : \Theta \to [0, 1]$. Any random variable is denoted as $X : \Theta \to \mathbb{R}$.

$\mathcal{Y}$ is Order-dense-in-itself

Take $a < a'$ in $\mathcal{Y}$. Since $a \neq a'$, by the definition of the icx order, there exists increasing convex function f such that $\mathbb{E}f(a) < \mathbb{E}f(a')$. By the convexity of f, $\mathbb{E}f(a) < \mathbb{E}f\left(\frac{a+a'}{2}\right) < \mathbb{E}f(a')$. We finish the proof by noticing that any mixing of two random variables is also a random variable, thus $\frac{a+a'}{2} \in \mathcal{F}$. Thus, $a < \frac{a+a'}{2} < a'$.

$\mathcal{Y}$ is Countably Chain-complete

Take the pointwise limit of the increasing chain $\{X_n\}_{n\in\mathbb{N}}$, $X(\theta) = \lim X_n(\theta)$. We attempt to show that X is the supremum. For any f increasing convex, for all n,

$$\mathbb{E}f(X) = \mathbb{E}f\left(\lim_n X_n\right) = \mathbb{E}\lim_n f(X_n) \geq \mathbb{E}f(X_n)$$

For any $\epsilon > 0$, there exists N such that for any $n > N$, and any $\theta \in \Theta$, $X(\theta) - X_n(\theta) < \epsilon$. So $X(\theta)$ is bounded and is a random variable with finite mean. Thus, it is in $\mathcal{Y}$.

Similar reasoning applies for the decreasing chain.

$\mathcal{Y}$ is Chain-separable

Take a strictly increasing convex function, f. Define $F : \mathcal{Y} \to \mathbb{R}$ as $F(X) = \mathbb{E}f(X)$. F is a strictly increasing function, $X >_{icx} Y \implies F(X) > F(Y)$. $F|_C^{-1}(\mathbb{Q} \cap F|_C(C))$ is the countable order-dense subset of C in $\mathcal{Y}$.

A.0.3 f is regular

Definition 6. The payoff function f is regular if:

1. The type derivative f_3 exists and is bounded, and $f_3(y, \cdot, t)$ is continuous for each $y \in \mathcal{Y}$ and $t \in [0, 1]$.

2. For every chain $\mathcal{C} \subseteq \mathcal{Y}$, f is jointly continuous on $\mathcal{C} \times \mathbf{R} \times [0, 1]$ when $\mathcal{C}$ has the relative topology inherited from the order topology on $\mathcal{Y}$.

[14] The order topology on $\mathcal{Y}$ is generated by the open order rays $\{y' \in \mathcal{Y} : y' < y\}$ and

$\{y' \in \mathcal{Y} : y < y'\}$ for each $y \in \mathcal{Y}$, where $<$ denotes the strict part of the order on $\mathcal{Y}$. [15] It is sufficient, but unnecessarily strong, to assume joint continuity on $\mathcal{Y} \times \mathbf{R} \times [0, 1]$.

Existence and Boundedness of f_3: By continuity of the distortion function, f_3 exists. For a fixed allocation y and payment τ, f_3 is bounded due to the continuity of δ. Additionally, $f_3(y, \cdot, t)$ is continuous for each $y \in \mathcal{Y}$ and $t \in [0, 1]$ due to the continuity of δ.

Joint Continuity of f on Chains: The derivative exists, imposing the functional derivative on $\mathcal{C}$. For a chain in the icx order, $X \geq_{icx} Y$, there exists a strictly convex increasing function k such that $\mathbb{E}k(X) > \mathbb{E}k(Y)$. It has been shown that a chain in the icx order is isomorphic to an interval in $\mathbb{R}$. We could index the chain as $\{X_i\}_{i \in I}$, where I is some interval in $\mathbb{R}$.

$$f(\mathbf{X}_i, \tau_i, \theta_i) = \frac{\mathbb{E}[u] \cdot [v \circ u](\mathbf{X}_i, \tau_i, \theta_i)}{\mathbb{E}[v \circ u](\mathbf{X}_i, \tau_i, \theta_i)}$$

The function $u(\mathbf{y}_i(\theta_{-i}), \tau_i, \theta) = q(\theta_{-i})\theta - p(\theta_{-i})\tau_i$ is jointly continuous in $(\mathbf{y}(\theta_{-i}), \tau, \theta)$ as $\mathbb{R}^4 \to \mathbb{R}$. The function v is continuous by assumption, so $v \circ u$ is continuous. Also, we assumed $G = F^{n-1}$ is continuous, so Eh($\mathbf{Z}$), where $\mathbf{Z} = (\mathbf{X}, \tau, \theta)$, is also continuous when h is continuous, and pointwise convergence is used. Let

$$\mathbf{Z}_i \to_{pt} \mathbf{Z} \Longrightarrow \mathbb{E}h(\mathbf{Z}_i) \to \mathbb{E}h(\mathbf{Z})$$

when h is continuous. icx convergence implies pointwise convergence (the function $k(\mathbf{y}, \mathbf{q}) = \mathbf{y} + \mathbf{q}$, pointwise convergence is stronger in the sense that convexity is not counted and also the underlying state ω is counted, while icx only cares about the distribution of the random variable).

$$\mathbf{Z}_i \to_{icx} \mathbf{Z} \Longrightarrow \mathbf{Z}_i \to_{pt} \mathbf{Z} \Longrightarrow \mathbb{E}h(\mathbf{Z}_i) \to \mathbb{E}h(\mathbf{Z})$$

So, f is jointly continuous, as long as $\mathbb{E}v(u(\mathbf{X}_i, \tau_i, \theta_i))$ does not equal zero, which by assumption is impossible, as v is a strictly positive increasing function.

Let $h_1 = u \cdot v \circ u$ and $h_2 = v \circ u$, where $f \cdot g$ is the pointwise product, $f \cdot g(x) = f(x) \cdot g(x)$ finishes the proof.

A.0.4 Generalized Proof of

1. **If** $\hat{\theta} = \theta$ **and** $u(\hat{\theta}, \theta, Q', T') \succ^s u(\hat{\theta}, \theta, Q, T)$**, then** $f_\theta(\theta, Q', T') > f_\theta(\theta, Q, T)$**;**

2. **FOSD preserved when (1.4);**

3. $f_\theta(\theta)$ **increases when the type increase.**

Notations: fix an type $\theta \in \Theta$, mechanism (Q, T) and (Q', T'),

$$\begin{aligned} q(\theta_{-i}) :=& Q(\hat{\theta} = \theta, \theta_{-i}), & q'(\theta_{-i}) :=& Q'(\hat{\theta} = \theta, \theta_{-i}), \\ t(\theta_{-i}) :=& T(\hat{\theta} = \theta, \theta_{-i}), \text{ and} & t'(\theta_{-i}) :=& T'(\hat{\theta} = \theta, \theta_{-i}). \\ u(\theta_{-i}) :=& u(\hat{\theta} = \theta, \theta_{-i}; \theta, Q, T), & u^e(\theta_{-i}) :=& u(\hat{\theta} = \theta, \theta_{-i}; \theta + \epsilon, Q, T), \\ u'(\theta_{-i}) :=& u(\hat{\theta} = \theta, \theta_{-i}; \theta, Q', T'), \text{ and} & u'^e(\theta_{-i}) :=& u(\hat{\theta} = \theta, \theta_{-i}; \theta + \epsilon, Q', T'). \end{aligned}$$

$\forall x_{-i} : \Theta_{-i} \to \mathbb{R}$: denote $\int_{-i} x_{-i} := \int_{\theta_{-i}\in\Theta_{-i}} x(\theta_{-i})G(d\theta_{-i})$. Also, $\phi(u) := v(u)u$.

$$\mu(\theta) := \frac{\int_{-i}\phi(u_{-i})}{\int_{-i}v(u_{-i})}, \qquad \mu(\theta+\epsilon) := \frac{\int_{-i}\phi(u^e_{-i})}{\int_{-i}v(u^e_{-i})}$$
$$\mu'(\theta) := \frac{\int_{-i}\phi(u'_{-i})}{\int_{-i}v(u'_{-i})}, \text{ and} \qquad \mu'(\theta+\epsilon) := \frac{\int_{-i}\phi(u'^e_{-i})}{\int_{-i}v(u'^e_{-i})}.$$

We study variation in mechanism (Q, T), (Q', T') mean-preserving spread of (Q, T) for type θ, that $q' \succ^s q$, and $u' \succ^s u$.

We define the following variation funciton:

$$u^e - u := \epsilon\eta^e, \qquad u' - u := \epsilon'\eta'_{-i}, \text{ and}$$
$$u'^e - u' := \epsilon\eta'^e. \qquad \text{Those implies } u'^e - u = \epsilon'\eta' + \eta'^e\epsilon.$$

By definition $\eta'^e = q'$, and $\eta^e = q$.

We take first order approximation of ϕ and v around u, with variation η's as defined as above. For $u^x = u + \eta^x\epsilon$, for small ϵ. $\frac{\int_{-i}\phi(u^x_{-i})}{\int_{-i}v(u^x_{-i})} = \frac{\int_{-i}\phi(u_{-i})+\int_{-i}\phi'(u_{-i})\eta^x_{-i}\epsilon}{\int_{-i}v(u_{-i})+\int_{-i}v'(u_{-i})\eta^x_{-i}\epsilon}$. With for random variable X, and Y, let C be the covariance function $C(X, Y) = \mathbb{E}XY - \mathbb{E}X\mathbb{E}Y$. Note that $\eta' = u' - u$, and u' is the mean-preserving variation of u, $\mathbb{E}\eta' = \int_{-i}\eta'_{-i} = 0$. Also, $\eta'^e = q'$ and $\eta^e = q$, with q' is an mean preserving spread of q, $\mathbb{E}\eta'^e = \mathbb{E}q' = \mathbb{E}q = \mathbb{E}\eta^e$.

We have

$$\mu(\theta) = \frac{a}{b}, \qquad \mu(\theta+\epsilon) = \frac{a+e_1}{b+e_2},$$
$$\mu'(\theta) = \frac{a+e_3}{b+e_4}, \text{ and} \qquad \mu'(\theta+\epsilon) = \frac{a+e_3+e_5}{b+e_4+e_6},$$

where $e_1 = [\mathbb{E}\phi'(u)\mathbb{E}q + C(\phi'(u), \eta^e)]\epsilon$, $e_2 = [\mathbb{E}v'(u)\mathbb{E}q + C(v'(u), \eta^e)]\epsilon$, $e_3 = C(\phi'(u), \eta')\epsilon'$, $e_4 = C(v'(u), \eta')\epsilon'$, $e_5 = [\mathbb{E}\phi'(u)\mathbb{E}q + C(\phi'(u), \eta'^e)]\epsilon$, and $e_6 = [\mathbb{E}v'(u)\mathbb{E}q + C(v'(u), \eta'^e)]\epsilon$.

Proof that $f_\theta(\theta|Q', T') > f_\theta(\theta|Q, T)$**.**

By definition $f_\theta(\theta|Q, T) = \lim_{\epsilon\to 0} \frac{\mu(\theta+\epsilon)-\mu(\theta)}{\epsilon}$, and $f_\theta(\theta|Q', T') = \lim_{\epsilon\to 0} \frac{\mu'(\theta+\epsilon)-\mu'(\theta)}{\epsilon}$. Thus, we show that $f_\theta(\theta|Q', T') \geq f_\theta(\theta|Q, T)$, adn ignore the 2nd order ϵ, ϵ' terms.

This turn out requires

$$a(e_6 - e_2) < b(e_5 - e_1).$$

$b > a$ since $u \leq 1$. The sufficent condition for the condition becomes

$$\begin{aligned}
& & e_5 - e_1 >& e_6 - e_2 \\
\Leftrightarrow & & C(\phi'(u) - v'(u), \eta'^e) >& C(\phi(u) - v'(u), \eta^e), \\
\Leftrightarrow & & C(\phi'(u) - v'(u), \eta'^e - \eta^e) >& 0, \\
\Leftrightarrow & & C(\phi'(u) - v'(u), q' - q) >& 0, \\
\Leftrightarrow & & \mathbb{E}\{(\phi'(u) - v'(u))(q' - q)\} >& (\mathbb{E}\{\phi'(u) - v'(u)\})(\mathbb{E}q' - \mathbb{E}q), \\
\Leftrightarrow & & \mathbb{E}\{(\phi'(u) - v'(u))(q' - q)\} >& 0.
\end{aligned}$$

Recall that by (1.3), $\phi'(u_i) - v'(u_i)$ is increasing in u_i, and (Q', T') is a mean-preserving spread of (Q,T) for θ, thus, for u_{-i} large $q'-q$ is also large. In general, the order of $\phi'-v'$ in line with $q'-q$. This shows that $C(\phi'(u)-v'(u), q'-q) > 0$. ■

Generalized A.1.: Preference is monotonic with respect to FOSD when (1.4)

Note that in previous we have $\eta^e_{-i} = q_{-i} \geq 0$; if we replace it by some function that is weakly positive, $\eta^e_{-i} = \eta_{-i} \geq 0$. $u^e = u + \eta\epsilon$ becomes any random utility that FOSD u by same variation in the direction η. Pick any η that is weakly positive and let $\eta^e = \eta$. We have to proof $\lim_{\epsilon \to 0} \frac{\mu(\theta+\epsilon)-\mu(\theta)}{\epsilon} \geq 0$.

$$\mu(\theta+\epsilon) \geq \mu(\theta),$$

$$\Leftrightarrow \qquad be_1 \geq ae_2.$$

Since $b \geq a$, $e_1 \geq e_2$ is suffice.

$$\begin{aligned} & e_1 \geq e_2, \\ \Leftrightarrow \quad & \mathbb{E}\phi'(u)\mathbb{E}\eta + C(\phi'(u),\eta) \geq \mathbb{E}v'(u)\mathbb{E}\eta + C(v'(u),\eta), \\ \Leftrightarrow \quad & \mathbb{E}\{(\phi'(u)-v'(u))\eta\} \geq 0. \end{aligned}$$

This is direct since $\forall\theta_{-i}, \eta(\theta_{-i}) \geq 0$ and by (1.4), $\forall u_i : \phi'(u_i) \geq v'(u_i)$.

Generalized proof of A.2.: $\mu'(\theta) \geq \mu(\theta)$. Change of distorted utility for type θ from replacing the mechanism (Q,T) by a variation mechanism (Q',T') that is mean-preserving spread of (Q,T) for θ. By definition, we are studying small variations:

$$\lim_{\epsilon'\to 0}\frac{\mu'(\theta)-\mu(\theta)}{\epsilon'} \geq 0$$
$$\Leftrightarrow \qquad be_3 \geq ae_4$$

This is ture since $b \geq a$. $e_3 \geq e_4$ is suffice. This is automatic by (1.4).

A.0.5 Proof of Lemma 5: Given a mechanism (Q,T) that give $V(\theta), f_\theta(\theta)$, the designer can reduce $V(\theta), f_\theta(\theta)$ by reducting $Q(\theta,\cdot)$ or increasing $T(\theta,\cdot)$. Both have different relatve effect on $V(\theta)$ and $f_\theta(\theta)$.

Fix an mechanism (Q,T) and a report $\hat{\theta}$ and a type θ. Let q,t be $q(\cdot) = Q(\hat{\theta},\cdot)$, $u(\cdot) = Q(\hat{\theta},\cdot)\theta - tT(\hat{\theta},\cdot)$, $v(\cdot) = v(u(\cdot))$, and $v'(\cdot) = \frac{\partial v}{\partial u}(\cdot)$. Instead of studying the effect of deviation of q,t on V, f_θ. We study the effect of (u,q) on V,q and use total derivative to find the effect of (q,t) on V, f_θ. Define the functional $V(u) = \frac{\mathbb{E}_{-i}\{uv\}}{\mathbb{E}_{-i}\{v\}}$ and $f_\theta(u,q) = \frac{\mathbb{E}_{-i}\{q(v+uv')\}\mathbb{E}_{-i}\{v\}-\mathbb{E}_{-i}\{qv'\}\mathbb{E}_{-i}\{uv\}}{(\mathbb{E}_{-i}\{v\}^2)}$; while $\hat{V}(q,t) = V(\theta q - t)$ and $\hat{f}_\theta(q,t) = f_\theta(q\theta - t, t)$.

Denote $B_1(u,q) = \mathbb{E}_{-i}\{q(v+uv')\}$, $B_2(u,q) = \mathbb{E}_{-i}\{v\}$, $B_3(u,q) = \mathbb{E}_{-i}\{qv'\}$, $B_4(u,q) = \mathbb{E}_{-i}\{uv\}$, and define the integrand $b_i(\alpha,\beta)$ by $B_i(u,q) = \int_{\theta_{-i}\in\Theta_{-i}} b_i(u(\theta_{-i}), q(\theta_{-i}))G(d\theta_{-i})$. Also let $B = B_1B_2 - B_3B_4$Using the notation we have $V = \frac{B_4}{B_2}$, $f_\theta = \frac{B}{(B_2)^2}$.

For functional F that take functions a,b as input, and deviation function η, denote the deviative of F with respect to a by function η as $\frac{\delta F}{\delta a}(a,b;\eta) = F_a^\eta(a,b) =$

$\lim_{\epsilon \to 0} \frac{F(a+\eta\epsilon,b)-F(a,b)}{\epsilon}$.

By A.1., we have for all $\eta \geq 0$, $V_u^\eta(u) > 0$, and by A.3. We have for all $\eta \geq 0$, $f_{\theta,u}^\eta(u,q) \geq 0$. Both are ture since if $\eta \geq 0$, $u + \eta$ FOSD u. By chain rule we have

$$\begin{aligned}
\hat{V}_q^\eta(q,t) =& \theta V_u^\eta(u) \\
\hat{V}_t^\eta(q,t) =& - V_u^\eta(u) \\
\hat{f}_{\theta,q}^\eta(q,t) =& \theta f_{\theta,u}^\eta(u,q) + f_{\theta,q}^\eta(u,q) \\
\hat{f}_{\theta,t}^\eta(q,t) =& - f_{\theta,u}^\eta(u,q)
\end{aligned}$$

, where $u = q\theta - t$.

Note that B_2 and B_4 are independent of q (with u fixed):

$$f_{\theta,q}^\eta = \frac{B_{1,q}^\eta B_2 - B_{3,q}^\eta B_4}{(B_2)^2}.$$

Recall that $B_2(u,q) = \mathbb{E}_{-i}\{v\}$ and $B_4 = \mathbb{E}_{-i}\{vu\}$, as we have $u \leq 1$, we know that $B_2 \geq B_4$. Also, $B_{1,q}^\eta(u,q) = \mathbb{E}_{-i}\{\eta(v + uv')\}$ and $B_{3,q}^\eta(u,q) = \mathbb{E}_{-i}\{\eta v'\}$. For $\eta \geq 0$, (1.4) implies $B_{1,q}^\eta \geq B_{3,q}^\eta$. Thus, we have

$$\forall \eta \geq 0 : f_{\theta,q}^\eta \geq 0.$$

Summarizing we have for every $\eta^q \leq 0$, $\eta^t \geq 0$.

$$\begin{aligned}
\hat{V}_q^{\eta^q}(q,t) =&\theta V_u^{\eta}(u) \le 0;\\
\hat{V}_t^{\eta^t}(q,t) =&-V_u^{\eta^t}(u) \le 0;\\
\hat{f}_{\theta,q}^{\eta^q}(q,t) =&\theta f_{\theta,u}^{\eta^q}(u,q) + f_{\theta,q}^{\eta^q}(u,q) \le 0;\\
\hat{f}_{\theta,t}^{\eta^t}(q,t) =&-f_{\theta,u}^{\eta^t}(u,q) \le 0.
\end{aligned}$$

Also, event it two variation have the same impact on utility, the relative effect is different

$$\frac{\hat{V}_q^{\theta\eta}}{\hat{f}_{\theta,q}^{\theta\eta}} \neq \frac{\hat{V}_t^{-\eta}}{\hat{f}_{\theta,t}^{-\eta}}.$$

A.0.6 Condition 5 $\implies$ Monotonicity

The single-crossing difference property states that if the expected utility of a change in allocation and prices is non-decreasing with respect to a bidder's type θ, then it remains non-decreasing for any higher type $\theta' \ge \theta$. Formally,

$$E^\delta U(\mathbf{q}',\mathbf{p}',\tau',\theta) - E^\delta U(\mathbf{q},\mathbf{p},\tau,\theta) \ge 0 \implies E^\delta U(\mathbf{q}',\mathbf{p}',\tau',\theta') - E^\delta U(\mathbf{q},\mathbf{p},\tau,\theta') \ge 0$$

for every $\theta' \ge \theta$, $(\mathbf{q}',\mathbf{p}') \ge (\mathbf{q},\mathbf{p})$, and any τ,τ'.

Consider a fixed report $\hat{\theta}$. Let $\mathbf{q}$ denote the random variable generated by $Q(\cdot,\hat{\theta})$, and similarly for $\mathbf{p}$. For notation ease, define $\mathbf{m}' := (\mathbf{q}',-\mathbf{p}') \ge_{\text{ICX}} (\mathbf{q},-\mathbf{p}) := \mathbf{m}$. Then there exist $\hat{\mathbf{m}}' =_{\text{ST}} \mathbf{m}'$ and $\hat{\mathbf{m}} =_{\text{ST}} \mathbf{m}$ defined on the same probability space such that $(\mathbf{m},\mathbf{m}')$ is a submartingale.

Taking the first-order approximation of each integrand in $E^\delta U$ function, we have

$$E^\delta U(\mathbf{m}',\tau',\theta) - E^\delta U(\mathbf{m},\tau,\theta) \approx \frac{\mathbb{E}\{\hat{\phi}(\mathbf{m},\tau,\theta) + \sum_x \hat{\phi}'_x(\mathbf{m},\tau,\theta)(x'-x)\}}{\mathbb{E}\{\hat{v}(\mathbf{m},\tau,\theta) + \sum_x \hat{v}'_x(\mathbf{m},\tau,\theta)(x'-x)\}} \quad \text{(A.1)}$$

$$- \frac{\mathbb{E}\{\hat{\phi}(\mathbf{m},\tau,\theta)\}}{\mathbb{E}\{\hat{v}(\mathbf{m},\tau,\theta)\}} \quad \text{(A.2)}$$

where $x \in \{q, p, \tau\}$.

We then take the first-order linear approximation of each integrand again for an increase in θ, yielding

$$E^\delta U'_\theta(\mathbf{m}',\tau',\theta) - E^\delta U'_\theta(\mathbf{m},\tau,\theta) \approx \quad \text{(A.3)}$$

$$\frac{\mathbb{E}\{\hat{\phi}'_\theta(\hat{\mathbf{m}},\tau,\theta) + \sum_x \hat{\phi}''_{x,\theta}(\hat{\mathbf{m}},\tau,\theta)(x'-x)\}}{\mathbb{E}\{\hat{v}'_\theta(\hat{\mathbf{m}},\tau,\theta) + \sum_x \hat{v}''_{x,\theta}(\hat{\mathbf{m}},\tau,\theta)(x'-x)\}} - \frac{\mathbb{E}\{\hat{\phi}'_\theta(\hat{\mathbf{m}},\tau,\theta)\}}{\mathbb{E}\{\hat{v}'_\theta(\hat{\mathbf{m}},\tau,\theta)\}} \quad \text{(A.4)}$$

We compare the difference before and after $\partial\theta$. With $\hat{\phi}'_\theta \geq \hat{v}'_\theta$ and $u(\theta) = \theta q - p\tau$, we have $\hat{\phi}''_{q,\theta} = \phi' + \phi'' u'_q u'_\theta$ and $\hat{\phi}''_{p,\theta} = \phi'' u'_p u'_\theta$. By assumption $\phi'' \geq v''$. With the submartingale property of $\geq_{\text{ICX}}$, $\mathbf{m}' \geq_{\text{ICX}} \mathbf{m}$. Moreover, $E\{\mathbf{q}' - \mathbf{q}|\mathbf{q}\} > 0$ and $E\{\mathbf{p}' - \mathbf{p}|\mathbf{p}\} < 0$. Thus, the increase in the numerator is greater than the increase in the denominator when

$$E\{[v(u) + uv'(u)](q'-q)\} \geq E\{[2v'(u) + uv''(u)]qp\}(\tau'-\tau) - E\{[v''(u)]qp\}(\tau'-\tau).$$

We impose this as an additional restriction to the order. The order represents the direction of preferred outcomes for increasing type. The RHS arises from the fact that when payment is conditional on some good outcome $q > 0$, higher types will distort towards the event more and have a higher perceived payment. When the budget is high enough, as the variation of u increases, $\{\theta_{-i} \in \Theta_{-i} | \mathbf{p}(\theta_{-i}) > 0\} \cap \{\theta_{-i} \in \Theta_{-i} | \mathbf{q}(\theta_{-i}) > 0\} = \emptyset$. Thus, $E\{[\phi'' - v''](\mathbf{u}(\theta))\mathbf{qp}\}$ will reduce to zero. The restriction becomes $E\{[v(u) + uv'(u)](q' - q)\} \geq 0$. A sufficient condition is $\mathbf{q}' \geq_{\text{ST}} \mathbf{q}$, which mirrors the traditional monotonicity condition $EQ(\theta') \geq EQ(\theta)$ when $\theta' \geq \theta$.

A.1 Proof: The Motivating Example

Results for auctions with SEU bidders are commonly known. For FPA, both bidders bids according to strategy $\beta_i^{FPA,SEU} = \frac{1}{2}\theta_i$. For SPA, both bidders bid their own value $\beta_i^{SPA,SEU} = \theta_i$. Both revenue are equivalent.

Now, consider the case of biased bidders. The correct belief on other's type is $\theta_j \sim G = U[0,1]$, with the density $g(\theta_j) = 1$ for every $\theta_j \in [0,1] := \Theta_j$. Cosider a simple distortion that place probability δ on the best case that the other have zero valuation. Formally, the distorted belief $G^\delta : \Theta_j \to [0,1]$ is

$$G^\delta(\theta_j) = \delta \text{ for } \theta_j = 0, \text{ and } g^\delta(\theta_j) = 1 - \delta \text{ for } \theta_j \in (0,1],$$

where g^δ is the denesity of the distorted belief G^δ when $\theta_i \in (0,1]$.

Consider the symmetric Bayesian Nash equilibrium (BNE) in an FPA. To make sure equilibrium exists, we assume that the minimum bid is $\underline{b}$ and consider the limit case that $\underline{b} \to 0$.

Given θ_1 and player 2's strategy $\beta_2(\theta_{-i})$ the expected utility for bidding $b_1 \geq \underline{b}$ is

$$\begin{aligned}
u_1(b_1|\theta_1) &= G^\delta(\beta_2(\theta_2) < b_1)(\theta_1 - b_1), \\
&= G^\delta(\theta_2 < \beta_2^{-1}(b_1))(\theta_1 - b_1), \\
&= (1-\delta)\beta_2^{-1}(b_1)\theta_1 - (1-\delta)\beta_2^{-1}(b_1) + \delta(\theta_1 - b_1).
\end{aligned}$$

FOC gives

$$(1-\delta)\beta_2'(\beta_2^{-1}(b_1))^{-1}\theta_1 = (1-\delta)[\beta_2'(\beta_2^{-1}(b_1))^{-1}b_1 + \beta_2^{-1}(b_1)] + \delta.$$

By the symmetry $\beta_1 = \beta_2 = \beta$, and $b_1 = \beta_1(\theta_1) = \beta(\theta)$. Thus, the FOC condition becomes

$$\begin{aligned}
\theta - \frac{\delta}{1-\delta} &= \beta'(\theta)\theta + \beta(\theta) = \frac{d}{d\theta}\beta(\theta)\theta, \\
\beta(\theta)\theta &= \int_0^\theta (x - \frac{\delta}{1-\delta})dx, \\
\beta(\theta) &= \frac{1}{2}\theta - \frac{\delta}{1-\delta}.
\end{aligned}$$

Thus, the BNE symmetric equibrium is

$$\beta^{FPA,WT}(\theta) = \begin{cases} \frac{1}{2}\theta - \frac{\delta}{1-\delta} & \text{if } \theta \geq \underline{b} \text{ and } \theta \geq \frac{2\delta}{1-\delta}, \\ \underline{b} & \text{if } \theta \geq \underline{b} \text{ and } \theta < \frac{2\delta}{1-\delta}, \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to see that $\beta^{FPA,WT} < \beta^{FPA,SEU}$, and the difference vanishes as $\delta = 0$. Thus, the revenue of the seller reduces when the bidders are subject to WT bias.

In SPA, the strategy $\beta(\theta) = \theta$ is weakly dominant regardless of belief.

Appendix B

Appendix for the paper "Conservative Updating"

B.1 Model and Setup

There is a (nonempty) finite set S of states of the world, a collection of events given by an algebra Σ over S, and a (nonempty) set of consequences X. Let $\mathcal{F}$ denote the set of functions $f : S \to X$, referred to as an act.[1] Following a standard abuse of notation, for any $x \in X$, by $x \in \mathcal{F}$ I mean the constant act that returns x in every state. Lastly, for any $f, g \in \mathcal{F}$ and for any $A \in \Sigma$, let fAg denote the act that returns $f(s)$ when $s \in A$ and returns $g(s)$ when $s \in A^c \equiv S \backslash A$. Following the literature, I assume that X is a convex subset of a vector space.[2] Thus, mixed acts can be defined point-wise, so that for every $f, g \in \mathcal{F}$ and $\lambda \in [0, 1]$, by $\lambda f + (1 - \lambda)g$ I mean the act

[1] It is not difficult to extend to an infinite state space by assuming Σ is a σ-algebra and $\mathcal{F}$ is the set of finite-valued Σ-measurable functions.

[2] For instance, X may be an interval of monetary prizes or, as in the classic [2] setting, a set of lotteries over some prize space.

that returns $\lambda f(s) + (1-\lambda)g(s)$ for each $s \in S$.

I assume that the agent has preferences over $\mathcal{F}$ conditional on her information. Formally, the agent has a collection of preference relations, $\{\succsim_A\}_{A\in\Sigma}$ over $\mathcal{F}$, where $\succsim_A$ are her preferences after observing A. Let $\succ_A$ and $\sim_A$ represent the asymmetric and symmetric parts of $\succsim_A$. The case when the agent has no information is represented by $\succsim_S:=\succsim$. For an event A, say that A is $\succsim$-null (or simply null) if $fAg \sim g$ for any $f, g \in \mathcal{F}$. Otherwise, A is non-null. Let $\Sigma_+ \subset \Sigma$ denote the collection of non-null events.

Let $\Delta(S)$ denote the set of probability measures over S. Typical elements, $\mu, \pi \in \Delta(S)$ are called beliefs. For any probability μ and non-null event $A \in \Sigma_+$, define the Bayesian update of μ given A by $\mathcal{B}(\mu, A)(B) = \frac{\mu(B\cap A)}{\mu(A)}$ for $B \in \Sigma$. Finally, for any two utility functions $u, v : X \to \mathbb{R}$, say $u \approx v$ if u is a positive affine transformation of v.

Definition B.1. Say that a collection of preferences $\{\succsim_A\}_{A\in\Sigma}$ admits a **conservative subjective expected utility** (conservative SEU) representation if there are a non-constant utility function $u : X \to \mathbb{R}$, a prior probability $\mu \in \Delta(S)$, and a function $\delta : \Sigma_+ \to [0, 1]$ such that for all $A \in \Sigma_+$,

$$f \succsim_A g \Longleftrightarrow \sum_{s\in S} u(f(s))\mu_A(s) \geq \sum_{s\in S} u(f(s))\mu_A(s)$$

and

$$\mu_\theta = \delta(\theta)\mu + (1-\delta(\theta))\mathcal{B}(\mu, \theta) \tag{B.1}$$

B.2 Behavioral Foundations

Axiom 10 (Conditional SEU)**.** For each $A \in \Sigma$, $\succsim_A$ admits a non-degenerate subjective expected utility representation.

Axiom 11 (Dynamic Consistency)**.** For any $A \in \Sigma_+$, and $f, g \in \mathcal{F}$,

$$fAg \succsim g \implies f \succsim_A g.$$

Axiom 12 (Dynamic Conservatism)**.** For any $A \in \Sigma_+$ and for all $f, g \in \mathcal{F}$,

$$\left.\begin{array}{ll}\text{(i)} & f \succsim g \\ \text{(ii)} & fAg \succsim g\end{array}\right\} \implies f \succsim_A g.$$

Further, if both (i) and (ii) are strict, then $f \succ_A g$.

Theorem B.2. *The following are equivalent:*

(i) The collection $\{\succsim_A\}_{A\in\Sigma}$ satisfies Conditional SEU and Dynamic Conservatism;

(ii) The collection $\{\succsim_A\}_{A\in\Sigma}$ admits a conservative subjective expected utility representation.

Further, if (u, μ, δ) and (u', μ', δ') both represent $\{\succsim_A\}_{A\in\Sigma}$, then (i) u' is a positive affine transformation of u, (ii) $\mu' = \mu$, and (iii) $\delta'(A) = \delta(A)$ for all A such that $A, A^c \in \Sigma_+$.

B.3 Degrees of Conservatism

B.3.1 A Conservatism Index

In the closing remarks of the main text, we discuss the pros and cons of the signal structure and the more general structure presented here. The main difference is that when the signal states begin segregated from the payoff-relevant states, acts (plans) with the same payoff within a signal become the same act. That is when $f, g \in \mathcal{F} \subset \mathcal{P} : f(\omega) = g(\omega)$ for every $\omega \in \theta$, then $f = g$. Thus, we could not discuss consequentialism in the main text.

We can discuss how conservatism conflicts with consequentialism in the more general structure here. The main difference between the two structures is the Indifferent Conservatism in the main text becomes the Weak Consequentialism in below. Both of them generate the constant conservative index for their corresponding structure.

Axiom 13 (Indifferent Conservatism)**.** For any $f, f' \in \mathcal{F}$, $\theta, \theta' \in \Theta$, and $x \in X$ such that $f \sim f'$ and $f\theta x \sim x \sim f'\theta' x$, there exists $y \in X$ such that $f \sim_\theta y \sim_{\theta'} f'$.

Axiom 4' (Weak Consequentialism)**.** For any $A, B, C \in \Sigma_+$ with $C \cap (A \cup B) = \varnothing$ and for all $f, y, z \in \mathcal{F}$,

$$fCy \succsim_A z \Longleftrightarrow fCy \succsim_B z.$$

Axiom 14 (Consequentialism)**.** For all $A \in \Sigma$, and $f, g \in \mathcal{F}$,

$$f \sim_A fAg.$$

Weak Consequentialism is a weak form of Consequentialism, suppose $C \cap (A \cup B) = \varnothing$

and that Consequentialism holds. Consider fCy for an arbitrary f. Since for all $s \in A \cup B$, $fCy(s) = fCy(s)$, it follows from Consequentialism that both $fCy \sim_A y$ and $fCy \sim_B y$. Thus, under Consequentialism, f is irrelevant, and under ordinal preference consistency, Weak Consequentialism always holds. On the other hand, Weak Consequentialism does not impose indifference between fCy and y but requires a consistent relative preference; if fCy is preferred to z after A, then it is also preferred after B. This highlight the fact that a conservative agent may violate consequentialism that the nullified state $\omega \in C$ being considered as possible.

In the signal information structure, Weak Consequentialism is being translated and weakened to Indifferent Conservatism. Observing that with the signal information structure Weak Consequentialism can be written as follows. For any $\theta_A, \theta_B, \theta_C \in \Theta$ with $\theta_C \cap (\theta_A \cup \theta_B) = \varnothing$ and for all $f, f', y, z \in \mathcal{F}$, and $f = f'$

$$f\theta_C y \succsim_{\theta_A} z \Longleftrightarrow f'\theta_C y \succsim_{\theta_B} z.$$

Which is equivalent to

$$f\theta_C y \sim_{\theta_A} x \sim_{\theta_B} f'\theta_C y.$$

The consequentialism require $f\theta_C y$ to be indifferent to y conditional on events that nullify θ_C. Weak Consequentialism weakens Indifferent Conservatism to allow for the updated preferences having $f\theta_C y$ not indifferent to y, but just requiring the value placed on $f\theta_C y$ conditional on any event nullifies θ_C to be equal. It means an equal effect on $f\theta_C y$ of information that are not compatible with θ_C. As discussed in the closing remark in the main text, when we focus on $\mathcal{F}$, this interpretation is not feasible. We weaken it by allowing for $f \neq f'$, but only that the two acts are ex-ante and contingently indifferent (i.e. $f \sim f'$ and $(f, \theta_A) \sim^c (f', \theta_B)$). This becomes a weakening of Dynamic Consistency. Dynamic Consistency applies to the $\sim$ relation gives $fAg \sim g \implies f \sim_A g$, and $f'Bg \sim g \implies f' \sim_B g$. In turn implies $fAg \sim g \sim f'Bg \implies f \sim_A g \sim_B f'$.

Proposition B.3. *Suppose the collection $\{\succsim_A\}_{A\in\Sigma}$ admits a Conservative SEU representation (u, μ, δ). Then the following are equivalent:*

(i) The collection $\{\succsim_A\}_{A\in\Sigma}$ satisfies Weak Consequentialism.

(ii) There is a unique $\delta \in [0,1]$ such that $\delta(A) = \delta$ for all A such that A and $S \setminus A$ are non-null.

B.3.2 Generalized Confirmation Bias

Definition B.4. For any $A, B \in \Sigma_+$, say that A is **more likely than** B, denoted $A \geq_l B$ if for all $x, y \in X$, $x \succsim y$ implies $xAy \succsim xBy$.

As in the main text, the following axiom makes sure that the same piece of bad news C has greater impact on xCy when it is delivered by a more likely event.

Axiom 15 (Generalized Confirmation Bias)**.** For any $A, B, C \in \Sigma_+$ with $C \cap (A \cup B) = \varnothing$ and for all $x, y, z \in \mathcal{F}$ if $x \succ y$ and $A \geq_l B$, then

$$xCy \succsim_A z \implies xCy \succsim_B z.$$

Proposition B.5. *Suppose the collection $\{\succsim_A\}_{A\in\Sigma}$ admits a Conservative SEU representation (u, μ, δ). Then the following are equivalent:*

(i) The collection $\{\succsim_A\}_{A\in\Sigma}$ satisfies Generalized Confirmation Bias.

(ii) $\mu(A) \geq \mu(B)$ if and only if $\delta(A) \leq \delta(B)$.

B.3.3 Comparative Conservatism

Similar as above section, the comparative conservatism can also defined without referring to the contextual ordering. The main intuition is to reserve the meaning that information affect the more conservative agent less.

Definition B.6. Say that $\{\succsim_A^1\}_{A\in\Sigma}$ is **more conservative** than $\{\succsim_A^2\}_{A\in\Sigma}$ if for all A and all $x, y_1, y_2 \in X$ satisfying (i) $x \succ^i y_i$ and (ii) $xAy_1 \succsim^1 z \Longleftrightarrow xAy_2 \succsim^2 z$ for all $z \in X$

$$xAy_1 \succsim_A^1 z \implies xAy_2 \succsim_A^2 z.$$

Proposition B.7. *Suppose* $(\{\succsim_A^i\}_{A\in\Sigma})_{i=1,2}$ *admit Conservative SEU representations* $(u_i, \mu_i, \delta_i)_{i=1,2}$ *where* $u_1 \approx u_2$ *and* $\Sigma_+^1 = \Sigma_+^2$*. Then the following are equivalent:*

(i) $\{\succsim_A^1\}_{A\in\Sigma}$ *is more conservative than* $\{\succsim_A^2\}_{A\in\Sigma}$*.*

(ii) $\delta_1(A) \geq \delta_2(A)$ *for every* A *such that* A *and* $S \setminus A$ *are non-null.*

B.4 Multiple Beliefs

Definition B.8. A preference relation $\succsim$ admits a multi-prior expected utility representation if there are a utility $u : X \to \mathbb{R}$ and a nonempty, closed, convex set of beliefs $\mathcal{M} \subseteq \Delta(S)$ such that for all acts $f, g \in \mathcal{F}$,

$$f \succsim g \Longleftrightarrow \sum_{s\in S} u(f(s))\mu(s) \geq \sum_{s\in S} u(g(s))\mu(s) \text{ for every } \mu \in \mathcal{M}. \tag{B.2}$$

Axiom 16 (Conditional Multi-prior Expected Utility)**.** For each $A \in \Sigma_+^*$, $\succsim_A^*$ admits a nondegenerate multi-prior expected utility representation. That is, there exists a pair $(u_A, \mathcal{M}_A)$ that represents $\succsim_A^*$ as in Equation B.2.

As in the main text, the following axiom makes sure the risk attitudes are unchanged by information.

Axiom 17 (Unambiguous Dynamic Conservatism)**.** For any $A \in \Sigma_+^*$, and for all $f, g \in \mathcal{F}$

$$\left.\begin{matrix} f \succsim^* g \\ fAg \succsim^* g \end{matrix}\right\} \implies f \succsim_A^* g.$$

Further, if both (i) and (ii) are strict, then $f \succ_A^* g$.

The following theorem is the direct counterpart of B.2 for multiple priors.

Theorem B.9. *The following are equivalent:*

(i) The collection $\{\succsim_A^*,\}_{A\in\Sigma}$ satisfies Conditional Multi-prior Expected Utility and Unambiguous Dynamic Conservatism.

(ii) There is a non-constant utility function $u : X \to \mathbb{R}$ such that $u_A = u$ for every $A \in \Sigma$, and for each $A \in \Sigma_+^*$,

$$\mathcal{M}_A \subseteq conv\left(\mathcal{M} \cup \mathcal{B}(\mathcal{M}, A)\right). \tag{B.3}$$

In this case, we say the agent admits a ***conservative multi-prior representation****.*

B.5 Proofs

The proof of theorem 1 and 2 are just replacing notations and also take care of the cases of null events. For proposition 1, 2, 3, we present an alternative proof that by pass the contextual ordering and the contingent preference.

B.5.1 Preliminary Results

Consider the following two properties for a binary relation $\succsim$ on $\mathcal{F}$.

C-Completeness: For any $x, y \in \mathcal{F}$, either $x \succsim y$ of $y \succsim x$.

Monotonicity: If $f(s) \succsim g(s)$ for all $s \in S$, then $f \succsim g$.

Lemma B.10. *Consider a collection of preferences $\{\succsim_A\}_{A\in\Sigma}$ such that (i) $\succsim$ satisfies C-Completeness and Monotonicity and (ii) the collection satisfies Dynamic Conservatism. Then for all $A \in \Sigma$ such that A is non-null, and any $x, y \in X$, $x \succsim y \Longleftrightarrow x \succsim_A y$.*

Proof. Since $\succsim$ is complete for constant acts, suppose $x \succsim y$. By Monotonicity of $\succsim$ this is equivalent to $xAy \succsim y$ for all A. Then by Dynamic Conservatism, $x \succsim_A y$. Suppose that $x \succsim_A y$ but $y \succ x$. Then it follows from Monotonicity and the fact that A is non-null that $yAx \succ x$. From Dynamic Conservatism it follows that $y \succ_A x$, a contradiction. Hence $x \succsim y$. □

Lemma B.11. *Suppose $\succsim^*$ admits a multi-prior expected utility representation $(u, \mathcal{M})$. For each $A \in \Sigma$ and all $f, g \in \mathcal{F}$,*

$$fAg \succsim^* g \Longleftrightarrow fAh \succsim^* gAh \text{ for all } h \in \mathcal{F}.$$

Proof. Let A be any event and let f, g, h be any acts in $\mathcal{F}$.

$$\begin{aligned} fAg \succsim^* g \Leftrightarrow & \sum_{s\in A} u(f(s))\mu(s) + \sum_{s\in S\setminus A} u(g(s))\mu(s) \\ & \geq \sum_{s\in A} u(g(s))\mu(s) + \sum_{s\in S\setminus A} u(g(s))\mu(s) \text{ for all } \mu \in \mathcal{M} \\ \Leftrightarrow & \sum_{s\in A} u(f(s))\mu(s) \geq \sum_{s\in A} u(g(s))\mu(s) \text{ for all } \mu \in \mathcal{M} \\ \Leftrightarrow & \sum_{s\in A} u(f(s))\mu(s) + \sum_{s\in S\setminus A} u(h(s))\mu(s) \\ & \geq \sum_{s\in A} u(g(s))\mu(s) + \sum_{s\in S\setminus A} u(h(s))\mu(s) \text{ for all } \mu \in \mathcal{M} \\ \Leftrightarrow & fAh \succsim^* gAh \end{aligned}$$

This lemma only relies on A being non-null, and it in fact holds for more general state spaces. □

Lemma B.12. *Suppose $\succsim^*$ admits a multi-prior expected utility representation $(u, \mathcal{M})$. For each $A \in \Sigma^*_+$ and all $f, g \in \mathcal{F}$, say $f \trianglerighteq^*_A g$ if $fAh \succsim^* gAh$ for some h. Then $\trianglerighteq_A$ admits a multi-prior expected utility representation $(u, \mathcal{B}(\mathcal{M}, A))$.*

Proof. By Lemma 2, $\trianglerighteq^*_A$ does not depend on the choice of h. It then follows that $f \trianglerighteq^*_A g$ if and only if for every $\mu \in \mathcal{M}$

$$\begin{aligned} \sum_{s\in A} u(f(s))\mu(s) + \sum_{s\in S\setminus A} u(h(s))\mu(s) \geq \sum_{s\in A} u(g(s))\mu(s) + \sum_{s\in S\setminus A} u(h(s))\mu(s) \\ \iff \frac{1}{\mu(A)} \sum_{s\in A} u(f(s))\mu(s) \geq \frac{1}{\mu(A)} \sum_{s\in A} u(g(s))\mu(s) \\ \iff \sum_{s\in A} u(f(s))\pi(s) \geq \sum_{s\in A} u(g(s))\pi(s) \text{ for all } \pi \in \mathcal{B}(\mathcal{M}, A), \end{aligned}$$

where $\mathcal{B}(\mathcal{M}, A) = \{\pi \in \Delta(S) \mid \pi = \mathcal{B}(\mu, A) \text{ for some } \mu \in \mathcal{M}\}$. □

Note that when $\succsim^*$ is complete, then we have the case of subjective expected utility and $\mathcal{M}$ and $\mathcal{B}(\mathcal{M}, A)$ are singleton sets.

B.5.2 Proof of B.2

Proof. Necessity is clear so only sufficiency is shown. By Conditional SEU, there exists a (u_A, μ_A) for each $A \in \Sigma$ that represents $\succsim_A$. Further, by Dynamic Conservatism, preferences satisfy ordinal preference consistency (see B.10): $x \succsim y$ if and only if $x \succsim_A y$. Hence we may assume $u = u_A$ for all A. Further, as X is convex, it is without loss to suppose $[-1, 1] \subset u(X)$, as u.

For each $A \in \Sigma_+$, define the binary relation $\trianglerighteq_A$ on $\mathcal{F}$ by $f \trianglerighteq_A g$ if and only if $fAg \succsim g$. Then by B.12, $\trianglerighteq_A$ has an expected utility representation $(u, \mathcal{B}(\mu, A))$.

Next, define the set $D_A := \{\pi \in \Delta(S) \mid \delta\mu + (1-\delta)\mathcal{B}(\mu, A) \text{ for } \delta \in [0, 1]\}$. By Dynamic Conservatism, it follows that $\mu_A \in D_A$. Suppose not, then as D_A and $\{\mu_A\}$ are closed, convex sets, there exists a separating hyperplane $a \in \mathbb{R}^{|S|}$ so that $\mu_A \cdot a > \hat{\mu} \cdot a$ for all $\hat{\mu} \in D_A$. Let $\bar{z} = \max_{\hat{\mu} \in D_A} \hat{\mu} \cdot a$ and let $\bar{a} = a - \bar{z}(1, \dots, 1)$. Then

$$\mu_A \cdot \bar{a} > 0 \geq \hat{\mu} \cdot \bar{a} \text{ for all } \hat{\mu} \in D_A. \tag{B.4}$$

We may suppose that $\bar{a} \in [-1, 1]^{|S|}$, since we can always multiply both sides of (B.4) by $\epsilon > 0$. Further, there are acts f, g such that $u(g(s)) - u(f(s)) = \bar{a}(s)$ for every $s \in S$. Consequently,

$$\sum_{s \in S} \mu_A(s) u(g(s)) > \sum_{s \in S} \mu_A(s) u(f(s)) \tag{B.5}$$

and

$$\sum_{s \in S} \hat{\mu}(s)u(f(s)) \geq \sum_{s \in S} \hat{\mu}(s)u(g(s)) \text{ for all } \hat{\mu} \in D_A. \tag{B.6}$$

By construction, $\mu, \mathcal{B}(\mu, A) \in D_A$ and so by (B.6) it follows that $fAg \succsim g$, $f \succsim g$. However, by (B.5) $g \succ_A f$, which contradicts Dynamic Conservatism. Hence $\mu_A \in D_A$. As A was arbitrary, the preceding argument applies to any non-null A. It is standard to show that u is unique up to a positive, affine transformation and, since $u(x) > u(y)$ for some $x, y \in X$, that μ and μ_A are also unique. Given uniqueness of μ and μ_A, it is trivial that there is a unique $\delta(A)$ that satisfies Equation B.1 whenever $\mu(A) < 1$. When $\mu(A) = 1$ (i.e., A^c is null), $\mu = \mathcal{B}(\mu, A) = \mu_A$ and $\succsim = \succsim_A$. When A, A^c are both non-null, define $\delta : \Sigma_+ \to [0,1]$ as the unique solution to

$$\mu_A = \delta(A)\mu + (1 - \delta(A))\mathcal{B}(\mu, A).$$

When A^c is null, define $\delta(A)$ arbitrarily. □

B.5.3 Proof of B.3

Proof. Theorem 1 shows that for each A, there is a $\delta(A) \in [0,1]$ satisfying the representation. Suppose Weak Consequentialism holds. It is sufficient to show that for any pair of non-null events $A, B \in \Sigma$, such that both $\mu(A) < 1$ and $\mu(B) < 1$, $\delta(B) = \delta(A)$.

Case 1: ($\mu(A \cup B) < 1$). Fix any non-null C satisfying $C \cap (A \cup B) = \varnothing$ and choose x, y, z such that $xCy \sim_A z$. By Weak Consequentialism, it follows that $xCy \sim_B z$.

Since preferences admit a conservative SEU representation, it follows that

$$\mu_A(C)u(x) + (1 - \mu_A(C))u(y) = u(z), \tag{B.7}$$

$$\mu_B(C)u(x) + (1 - \mu_B(C))u(y) = u(z). \tag{B.8}$$

Since x, y are arbitrary, it is without loss to suppose that $u(x) > u(y)$. Then, combining (B.7) and (B.8), it is clear that $\mu_A(C) = \mu_B(C)$. Since $C \cap (A \cup B)$, it follows that $\mathcal{B}(\mu, A)(C) = 0 = \mathcal{B}(\mu, B)(C)$, and hence $\mu_A(C) = \delta(A)\mu(C)$ and $\mu_B(C) = \delta(B)\mu(C)$. Hence equality is true if and only if $\delta(A) = \delta(B)$.

Case 2: ($\mu(A \cup B) = 1$). Suppose $A \cap B \neq \varnothing$. Then notice that $A, A \cap B$ ($B, A \cap B$) satisfy the conditions of Case 1. It follows then that $\delta(A) = \delta(A \cap B) = \delta(B)$. Now suppose $A \cap B = \varnothing$. Since there are at least three non-null events, without loss there exists some non-null set A' such that $A' \subset A$ and $\mu(A') < \mu(A)$.[3] Then we may again apply Case 1 to $A, A', A \setminus A'$, showing that $\delta(A) = \delta(A') = \delta(A \setminus A')$. Further, since $A \cap B = \varnothing$, it follows that $\mu(A' \cup B) < 1$, and so $\delta(A') = \delta(B)$. □

B.5.4 Proof of B.5

Proof. Let $A \geq_l B$ and fix any non-null C satisfying $C \cap (A \cup B) = \varnothing$. Choose any x, y, z such that $x \succ y$ and suppose $xCy \sim_B z$. It follows that

$$\mu_B(C)u(x) + (1 - \mu_B(C))u(y) \geq \mu_A(C)u(x) + (1 - \mu_A(C))u(y). \tag{B.9}$$

[3] Alternatively, there exists A' such that $A \subset A'$ and $\mu(A) < \mu(A')$, but this is just a relabeling. Note that such a pair of nested events must exist for at least one of A or B.

Since $u(x) - u(y) > 0$, simplifying the above yields $\mu_B(C) \geq \mu_A(C)$. The result then follows simply from the fact that $\mu_B(C) = \delta(B)\mu(C)$ and $\mu_A(C) = \delta(A)\mu(C)$. The reverse direction is routine. □

B.5.5 Proof of B.7

Proof. Suppose $\{\succsim_A^i\}_{A\in\Sigma}$ admit representations $(u_i, \mu_i, \delta_i)_{i=1,2}$ where $u_1 \approx u_2$. Then without loss $u_1 = u_2 = u$. Suppose agent 1 is more conservative than agent 2. Note that A and A^c are $\succsim_i$-non-null if and only if $\mu_i(A) \in (0,1)$. Pick some x, y_1, y_2 satisfying the conditions of B.6; then $u(x)\mu_1(A) + u(y_1)(1 - \mu_i(A)) = u(x)\mu_2(A) + u(y_2)(1 - \mu_2(A))$.

Case 1: ($\mu_1(A) = \mu_2(A)$). It is without loss to suppose $y_1 = y_2 = y$ for some y and that $u(x) = 0$. Further, we may ignore the dependence of μ on i. If $\{\succsim_A^1\}_{A\in\Sigma}$ is more conservative than $\{\succsim_A^2\}_{A\in\Sigma}$, it follows that $\delta_1(A)(1-\mu(A))u(y) \leq \delta_2(A)(1-\mu(A))u(y)$, or $\delta_1(A) \geq \delta_2(A)$. The reverse direction is similar.

Case 2: ($\mu_1(A) \neq \mu_2(A)$). Note that

$$\mu_1(A) - \mu_2(A) = (1 - \mu_2(A)) - (1 - \mu_1(A)) \neq 0. \tag{B.10}$$

By hypothesis, $u(x)\mu_1(A) + u(y_1)(1 - \mu_i(A)) = u(x)\mu_2(A) + u(y_2)(1 - \mu_2(A))$ from which, when combined with (B.10), it follows that

$$(1 - \mu_1(A))[u(y_1) - u(x)] = (1 - \mu_2(A))[u(y_2) - u(x)]. \tag{B.11}$$

From $\{\succsim_A^1\}_{A\in\Sigma}$ is more conservative than $\{\succsim_A^2\}_{A\in\Sigma}$, we conclude that

$$V_A^1(xAy_1) = \delta_1(A)[\mu_1(A)u(x) + (1-\mu_1(A))u(y_1)] + (1-\delta_1(A))u(x)$$

$$\leq \delta_2(A)[\mu_2(A)u(x) + (1-\mu_2(A))u(y_2)] + (1-\delta_2(A))u(x) = V_A^2(xAy_1).$$

Simplifying the above expression yields

$$\delta_1(A)(1-\mu_1(A))u(y_1) - \delta_1(A)(1-\mu_1(A))u(x)$$

$$\leq \delta_2(A)(1-\mu_2(A))u(y_2) - \delta_2(A)(1-\mu_2(A))u(x),$$

which directly implies

$$\delta_1(A)(1-\mu_1(A))[u(y_1)-u(x)] \leq \delta_2(A)(1-\mu_2(A))[u(y_2)-u(x)]. \tag{B.12}$$

The result that $\delta_1(A) \geq \delta_2(A)$ then follows by combining (B.12) with (B.11) and the facts that $u(y_i) - u(x) < 0$ and $0 < \mu_i(A) < 1$ for $i = 1, 2$.

□

B.5.6 Proof of B.9

Proof. The proof of this theorem is similar to the proof of B.2. First, for each $A \in \Sigma$, there exists a pair $(u_A, \mathcal{M}_A)$ such that $\succsim_A^*$ is represented by Equation B.2. For any $A \in \Sigma_+^*$, define the binary relation $\trianglerighteq_A^*$ on $\mathcal{F}$ by $f \trianglerighteq_A^* g$ if and only if $fAg \succsim^* g$. Then by B.12, $\trianglerighteq_A^*$ has a multi-prior expected utility representation $(u, \mathcal{B}(\mathcal{M}, A))$.

As in the proof of B.2, since $\succsim_A^*$ is complete on constant acts for every A it follows from B.10 that for any $A \in \Sigma_+^*$, $x \succsim y$ if and only if $x \trianglerighteq_A^* y$ if and only if $x \succsim_A y$. Hence it is without loss to suppose that $u = u_A$.

Next, let $D_A := conv(\mathcal{M}, \mathcal{B}(\mathcal{M}, A))$. Since $\mathcal{M}$ is a closed subset of $\Delta(S)$, it is compact. Further since A is unambiguously non-null, $\mathcal{B}(\mathcal{M}, A)$ is closed and hence also compact. Further, they are both convex. Then D_A is also compact and convex.

Now, suppose for contradiction that $\mathcal{M}_A \subseteq D_A$ is false. Then there exists some $\tilde{\mu}_A \in \mathcal{M}_A \setminus D_A$. Following an argument similar to a B.2, there exists a separating hyperplane $a \in \mathbb{R}^{|S|}$ so that $\tilde{\mu}_A \cdot a > \hat{\mu} \cdot a$ for all $\hat{\mu} \in D_A$. Let $\bar{z} = \max_{\hat{\mu} \in D_A} \hat{\mu} \cdot a$ and let $\bar{a} = a - \bar{z}(1, \ldots, 1)$. Then

$$\tilde{\mu}_A \cdot \bar{a} > 0 \geq \hat{\mu} \cdot \bar{a} \text{ for all } \hat{\mu} \in D_A. \tag{B.13}$$

We may suppose that $\bar{a} \in [-1, 1]^{|S|}$, since we can always multiply both sides of (B.13) by $\epsilon > 0$. Further, there are acts f, g such that $u(g(s)) - u(f(s)) = \bar{a}(s)$ for every $s \in S$. Consequently,

$$\sum_{s \in S} \tilde{\mu}_A(s) u(g(s)) > \sum_{s \in S} \tilde{\mu}_A(s) u(f(s)) \tag{B.14}$$

and

$$\sum_{s \in S} \hat{\mu}(s) u(f(s)) \geq \sum_{s \in S} \hat{\mu}(s) u(g(s)) \text{ for all } \hat{\mu} \in D_A. \tag{B.15}$$

By construction, $\mathcal{M}, \mathcal{B}(\mathcal{M}, A) \subset D_A$ and so by (B.15) it follows that $fAh \succsim^* gAh$ and $f \succsim^* g$. However, by (B.14) $f \not\succsim_A^* g$, which contradicts Unambiguous Dynamic Conservatism. Hence $\mathcal{M}_A \subseteq D_A$.

□

References

[1] Anderson, C. and A. D. Galinsky (2006): "Power, optimism, and risk-taking," *European Journal of Social Psychology*, 36, 511–536.

[2] Anscombe, F. J. and R. J. Aumann (1963): "A Definition of Subjective Probability," *The Annals of Mathematical Statistics*, 34, 199–205.

[3] Anscombe, F. J., R. J. Aumann, et al. (1963): "A definition of subjective probability," *Annals of mathematical statistics*, 34, 199–205.

[4] Aydogan, I., A. Baillon, E. Kemel, and C. Li (2017): "Signal perception and belief updating," Working paper.

[5] Camerer, C. (1995): "Individual Decision Making," in *Handbook of Experimental Economics*, ed. by A. Roth and J. Kagel, Princeton University Press.

[6] Charness, G. and C. Dave (2017): "Confirmation bias with motivated beliefs," *Games and Economic Behavior*, 104, 1–23.

[7] Cohen, B. L. and T. S. Wallsten (1992): "The effect of constant outcome value on judgments and decision making given linguistic probabilities," *Journal of Behavioral Decision Making*, 5, 53–72.

[8] Cohen, C. and A. Sela (2005): "Manipulations in contests," *Economics Letters*, 86, 135–139.

[9] Cox, J. C., V. L. Smith, and J. M. Walker (1988): "Theory and individual

behavior of first-price auctions," *Journal of Risk and Uncertainty*, 1, 61–99.

[10] de Clippel, G. and X. Zhang (2020): "Non-Bayesian Persuasion," *mimeo*.

[11] Dillenberger, D., J. S. Lleras, P. Sadowski, and N. Takeoka (2014): "A theory of subjective learning," *Journal of Economic Theory*, 153, 287–312.

[12] Dominiak, A., P. Duersch, and J.-P. Lefort (2012): "A Dynamic Ellsberg Urn Experiment," *Games and Economic Behavior*.

[13] Easley, D. and A. Ghosh (2015): "Behavioral mechanism design: Optimal crowdsourcing contracts and prospect theory," in *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 679–696.

[14] Edwards, W. (1982): "Conservatism in human information processing," in *Judgment under Uncertainty: Heuristics and Biases*, ed. by D. Kahneman, P. Slovic, and A. Tversky, Cambridge University Press, 359–369.

[15] Eil, D. and J. M. Rao (2011): "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself," *American Economic Journal: Microeconomics*, 3, 114–138.

[16] El-Gamal, M. A. and D. M. Grether (1995): "Are People Bayesian? Uncovering Behavioral Strategies," *Journal of the American Statistical Association*, 90, 1137–1145.

[17] Engelbrecht-Wiggans, R. (1989): "The effect of regret on optimal bidding in auctions," *Management Science*, 35, 685–692.

[18] Engelmann, J. B., M. Lebreton, N. A. Salem-Garcia, P. Schwardmann, and J. J. van der Weele (2024): "Anticipatory anxiety and wishful thinking," *American Economic Review*, 114, 926–960.

[19] Epley, N. and T. Gilovich (2016): "The mechanics of motivated reasoning," *Journal of Economic perspectives*, 30, 133–40.

[20] Epstein, L. G. (2006): "An Axiomatic Model of Non-Bayesian Updating,"

Review of Economic Studies.

[21] Epstein, L. G., J. Noor, and A. Sandroni (2008): "Non-Bayesian Updating: A Theoretical Framework," *Theoretical Economics*.

[22] Gagnon-Bartsch, T., M. Pagnozzi, and A. Rosato (2021): "Projection of private values in auctions," *American Economic Review*, 111, 3256–3298.

[23] Ghirardato, P. (2002): "Revisiting Savage in a Conditional World," *Economic Theory*.

[24] Gilboa, I. and D. Schmeidler (1989): "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics*, 18, 141–153.

[25] Granberg, D. and E. Brent (1983): "When prophecy bends: The preference–expectation link in US presidential elections, 1952–1980." *Journal of Personality and Social Psychology*, 45, 477.

[26] Gul, F. and W. Pesendorfer (2001): "Temptation and Self-Control," *Econometrica*.

[27] Hogarth, R. M. (1987): *Judgement and choice: The psychology of decision*, Chichester: John Wiley & Sons.

[28] Kagel, J. H. and D. Levin (2008): "Auctions: A survey of experimental research, 1995-2008," *Handbook of Experimental Economics*, 2.

[29] Kahneman, D. and A. Tversky (1972): "Subjective Probability: A Judgment of Representativeness," *Cognitive Psychology*.

[30] ——— (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263–291.

[31] Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

[32] Kőszegi, B. and M. Rabin (2006): "A model of reference-dependent preferences," *The Quarterly Journal of Economics*, 121, 1133–1165.

[33] KOVACH, M. (2020): "Twisting the truth: Foundations of wishful thinking," *Theoretical Economics*, 15, 989–1022.

[34] ——— (2021): "Conservative Updating," *working paper* .

[35] KRIZAN, Z. AND P. D. WINDSCHITL (2007): "The influence of outcome desirability on optimism." *Psychological Bulletin*, 133, 95.

[36] KUNDA, Z. (1990): "The case for motivated reasoning." *Psychological bulletin*, 108, 480.

[37] LEE, Y.-J., W. LIM, AND C. ZHAO (2019): "Cheap Talk with Non-Bayesian Updating," *working paper.*

[38] MATROS, A. (2012): "Sad-loser contests," *Journal of Mathematical Economics*, 48, 155–162.

[39] MAYRAZ, G. (2011): "Wishful thinking," *Available at SSRN 1955644.*

[40] MENGEL, F. AND P. CAMPOS MERCADE (2021): "Non Bayesian Statistical Discrimination," *Available at SSRN 3843579.*

[41] MILGROM, P. AND I. SEGAL (2002): "Envelope theorems for arbitrary choice sets," *Econometrica*, 70, 583–601.

[42] MINCHUK, Y. (2018): "Effect of reimbursement on all-pay auction," *Economics Letters*, 172, 28–30.

[43] MÖBIUS, M. M., M. NIEDERLE, P. NIEHAUS, AND T. S. ROSENBLAT (2022): "Managing self-confidence: Theory and experimental evidence," *Management Science.*

[44] MORGAN, J., K. STEIGLITZ, AND G. REIS (2003): "The spite motive and equilibrium behavior in auctions," *Contributions to Economic Analysis & Policy*, 2, 1102.

[45] MYERSON, R. B. (1981): "Optimal auction design," *Mathematics of operations research*, 6, 58–73.

[46] MÃ¼ller, A. and D. Stoyan (2002): *Comparison Methods for Stochastic Models and Risks*.

[47] Neugebauer, T. and R. Selten (2006): "Individual behavior of first-price auctions: The importance of information feedback in computerized experimental markets," *Games and Economic Behavior*, 54, 183–204.

[48] Nickerson, R. S. (1998): "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises," *Review of General Psychology*, 2, 175–220.

[49] Nikolova, E., E. Pountourakis, and G. Yang (2018): "Optimal mechanism design with risk-loving agents," in *International Conference on Web and Internet Economics*, Springer, 375–392.

[50] Ortoleva, P. (2012): "Modeling the Change of Paradigm: Non-Bayesian Reactions to Unexpected news," *The American Economic Review*.

[51] ——— (2012): "Modeling the change of paradigm: Non-Bayesian reactions to unexpected news," *American Economic Review*, 102, 2410–2436.

[52] Phillips, L. D. and W. Edwards (1966): "Conservatism in a simple probability inference task," *Journal of Experimental Psychology*, 72, 346–354.

[53] Rego, A., F. Sousa, C. Marques, and M. P. e. Cunha (2012): "Optimism predicting employees' creativity: The mediating role of positive affect and the positivity ratio," *European Journal of Work and Organizational Psychology*, 21, 244–270.

[54] Riley, J. G. and W. F. Samuelson (1981): "Optimal auctions," *The American Economic Review*, 71, 381–392.

[55] Rosato, A. (2023): "Loss aversion in sequential auctions," *Theoretical Economics*, 18, 561–596.

[56] Savage, L. J. (1954): *The Foundations of Statistics*, John Wiley and Sons.

[57] Shaked, M. and J. G. Shanthikumar (1994): *Stochastic orders and their*

applications, Academic press.

[58] Sigall, H., A. Kruglanski, and J. Fyock (2000): "Wishful thinking and procrastination," *Journal of Social Behavior and Personality*, 15, 283–296.

[59] Sinander, L. (2022): "The converse envelope theorem," *Econometrica*, 90, 2795–2819.

[60] Tversky, A. and D. Kahneman (1991): "Loss aversion in riskless choice: A reference-dependent model," *The Quarterly Journal of Economics*, 106, 1039–1061.

[61] ——— (1992): "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and Uncertainty*, 5, 297–323.