

Eliciting Informed Preferences

Modibo Camara

Nicole Immorlica

Brendan Lucier

February 9, 2025

Abstract

If people find it costly to evaluate the options available to them, their choices may not directly reveal their preferences. Yet, it is conceivable that a researcher can still learn about a population's preferences with careful experiment design. We formalize the researcher's problem in a model of robust mechanism design where it is costly for individuals to learn about how much they value a product. We characterize the statistics that the researcher can identify, and find that they are quite restricted. Finally, we apply our positive results to social choice and propose a way to combat uninformed voting.

Contents

1	Introduction	1
2	Model	6
2.1	Information Acquisition	7
2.2	Assumptions	9
2.3	Mechanism Design	12
2.4	Elicitable Statistics	13
3	Strong Elicitation	16
3.1	Intuition for Theorem 1	16
4	Weak Elicitation	18
4.1	Mechanism	19
4.2	Intuition for Theorem 2	21
5	Application to Social Choice	25
5.1	Model	26
5.2	Challenges	28
5.3	Statement of Theorem 3	29
5.4	Mechanism	30
5.5	Intuition for Theorem 3	32
6	Related Literature	33
7	Conclusion	38
A	Proof Outlines	39
A.1	Proof of Theorem 1	39
A.2	Special Cases of Theorem 2	42
A.3	Proof of Theorem 2	47
A.4	Proof of Theorem 3 (Special Case)	51
B	Supplemental Appendix	S1
B.1	Proof of Theorem 3 (General Case)	S1
B.2	Proof of Claim 1	S5

B.3	Proof of Lemma 1	S5
B.4	Proof of Claim 2	S8
B.5	Proof of Claim 3	S9
B.6	Proof of Lemmas 2, 6, and 8	S9
B.7	Proof of Corollary 2	S15
B.8	Proof of Lemmas 3 and 9	S16
B.9	Proof of Lemmas 4 and 10	S22
B.10	Proof of Lemma 5	S25
B.11	Proof of Lemma 7	S26
B.12	Proof of Claims 9 and 10	S28
B.13	Proof of Claim 11	S30
B.14	Proof of Claim 12	S31
B.15	Proof of Claim 13	S37
B.16	Proof of Claim 14	S39
B.17	Proof of Lemma 11	S41
B.18	Proof of Claim 15	S42

1 Introduction

People often find it costly to process information about the goods they consume, the services they receive, and the policies that affect them. As a result, their choices may not directly reveal the preferences they would have had if they had processed all the information available to them. To what extent can researchers still learn about a population’s preferences through careful experiment design?

Learning about a population’s preferences – i.e., their informed preferences – is important for a number of applications. For example, take market research. Suppose that OpenAI is about to launch a product called GPT-5 and wants to forecast demand. In the present, consumers do not know how much value they would get from GPT-5, and find it costly to test the product. In the future, consumers become better informed (as subscribers gain first-hand experience and others rely on second-hand experience). Can we design an experiment that helps OpenAI forecast future demand?

We formalize the researcher’s problem as robust mechanism design with information acquisition. There is a sample of agents, and they do not necessarily know their values from some product. The researcher wants to learn about these values, but does not know their distribution, including whether and how they are correlated across agents. Agents can learn about their own and potentially about others’ values by acquiring costly signals. We impose little structure on the information acquisition technology, but we always assume that agents can learn their own values at some finite cost.

We use this model to tightly characterize what researchers can learn about the population’s preferences. The researcher is able to design any experiment that he wishes, recruit as many agents as he wishes, and spend as much money as he wishes. He *strongly elicits* a given statistic (i.e., moment) of the population’s values if, for every equilibrium, the agents’ choices in the experiment identify that statistic. He *weakly elicits* that statistic if (i) there is an equilibrium that identifies the statistic and (ii) there is no equilibrium where the researcher’s estimate is worse than if the agents had not acquired any information at all.

It turns out that what the researcher can learn is quite limited. Theorem 1 says that a statistic is strongly elicitable if and only if it is constant. That is, the statistic must take on the same value in all cases. Theorem 2 says that a statistic is weakly elicitable if and only if it is affine. That is, the statistic must be equivalent to the population’s

average value (up to affine transformation). Without additional assumptions, many basic statistics – like quantity demanded by an informed population – are not revealed by choices, no matter how carefully-designed those choices are.

We illustrate the value of our positive results in an application to social choice. Specifically, consider the problem of uninformed voting. Empirical evidence suggests many voters are poorly informed (e.g., Delli Carpini and Keeter 1996, Angelucci and Prat 2024) and would vote differently if they were better informed (e.g., Bartels 1996, Fowler and Margolis 2014). Building on our previous results, we design a mechanism that incentivizes voters to become better informed. This mechanism obtains nearly first-best welfare in favorable equilibria, and provides meaningful guarantees even in unfavorable equilibria (Theorem 3).

We now discuss the model and results in more detail.

Model. We formalize the researcher’s problem as robust mechanism design with information acquisition.

The researcher asks a random sample of n agents to participate in a mechanism. There is a product and each agent i derives value v_i from that product. The distribution of values in the population may feature complex patterns of heterogeneity and correlation. The researcher does not know this distribution. He commits to a prior-independent mechanism (e.g., Devanur et al. 2011) where the agents send messages to the researcher and then receive individualized allocations and transfers.

Initially, agents do not know their own values. Instead, each agent i has some private information and may learn more – both about her own value v_i and perhaps about others’ values v_{-i} – by acquiring costly signals. We assume that each agent can learn her value v_i at some finite cost. Agents may also have access to additional, more complex signals that are unknown to the researcher

The researcher wants to learn about the distribution of values in the population. More precisely, he wants to learn about demand for the product in the counterfactual where every agent i knew her value v_i . We focus on statistics that can be represented as moments (e.g., average willingness to pay, or quantity demanded at a fixed price). To learn about a statistic, the researcher can recruit as many agents as he likes, spend as much money as he likes, and allocate as many products as he likes.

A statistic is *strongly elicitable* if there exists a mechanism where the observed choices identify that statistic in every equilibrium. More precisely, there exists a

sequence of mechanisms, indexed by the sample size n , where messages sent in equilibrium pin down the statistic's value in the limit as $n \rightarrow \infty$. This must hold for any value distribution and every Bayes-Nash equilibrium, and requires exact (not just partial) identification.

A statistic is *weakly elicitable* if there exists a mechanism where (i) the observed choices identify that statistic in the researcher's preferred equilibrium and (ii) the researcher's estimate is not unreasonably bad in any equilibrium. By "not unreasonably bad", we mean that the researcher's estimate should not be worse than in a counterfactual where agents did not acquire any information at all.

Next, we completely characterize the set of elicitable statistics.

Strong Elicitation. It turns out that any non-trivial statistic is not strongly elicitable. More precisely, Theorem 1 says that a strongly elicitable statistic must take on the same value regardless of the distribution.

To overturn this negative result, the researcher would have to impose stronger assumptions or give up on exact identification, neither of which is needed in classical models of preference elicitation (e.g., Becker et al. 1964; Clarke 1971). However, we find that the result is robust to imposing certain assumptions. It remains true even if we restrict attention to *basic instances* where information acquisition is binary and a fraction of the population is intrinsically well-informed.

Weak Elicitation. Since strong elicitation is generally impossible, we turn to weak elicitation. Our results in this case are more positive, but only slightly.

Theorem 2 says that a statistic is weakly elicitable if and only if it is equivalent to the population's average value (up to affine transformation). This implies that many natural and economically-relevant statistics are not even weakly elicitable. In our motivating example, it means that OpenAI would not be able to forecast the quantity demanded at a given price.

To weakly elicit the average, the researcher presents agents with an incentivized survey and asks them, before completing the survey, to predict the survey results. We call this mechanism *BDM-with-betting*. It has two stages. In the second stage, agents report their willingness to pay in a standard Becker–DeGroot–Marschak (BDM) mechanism. In the first stage, agents predict the average reported willingness to pay. Prediction accuracy is rewarded according to a proper scoring rule, and the researcher

uses the average reported value to estimate the population’s average value.

The BDM-with-betting mechanism is relatively simple – even naive – and this begs the question: why does it work? It is simpler than than other mechanisms with comparable features (e.g., Crémer and McLean 1985; Miller et al. 2005). It is naive because it incentivizes agent i to learn about others’ willingness to pay, not her own value v_i . Nonetheless, BDM-with-betting works because it incentivizes agents to acquire information up until the point where their reporting errors are uncorrelated. If these errors – the difference between an agent’s value and her reported value – are uncorrelated, they vanish in the aggregate. Essentially, betting restores the “wisdom of the crowd”, which usually requires strong distributional assumptions.

All in all, our results suggest that what the researcher can learn is quite limited. However, there is a setting where learning the population’s average value is all that is needed. That setting is social choice.

Application to Social Choice. We turn to the problem of social choice, where a planner chooses one of two alternatives on behalf of a population. He forms a committee consisting of n agents sampled from the population (e.g., a citizen’s assembly). As before, we consider an optimistic case where the planner can spend as much money and recruit as many agents as he likes.

We are motivated by the problem of uninformed voting. Existing electoral systems do not give voters much of a reason to become informed about the policies and candidates that appear on their ballot. This is not just a theoretical concern. Empirically, many voters are poorly informed (e.g., Delli Carpini and Keeter 1996, Angelucci and Prat 2024), and making them better-informed could plausibly affect vote margins (e.g., Lau and Redlawsk 1997, Fowler and Margolis 2014).

It turns out that the methods we developed to elicit preferences can be used to design social choice mechanisms that are robust to costly information processing. Theorem 3 finds mechanisms that are (i) efficient in favorable equilibria and, for any equilibria, (ii) do no worse in than simply choosing the ex-ante optimal alternative. By comparison, standard voting procedures do not even guarantee (ii) in favorable equilibria (see e.g., Ali et al. 2025).

Under a strong simplifying assumption, we can motivate a mechanism – *majority-rule-with-betting* – that is particularly simple. There are two stages. In the second stage, agents report their preferred alternative. In the first stage, they predict

the vote margin. The planner chooses the alternative with majority support.¹

Transfers are only used to reward accurate predictions; there is no vote buying.

There are two key reasons why this mechanism works. The first reason is the same as before: betting restores the “wisdom of the crowd”, assuming truthful voting. The second reason is new. It is well-known that uninformed voters may not vote truthfully if they believe that how others vote may be informative about their own preferences (e.g., Austen-Smith and Banks 1996). It turns out that betting restores incentives for truth-telling. Intuitively, an uninformed agent i that believes her value v_i is correlated with the vote margin is leaving money on the table. If the betting stakes are large, she prefers to learn her value v_i in order to better predict the vote margin.

Related Literature. We build on several research areas. We briefly highlight our contributions now, and leave a fuller discussion to Section 6.

First, we contribute to research on preference elicitation (e.g., Becker et al. 1964; Clarke 1971) by relaxing the assumption that individuals know their own preferences. We identify sharp limits on our ability to learn about preferences through observed choices. Moreover, BDM-with-betting follows a line of work on eliciting private information that (implicitly) asks agents to bet on each others’ reports. Most existing work uses betting to incentivize information revelation (e.g., Crémer and McLean 1985; Prelec 2004; Miller et al. 2005; Pakzad-Hurson 2022). BDM-with-betting, in contrast, only uses betting to incentivize information acquisition.

Second, we contribute to research on mechanism design with information acquisition (e.g., Persico 2000; Bergemann and Välimäki 2002). We propose a model of unstructured information acquisition that builds on and complements prior work (e.g., Carroll 2015; Carroll 2019; Denti and Ravid 2024). We find mechanisms that are robust to the details of how agents acquire information.

Third, we contribute to the theory of voting, and social choice more broadly. Historically, this literature has focused on mitigating the negative effects of strategic voting (e.g., Arrow 1950; Groves and Ledyard 1977). We contribute to more recent research on mitigating the negative effects of uninformed voting (e.g., Persico 2004; Feddersen and Sandroni 2006; Gerardi and Yariv 2008; Gershkov and Szentes 2009), largely by relaxing restrictive assumptions (e.g., common values). We also propose a mechanism that alleviates the swing voter’s curse (e.g., Feddersen and Pesendorfer

¹More precisely, the planner chooses this alternative with high probability. See Section 5.3 for details.

1996; Jaehoon Kim and Fey 2007) when information is endogenous.

Organization. The rest of this paper is organized as follows. Section 2 presents the model. Section 3 characterizes strongly elicitable statistics. Section 4 characterizes weakly elicitable statistics. Section 5 applies our results to social choice. Section 6 discusses the related literature. Section 7 concludes. Appendix A outlines the proofs of our main results, and the Supplemental Appendix fills the gaps in the outlines.

2 Model

We develop a model of robust mechanism design with information acquisition. This model accommodates the wide range of experiments that researchers might run.

In the model, there is a researcher and n agents. Agent i receives an *alternative*

$$x_i \in \mathcal{X} = \{0, 1\}$$

Given alternative x_i , agent i receives *value*

$$v_{ix_i} \in \mathcal{V} \subseteq [-v_L, v_H] \subseteq \mathbb{R}$$

where we normalize $v_{i0} = 0$ and let v_i refer to v_{i1} . We assume that \mathcal{V} is finite. We often refer to alternative $x = 1$ as a product and $x = 0$ as an outside option.

As a running example, suppose the firm OpenAI develops a new product called GPT-5. The outside option is an existing product called GPT-4, and each consumer i 's value v_i captures how much she would benefit from upgrading GPT-4 to GPT-5. Before GPT-5 is launched, the consumer does not know her value v_i , because she does not know enough about the new product's functionalities, use cases, flaws, etc.

Naturally, an uninformed consumer's willingness to pay for GPT-5 may differ from her value v_i . Although OpenAI ultimately cares about willingness to pay, it is likely to converge to the value v_i in the long-run. Consumers that do not value a product are unlikely to buy it again or renew their subscription. Even a first-time customer is likely to become better-informed in the years after a product launch, as early adopters and professional critics share their experiences.

Next, we turn to agents' information about their own and each others'

preferences, and the process by which they acquire more information.

2.1 Information Acquisition

We formalize the agents' information and information acquisition in an unstructured way. In particular, we aim to accommodate the complicated patterns of heterogeneity and correlation that are likely to arise in practice.

We begin by briefly summarizing the model. Initially, agent i does not know much about her own value v_i . She understands that it is drawn from a distribution that belongs to family of distributions parameterized by a *hidden state* ω , and she shares a common prior over the state ω with the other agents $j \neq i$. She can learn more by making use of an *information acquisition technology* θ_i , which determines the cost of different *signals* s_i . In turn, signals s_i may tell her something about her own value v_i , or about the distribution $F(\omega)$ of values in the population.

Hidden States. Let Ω be a finite set of hidden states ω . Neither the researcher nor the agents know the state initially. For example, the hidden state might indicate how useful GPT-5 is for a given task.

Technologies. In order to learn more about themselves and about the population, each agent i can use her information acquisition technology θ_i to acquire a signal s_i . Let Θ be a finite set of technologies and let \mathcal{S} be a finite set of signals. A *cost function* c determines agent i 's cost of acquiring signal s_i if her technology is θ_i , i.e.,

$$c : \mathcal{S} \times \Theta \rightarrow \mathbb{R}_+ \cup \{\infty\}$$

For example, a technology could indicate how long it takes a consumer to set up a coding task for GPT-5. Some consumers may be faster than others.

Let agent i 's *type* be her technology-value pair (θ_i, v_i) . We do not assume that agent i knows anything about her type up front. For example, a consumer may not know how long it will take her to set up a coding task for GPT-5. This modeling choice may seem to rule out many natural cases, but Remark 2 explains why it does not.

Joint Distribution. The *joint distribution* F determines how states, technologies, and values are distributed. Formally, let

$$F \in \Delta(\Omega \times \Theta \times \mathcal{V})$$

where F has a finite support. Let the *realized distribution* $F(\omega)$ be the marginal distribution over $\Theta \times \mathcal{V}$ after conditioning on the state ω . We assume that the agents know the joint distribution F , but not necessarily the realized distribution $F(\omega)$.

There are two steps to the sampling process. First, the state ω is drawn according to the marginal distribution F_Ω . This determines the realized distribution $F(\omega)$, which describes the frequency of values and technologies in the population.

Second, the researcher randomly draws n agents from the population. Formally, the types (θ_i, v_i) are drawn independently across agents $i = 1, \dots, n$ according to the realized distribution $F(\omega)$. To be clear, this means that for any two agents $i \neq j$, (θ_i, v_i) is independent of (θ_j, v_j) conditional on the state ω . It does not mean that v_i is independent of θ_i conditional on the state ω . It also does not mean that (θ_i, v_i) is unconditionally independent of (θ_j, v_j) .

Signals. A signal s_i for agent i is a function of the state ω , her technology θ_i , and her value v_i . In principle, signals could output an arbitrary message, but it is convenient to let the output be a finite sequence of real numbers. Formally,

$$s_i : \Omega \times \Theta \times \mathcal{V} \rightarrow \mathbb{R}^*$$

is a random variable that maps samples $(\omega, \theta_i, v_i) \sim F$ to *signal realizations* $s_i(\omega, \theta_i, v_i)$.

After observing the signal realization, agent i may update her beliefs about the state ω , her technology θ_i , and her value v_i . If she learns about the state ω , she may also update her beliefs about the realized distribution $F(\omega)$ and thereby learn about the technologies θ_{-i} and values v_{-i} of other agents.

For example, a consumer could learn about her value from GPT-5 by using it. They might think of different tasks (e.g., writing, coding, translation, search, etc.) and evaluate how well it performs in each case. Each task corresponds to a different signal, and GPT-5's performance at the task corresponds to a signal realization. Learning about GPT-5's performance could cause consumer to learn directly about her own value v_i , and indirectly about other consumers' values v_{-i} .

Remark 1. Technically, each agent i acquires exactly one signal. However, this model is rich enough to accommodate settings where agents acquire multiple signals. For any two signals $s_i, s'_i \in \mathcal{S}$, we can enrich the signal space by adding a combined signal

$$s''_i(\omega, v_i, \theta_i) = (s_i(\omega, \theta_i, v_i), s'_i(\omega, v_i, \theta_i)) \quad (1)$$

Similarly, this model can accommodate dynamic information acquisition. Suppose that agent i wants to acquire signal s'_i if and only if the realization of signal s_i is in a set $U \subseteq s_i(\Omega, \mathcal{V}, \Theta)$. We can enrich the signal space by adding a dynamic signal

$$s''_i(\omega, v_i, \theta_i) = \begin{cases} (s_i(\omega, \theta_i, v_i), s'_i(\omega, v_i, \theta_i)) & s_i(\omega, \theta_i, v_i) \in U \\ s_i(\omega, \theta_i, v_i) & \text{otherwise} \end{cases} \quad (2)$$

We can do so for every possible combination of (s_i, s'_i, U) without violating the assumption that the signal space \mathcal{S} is finite. Along these lines, we can also accommodate dynamic signals that combine more than two signals.

Remark 2. Initially, agent i does not know anything about her type (v_i, θ_i) . However, it is easy to accommodate initial information as follows. First, define a signal s_i with cost $c(s_i, \theta_i) = 0$ for all types θ_i . This signal realization represents the agent's initial information. Second, expand the signal space \mathcal{S} by allowing all dynamic signals that combine s_i with another signal $s'_i \in \mathcal{S}$, following the construction (2).

Model Instances. A *model instance* I combines the model primitives, i.e.,

$$I = (\Omega, \mathcal{V}, \Theta, \mathcal{S}, F, c)$$

We assume that the agents know the instance I , but the researcher does not. Let \mathcal{J} be the collection of all instances that the researcher considers plausible.

2.2 Assumptions

We maintain three assumptions. The first assumption says that the collection \mathcal{J} is sufficiently rich. It refers to highly-structured instances that we call *basic*.

Definition 1. A basic instance I consists of a binary state space Ω , a ternary value space \mathcal{V} , a binary technology space Θ , a binary signal space \mathcal{S} , a joint distribution F , and a cost function c . There is a good technology t_i^G and a bad technology t_i^B , where the good technology is in the support of the realized distribution $F(\omega)$ for any state $\omega \in \Omega$. There is also a good signal s_i^G and a bad signal s_i^B . The good signal reveals everything, i.e.,

$$s_i^G(\omega, \theta_i, v_i) = (\omega, \theta_i, v_i)$$

The bad signal reveals everything if agent i has the good technology. If agent i has the bad technology, it only reveals the technology. That is,

$$s_i^B(\omega, \theta_i, v_i) = \begin{cases} (\omega, \theta_i, v_i) & \theta_i = \theta_i^G \\ \theta_i & \theta_i = \theta_i^B \end{cases}$$

Finally, only the good signal under the bad technology can have a non-zero cost. That is, the cost function satisfies $c(s_i^G, \theta_i^G) = c(s_i^B, \theta_i^G) = c(s_i^B, \theta_i^B) = 0$.

In a basic instance, agents face a simple information acquisition problem. If they have the good technology, it does not matter which signal they acquire: they learn their value at zero cost either way. For that reason, agents might as well assume that they have the bad technology. In that case, they can either (i) learn their value at a cost or (ii) not learn their value and avoid that cost.

One notable feature of basic instances is that a positive fraction of the population will always be well-informed. This follows from the fact that the good technology is always in the support of the realized distribution. We add this feature to emphasize that our negative results hold even when some agents are intrinsically well-informed.

Assumption 1. The collection \mathcal{I} includes (but is not limited to) all basic instances.

We stress that Assumption 1 does not restrict attention to basic instances, and that our positive results must hold for any collection of instances (basic or not).

There are two reasons why we draw attention to basic instances. First, it makes it easier to judge whether the kinds of instances that drive our negative results are likely to arise in a given application of interest. Second, it highlights that our negative results are not driven by cases that are obviously pathological.

The second assumption says that agents can combine signals. That is, acquiring one signal does not prevent an agent from acquiring another one.

Assumption 2. *Restrict attention to instances I with the following structure. The set of signals \mathcal{S} consists of base signals and combined signals. For every set of base signals $\{s_i^1, \dots, s_i^\ell\} \subseteq \mathcal{S}$, there is a combined signal $s_i \in \mathcal{S}$ where*

$$\forall (\omega, v_i, \theta_i) \in \Omega \times \mathcal{V} \times \Theta : s_i(\omega, \theta_i, v_i) = (s_i^1(\omega, \theta_i, v_i), \dots, s_i^\ell(\omega, \theta_i, v_i)) \quad (3)$$

If these base signals have a finite cost, then the combined signal also has a finite cost, i.e.,

$$\forall \theta_i \in \Theta : \sum_{j=1}^{\ell} c(s_i^j, \theta_i) < \infty \implies c(s_i, \theta_i) < \infty \quad (4)$$

For example, Assumption 2 says that evaluating GPT-5 on a coding task does not prevent a consumer from also evaluating GPT-5 on a translation task. This may be violated in practice if, say, the consumer only has time to prepare one task.

The third assumption says that agent i can learn her own value v_i at a finite cost. It refers to a particular kind of signal that we call a *revealing signal*.

Definition 2. *A signal s_i is revealing if agent i learns her value v_i after acquiring it, i.e.,*

$$\text{Var}[v_i \mid s_i(\omega, \theta, v_i)] = 0 \quad (5)$$

A revealing signal for agent i does not necessarily reveal everything. It does not necessarily reveal the hidden state ω , the values v_{-i} of other agents, or the technology profile $\theta = (\theta_1, \dots, \theta_n)$. An agent with access to a revealing signal can precisely learn about her own preferences, but not necessarily more than that.

Assumption 3. *Restrict attention to instances I that feature a revealing signal $s_i \in \mathcal{S}$ that has finite cost. That is, for all technologies θ_i , $c(s_i, \theta_i) < \infty$.*

For example, a consumer that dedicates an entire week to testing GPT-5 may get a pretty good sense of her value from that product, even if she does not necessarily learn how much value other consumers will derive from that product. Even if this is quite costly to the consumer, Assumption 3 says that it is feasible. Of course, this assumption is at best an approximation of reality. Regardless of how much effort they put in, real consumers will never learn their values with perfect precision.

Assumption 3 follows work in mechanism design that assumes there is a known action available to agents but seeks robustness to other unknown actions (e.g., Carroll 2015; Carroll 2019). In this case, the revealing signal is the known action. Of course, the revealing signal may be quite costly. Agents may prefer to acquire other signals that the researcher may not be aware of (e.g., signals that are cheaper but provide less information about v_i , or signals that only provide information about the state ω).

2.3 Mechanism Design

To learn about preferences, the researcher can commit to a *mechanism*. As usual, agents send messages to the researcher and the researcher uses those messages to determine the alternative x_i and transfers t_i for each agent i .

Definition 3. Fix a set of message profiles $M = (M_i)_{i=1}^n$. A mechanism (\mathbf{x}, \mathbf{t}) consists of:

1. Allocation rules that maps message profiles to distributions over alternatives x , i.e.,

$$\mathbf{x}_i : M \rightarrow \Delta(\mathcal{X})$$

2. Transfer rules that map message profiles to distributions over transfers, i.e.,

$$\mathbf{t}_i : M \rightarrow \Delta(\mathbb{R})$$

For convenience, let $\mathbf{x}_i(m)$ refer to both the distribution over alternatives as well as the realized alternative. Do the same for transfers $\mathbf{t}_i(m)$.

After the mechanism is announced, the game proceeds in three steps. First, each agent i chooses a signal s_i to acquire. Second, each agent i observes her signal realization $s_i(\omega, \theta_i, v_i)$. Third, each agent i sends a message m_i . Agents' strategies specify what signals they acquire and what messages they send to the researcher.

Definition 4. Agent i 's strategy $(\mathbf{s}_i, \mathbf{m}_i)$ pairs a signal rule $\mathbf{s}_i \in \Delta(\mathcal{S})$ with a message rule \mathbf{m}_i that maps the signal and signal realization to a distribution over messages, i.e.,

$$\mathbf{m}_i : \mathcal{S} \times \mathbb{R}^* \rightarrow \Delta(M_i)$$

For convenience, let \mathbf{s}_i refer to both the distribution over signals chosen as well as the realized signal s_i . Do the same for messages $\mathbf{m}(s_i, s_i(\omega, \theta_i, v_i))$.

Agent i 's utility depends on her values v_i , her alternative x_i , the transfers t_i she receives, and the cost $c(s_i, \theta_i)$ of her chosen signal s_i given her technology θ_i . That is,

$$u_i(v_i, x_i, t_i, s_i, \theta_i) = v_{ix_i} + t_i - c(s_i, \theta_i)$$

In turn, agent i 's expected utility from a strategy profile (\mathbf{s}, \mathbf{m}) is

$$U_i(\mathbf{s}, \mathbf{m}, \mathbf{x}, \mathbf{t}) = E[u_i(v_i, \mathbf{x}_i(m), \mathbf{t}_i(m), s_i, \theta_i)] \quad \text{where} \quad m_i = \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i))$$

The expectation is taken with respect to the joint distribution F , any randomness in the strategy profile (\mathbf{s}, \mathbf{m}) , and any randomness in the mechanism (\mathbf{x}, \mathbf{t}) .

With this notation in hand, we can specify our solution concept.

Definition 5. A strategy profile (\mathbf{s}, \mathbf{m}) is a Bayes-Nash equilibrium of mechanism (\mathbf{x}, \mathbf{t}) if every agent i prefers her strategy (s_i, \mathbf{m}_i) to every alternative strategy (s'_i, \mathbf{m}'_i) . That is,

$$U_i(\mathbf{s}, \mathbf{m}, \mathbf{x}, \mathbf{t}) \geq U_i((s'_i, \mathbf{s}_{-i}), (\mathbf{m}'_i, \mathbf{m}_{-i}), \mathbf{x}, \mathbf{t})$$

Having specified a solution concept, we can formalize the researcher's problem.

2.4 Elicitable Statistics

The researcher wants to learn about the population's values over alternatives.

Specifically, the researcher wants to learn about $F_{\mathcal{V}}(\omega)$, the marginal distribution of values conditional on the realized state ω . This is the distribution he cares about because it describes realized demand by an informed population. For example, if the realized state ω indicates that GPT-5 is useful, the researcher may learn that demand is high. Typically, he does not care whether demand would also have been high in a counterfactual state ω' where GPT-5 was not useful.

We focus on learning *statistics* that are moments of the realized marginal distribution $F_{\mathcal{V}}(\omega)$. Two important objects that this leaves out are the values (v_1, \dots, v_n) of the sampled agents and the distribution $F_{\mathcal{V}}(\omega)$ itself. However, our negative results immediately imply that these objects cannot be elicited.

Definition 6. A σ -statistic is a moment of the realized marginal distribution $F_{\mathcal{V}}(\omega)$.

Given a function $\sigma : [-\bar{v}, \bar{v}] \rightarrow \mathbb{R}$ with a finite discontinuity set, the σ -statistic is

$$E[\sigma(v_i) \mid \omega]$$

Put differently, the σ -statistic is the probability limit of the sample moment, i.e.,

$$\text{plim}_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \sigma(v_i)$$

The researcher's data consists of the messages m_1, \dots, m_n sent by the sample of n agents in equilibrium. Traditionally, a statistic is identified if there exists an estimator whose estimate converges to said statistic in the limit as the sample size n grows. The same is true in our model, except that whether a statistic is identified also depends on the sequence of mechanisms that generates the data.

Definition 7. Fix a sequence of mechanisms $(\mathbf{x}^n, \mathbf{t}^n)$ and estimators $f_n : (M_i)_{i=1}^n \rightarrow \mathbb{R}$. A σ -statistic is strongly identified if, for all instances $I \in \mathcal{I}$ and equilibria $(\mathbf{s}^n, \mathbf{m}^n)$,

$$f_n(m_1, \dots, m_n) \rightarrow_p E[\sigma(v_i) \mid \omega] \quad \text{where} \quad m_i = \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i)) \quad (6)$$

A σ -statistic is weakly identified if, for all instances $I \in \mathcal{I}$, condition (6) holds for some sequence of equilibria $(\mathbf{s}^n, \mathbf{m}^n)$.

Just as the researcher is free to recruit as many agents as he wishes, he is free to choose any mechanism he wishes. He can *strongly elicit* a statistic if he can design some sequence of mechanisms for which the statistic is strongly identified.

Definition 8. A σ -statistic is strongly elicitable if there is a sequence of mechanisms and estimators that strongly identify it.

We also consider *weak elicitation*, which allows for “bad equilibria” where the researcher's estimate does not converge to the σ -statistic. It has two conditions. First, the researcher must weakly identify the statistic. That is, there must exist a “good equilibrium” where the researcher's estimate converges to the σ -statistic. Second, the researcher's estimate can never be less accurate than a *no-information benchmark*. In other words, even if there are bad equilibria, they should not be unreasonably bad.

The no-information benchmark compares the mean-square error of the

researcher's estimate with the ex ante variance of the σ -statistic, i.e.,

$$\text{Var}[\mathbb{E}[\sigma(v_i) \mid \omega]]$$

We call this the no-information benchmark because it represents the best-attainable mean-square error in an alternative model where agents cannot acquire signals.

Definition 9. A sequence of mechanisms $(\mathbf{x}^n, \mathbf{t}^n)$ and estimators f_n satisfy the no-information benchmark if, for all instances $I \in \mathcal{I}$ and equilibria $(\mathbf{s}^n, \mathbf{m}^n)$,

$$\limsup_{n \rightarrow \infty} \mathbb{E} \left[(f_n(m_1, \dots, m_n) - \mathbb{E}[\sigma(v_i) \mid \omega])^2 \right] \leq \text{Var}[\mathbb{E}[\sigma(v_i) \mid \omega]]$$

with messages $m_i = \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i))$.

Definition 10. A σ -statistic is weakly elicitable if there is a sequence of mechanisms and estimators that (i) weakly identify it and (ii) satisfy the no-information benchmark.

Remark 3. It is certainly possible to define an even weaker form of elicitation that only requires weak identification. This would entirely ignore bad equilibria, including “unreasonably bad” equilibria that lead to worse estimates than what could be achieved if agents did not acquire any information.

Absent a compelling criterion for equilibrium selection, we do not find this weaker form of elicitation to be especially convincing or robust. Putting the merits aside, we do not study this weaker form of elicitation because (i) doing so would not affect our existing positive results, and (ii) we expect that any new positive results would involve mechanisms that we consider particularly unrealistic.

To illustrate (ii), consider a mechanism where the researcher first asks agents to report the instance $I = (\Omega, \mathcal{V}, \Theta, \mathcal{S}, F, c)$ and then harshly penalizes all of the agents if their reports are not identical. There is an equilibrium where agents are truthful, alongside many other equilibria where agents are not truthful. Learning the instance I does not solve or trivialize the researcher's problem, since knowing the joint distribution F is different from knowing the realized distribution $F(\omega)$. However, it may expand the set of statistics that the researcher can elicit.

Finally, we stress that a researcher is likely to face challenges in practice that are not captured in this model. He may only have access to a small number n of agents, or find it costly to recruit more. He may not be able to design any mechanism he wishes.

He may face budget constraints. As a first step, we ask what an empowered researcher can achieve when these particular challenges are not present. Even so, we find that what the researcher can achieve is quite limited.

3 Strong Elicitation

It turns out that non-trivial statistics are not strongly elicitable.

Theorem 1. *A σ -statistic is strongly elicitable if and only if σ is constant.*

The strength of Theorem 1 reflects three features of our model. First, we insist on point identification rather than partial identification. It may be possible to obtain informative bounds for a given statistic, even if it is not point identified. Second, strong elicitation requires identification in every equilibrium. Theorem 2 shows how our results change when we relax this requirement. Third, following Assumption 1, we insist on elicitation for a rich set of instances (i.e., the basic instances).

With that said, these three features would all be without loss in classical models of preference elicitation where agents know their values (e.g., Becker et al. 1964; Clarke 1971). Since we obtain a negative result in settings where classical models obtain positive results, Theorem 1 tells us that stronger assumptions are needed to learn about a population’s preferences when processing information is costly.

Next, we provide a high-level intuition for Theorem 1. To understand how to turn this high-level intuition into a proof, see the proof outline in Appendix A.1.

3.1 Intuition for Theorem 1

We provide intuition using a simple example. The technology spaces is a singleton. The state space is $\Omega = \{-1, 1\}$. For coefficients $\alpha, \beta \in \mathbb{R}$, agent i ’s value is

$$v_i = \beta\omega + \epsilon_i \quad \text{where} \quad \epsilon_i \sim \text{UNIFORM}\{-\alpha, \alpha\} \text{ i.i.d.}$$

There are two signals: a revealing signal that costs $\bar{c} \gg \alpha$ and reveals the value v_i ; and an uninformative signal that reveals nothing and costs nothing.

There are two cases to consider. We begin with σ -statistics where σ is non-affine, and then turn to the case where σ is affine (but not constant).

3.1.1 Why Non-Affine Statistics are Not Strongly Elicitable

Suppose that $\beta = 0$, so that the values v_i are uncorrelated across agents i .

This is an easy case when it comes to eliciting affine statistics, like the population's average value. The researcher can ask each agent to make a report \hat{v}_i equal to their willingness to pay.² If they do not acquire information, each agent i reports $\hat{v}_i = 0$ and the average reported value is 0. If they acquire information, each agent i reports $\hat{v}_i = v_i$ and the average reported value converges to the average $E[v_i | \omega] = 0$ by the law of large numbers. The average is identified either way.

Interestingly, this setting is a hard case when it comes to eliciting non-affine σ -statistics. There are two reasons for this.

1. *In this example, most statistics are hard to identify if agents are not informed.*

To see this, consider a σ -statistic that is not an average. Typically, the σ -statistic will be sensitive to the parameter α . If agents do not acquire information, their willingness to pay does not pin down the σ -statistic.

2. *In this example, it is difficult to incentive agents to acquire information.*

There are two broad reasons for agent i to acquire information. The first reason is to decide whether she is willing to “buy” the product at a given price. This is not sufficient because the information cost \bar{c} far exceeds the agent's uncertainty over her value v_i . The second reason is to better predict the messages m_{-i} sent by other agents. This is not sufficient because agent i 's value v_i is independent of any private information that the other agents might have.

One way to circumvent these difficulties is by having agents directly report the parameter α – or more generally, the instance I – and penalize them if any of their reports disagree. This does not require agents to acquire information because the instance I (and therefore, α) is common knowledge. The problem with this mechanism is that it has many equilibria, most of which are non-truthful.

It is possible to imagine many other mechanisms, including ones that take advantage of the special structure of basic instances. Recall that, in a basic instance, a positive fraction of agents learn their value v_i at zero cost. If this were true in our simple example, the researcher could identify the parameter α in two steps. First, ask

²It is easy to incentivize truth-telling using the Becker–DeGroot–Marschak (BDM) mechanism.

agents to report their willingness to pay. Second, estimate the parameter α by excluding all apparently-uninformed reports $\hat{v}_i = 0$ from the data.

The challenge at this point is to rule out all possible mechanisms, for any non-affine σ -statistic, and for any collection of instances \mathcal{J} that includes the basic instances. We leave this challenge to the proof outline in Appendix A.1.

3.1.2 Why Affine Statistics are Not Strongly Elicitable

Suppose that $\beta = 1$, so that the values v_i are correlated across agents i .

The issues that arise are similar to the non-affine case. First, observe that affine statistics – say, the population’s average value – are proportional to the state ω . If no agent acquires information, then agents’ willingness to pay is not correlated with the state ω . Furthermore, when no one acquires information, acquiring information does not help agent i predict the messages m_{-i} that other agents send. This makes it hard to sustain information acquisition in every equilibrium, unless the costs \bar{c} are low relative to the agents’ uncertainty about their values (captured by α).

In this example, the uninformative equilibrium could be ruled out by assuming that a positive fraction of agents learn their value v_i or the state ω at zero cost. In basic instances, which have this property, one could construct a mechanism where agents with the bad technology acquire information in order to predict how agents with the good technology will behave. Even here, bad equilibria arise if (i) agents with the bad technology care about the state, while (ii) agents with the good technology do not.

Finally, note that the $\beta = 1$ case is also problematic for many non-affine statistics. In contrast, the $\beta = 0$ case is not problematic for affine statistics. This asymmetry hints at the possibility that it may be easier to learn about affine statistics than about non-affine statistics. Our next result confirms that this is the case.

4 Weak Elicitation

Since strong elicitation is impossible, we turn to weak elicitation. Fortunately, we find that the researcher can weakly elicit the population’s average value, and other affine statistics. Unfortunately, he cannot weakly elicit anything else.

Theorem 2. *A σ -statistic is weakly elicitable if and only if σ is affine.*

Many natural and economically-relevant statistics are not affine. For example, consider the quantity demanded by an informed population at some fixed price p . This is proportional to the σ -statistic where

$$\sigma(v_i) = \mathbf{1}(v_i \geq p)$$

This statistic is likely to be relevant for firms trying to assess market size or develop pricing strategies for new products. In our running example, it captures long-run demand as consumers gain first- and second-hand experience with GPT-5. Theorem 2 says that the firm cannot even weakly elicit this statistic.

With that said, the population’s average value is affine and is an important statistic in its own right (see Section 5). Theorem 1 says that the population’s average value is not strongly elicitable, and therefore “bad equilibria” are inevitable. However, Theorem 2 says that it is weakly elicitable. That is, the researcher can guarantee the existence of “good equilibria” where he learns the population’s average value, while also ensuring that his estimates are not unreasonably bad in any equilibrium.

The rest of this section provides intuition for Theorem 2. In Section 4.1, we describe the mechanisms we use to weakly elicit the population’s average value. In Section 4.2, we explain at a high level why the mechanism works, and why we cannot weakly elicit other statistics. Readers interested in a deeper understanding should refer to the examples in Appendix A.2 and the proof outline in Appendix A.3.

4.1 Mechanism

We propose a mechanism for weakly eliciting the population’s average value (and equivalently, any other affine statistic). It presents agents with an incentivized survey and asks them, before completing the survey, to predict the survey results.

More precisely, we build on the *Becker–DeGroot–Marschak (BDM) mechanism*. This mechanism is widely-used in behavioral experiments to elicit willingness to pay. It asks each agent to report their values \hat{v}_i , draws a random price p , and sells the product to the agent if and only if the reported value exceeds the price.

We also rely on *proper scoring rules*. These are rules that can be used to incentivize an agent i to report her beliefs about some random variable.³ Here, the

³Proper scoring rules tend to give higher scores to beliefs that assign higher probability to the observed value. There are many known proper scoring rules and they are easy to construct (e.g., McCarthy 1956).

random variable is the average reported value in the BDM mechanism, among agents other than i , i.e.,

$$\tilde{v}_i = \frac{1}{n-1} \sum_{j \neq i} \hat{v}_j$$

We denote agent i 's beliefs over the average reported value \tilde{v}_i by

$$b_i \in \mathcal{B} = \Delta(\mathbb{R})$$

With this notation, we can define proper scoring rules in our setting.

Definition 11. A scoring rule for agent i maps her reported belief $\hat{b}_i \in \mathcal{B}$ and the average reported value $\tilde{v}_i \in \mathbb{R}$ to a numerical score, i.e.,

$$\text{SR} : \mathcal{B} \times \mathbb{R} \rightarrow \mathbb{R}$$

It is proper if she maximizes the expected score by reporting her beliefs truthfully, i.e.,

$$\forall b_i \in \mathcal{B}, \quad b_i \in \arg \max_{\hat{b}_i} E_b[\text{SR}(\hat{b}_i, \tilde{v}_i)]$$

We can now introduce the *BDM-with-betting mechanism*. It involves two stages. The second stage is simply the BDM mechanism. The first stage asks agents to predict the average of the reported values in the second stage. Agents are paid more if their predictions turn out to be more accurate.

Definition 12. The BDM-with-betting mechanism (\mathbf{x}, \mathbf{t}) is parameterized by a proper scoring rule SR and scaling parameter λ . Each agent i sends a message

$$m_i = (\hat{v}_i, \hat{b}_i) \in \mathbb{R} \times \mathcal{B}$$

that consists of a reported value \hat{v}_i and a reported belief \hat{b}_i . She receives the product if her reported value exceeds the random price p , i.e.,

$$\mathbf{x}_i(\hat{v}, \hat{b}) = \mathbf{1}(\hat{v}_i \geq p) \quad \text{where} \quad p \sim \text{UNIFORM}[-\bar{v}, \bar{v}]$$

We mostly rely on the *continuous ranked probability score* and the *quadratic scoring rule*.

She is charged if she receives the product and earns a bonus from the scoring rule, i.e.,

$$t_i(\hat{v}, \hat{b}) = \lambda \cdot \text{SR}(\hat{b}_i, \hat{v}_i) - p \cdot \mathbf{1}(\hat{v}_i \geq p)$$

We argue that BDM-with-betting is simple compared to other mechanisms that elicit preferences by having agents bet on each others' reports (e.g., Crémer and McLean 1988, Miller et al. 2005). There are two ways in which it is simple.

1. *Estimation is straightforward.*

The researcher's estimate of the population's average value is the average reported value, i.e.,

$$f_n(m_1, \dots, m_n) = \frac{1}{n} \sum_{i=1}^n \hat{v}_i$$

In particular, the researcher does not try to estimate the population's average value from their reported beliefs.

2. *Incentives for truth-telling are straightforward.*

The second stage (BDM) incentivizes agents to truthfully report their willingness to pay given any information they acquired. The first stage (betting) plays no role when it comes to incentivizing truthfulness; the researcher does not use bets in the first stage to detect or disincentive lies in the second stage.

In contrast, other mechanisms that elicit preferences by having agents bet on each others' reports tend to use those bets in more complicated ways.⁴

What is interesting about BDM-with-betting is not the mechanism itself. It is the fact that it can incentivize agents to acquire precisely the kind of information needed to identify the population's average value (and not much else). Next, we explain why.

4.2 Intuition for Theorem 2

We first explain why the BDM-with-betting mechanism can be used to weakly elicit the population's average value (and, equivalently, any other affine statistic). Then we

⁴Full surplus extraction mechanisms implicitly use bets to make inferences about agents' willingness to pay (e.g., Crémer and McLean 1988). Peer prediction mechanisms have agents report signal realizations, which are implicitly treated as bets on others' reported signal realizations (e.g., Miller et al. 2005). In both cases, the designer uses betting as a way to incentivize truth-telling. In our setting, incentivizing truth-telling is straightforward, and betting is simply a way to incentivize information acquisition.

briefly explain why non-affine statistics cannot be weakly elicited.

4.2.1 Why BDM-with-Betting Works

The BDM-with-betting mechanism may seem naive. It does not incentivize agents to learn their values v_i , at least not in general. Instead, it incentivizes agents to learn about others' reported values \hat{v}_{-i} . Why is that enough to weakly elicit the population's average value? Why is it not enough to weakly elicit any other statistic?

Recall the simple example from Section 3.1. The technology space is a singleton and the state space is $\Omega = \{-1, 1\}$. For coefficient $\beta \in \mathbb{R}$, agents' values are

$$v_i = \beta\omega + \epsilon_i \quad \text{where} \quad \epsilon_i \sim N(0, 1) \text{ i.i.d.}$$

There are two signals: a revealing signal that costs $\bar{c} \gg \alpha$ and reveals the value v_i ; and an uninformative signal that reveals nothing and costs nothing. There are two cases to consider.

Uncorrelated Values. Let values v_i be uncorrelated across agents (i.e., $\beta = 0$). Recall from Section 3.1 that this is the hard case for eliciting statistics other than the average, because it is generally impossible to incentivize agents to acquire the revealing signal. However, the average has a special property: it is not necessary for agents to acquire information when their values are uncorrelated. To see this, compare the researcher's estimate if the agents acquire information, i.e.,

$$\frac{1}{n} \sum_{i=1}^n v_i = \frac{1}{n} \sum_{i=1}^n \epsilon_i \rightarrow_p 0$$

to the researcher's estimate if the agents do not acquire information, i.e.,

$$\frac{1}{n} \sum_{i=1}^n E[v_i] = 0$$

Either way, the researcher's estimate is correct in the limit.

More generally, when values are uncorrelated, the researcher can take advantage of the *wisdom of the crowd*. It is not necessary that agents learn their values perfectly. What is important is that the errors $\hat{v}_i - v_i$ they make are uncorrelated across agents i .

This property is what sets the average apart from other statistics.

Correlated Values. Let values v_i be correlated across agents (say, $\beta = 1$). If the betting stakes λ are sufficiently high, then agent i will acquire the revealing signal if it helps her predict the average reported value \tilde{v}_i . Suppose that agent i expects all other agents $j \neq i$ to acquire their revealing signals and learn their values v_j . Then she expects the average reported value \tilde{v}_i to be

$$\tilde{v}_i = \frac{1}{n-1} \sum_{j \neq i} v_j \rightarrow_p \omega$$

Ideally, agent i would learn directly about the state ω , but that is not feasible in this example. But she can learn about her value v_i , which is correlated with the state ω . Therefore, betting incentivizes her to acquire the revealing signal.

There are two equilibria in this example, provided that the betting stakes λ are sufficiently large. In the informed equilibrium, all agents acquire the revealing signal and report their values $\hat{v}_i = v_i$. The researcher's estimate is consistent, i.e.,

$$\frac{1}{n} \sum_{i=1}^n \hat{v}_i = \frac{1}{n} \sum_{i=1}^n v_i = \omega$$

In the uninformed equilibrium, all agents acquire the uninformative signal and report their expected values $\hat{v}_i = E[\omega]$. Now, the researcher's estimate is not consistent, i.e.,

$$\frac{1}{n} \sum_{i=1}^n \hat{v}_i = \frac{1}{n} \sum_{i=1}^n E[v_i] = E[\omega]$$

There is no wisdom of the crowd in the uninformed equilibrium, since errors $\hat{v}_i - v_i$ are correlated across agents i .

This multiplicity does not always arise (e.g., Example 1 in Appendix A.2) but it is unavoidable in general (Theorem 1). Weak elicitation allows for the existence of the uninformed equilibrium as long as the researcher's estimate is at least as good as the no-information benchmark. In this case, the mean-square error of the researcher's estimate is $\text{Var}[\omega]$. This precisely matches the no-information benchmark.

Generalizing the Example. In general, there is no clean separation between “uncorrelated values” and “correlated values”. Whether values are correlated depends on what information agents begin with and the properties of the signals they acquire. This is especially true for non-binary signal spaces, where agents may deviate from the revealing signal in many ways. For example, they may prefer to learn directly about the state ω , or only acquire partial information about their values v_i .

Furthermore, a general proof must account for counter-intuitive phenomena that can arise for highly-unstructured signal spaces and joint distributions. For example, it is possible for values to be uncorrelated when agents acquire the revealing signal, but correlated when agents do not (e.g., Example 3 in Appendix A.2). This makes the best response map unstable – since each agent acquires information only if other agents do not – and suggests that pure-strategy equilibria may not exist. Rather than attempting to construct mixed-strategy equilibria in strategy spaces of arbitrary size and dimension, we take a non-constructive approach to the proof.

Despite these difficulties, we find that BDM-with-betting can always sustain a sequence of equilibria where a wisdom-of-the-crowd effect holds in the limit. That is, betting incentivizes agents to acquire information until the point where errors $\hat{v}_i - v_i$ are (approximately) uncorrelated. We outline the proof in Appendix A.3.

4.2.2 Why Non-Affine Statistics are Not Weakly Elicitable

We have already explained why the BDM-with-betting mechanism, as currently constructed, cannot be used to weakly identify non-affine statistics. To recap, in the uncorrelated case of the example discussed above, BDM-with-betting does not incentivize information acquisition. The wisdom of the crowd means that this is not a problem for affine statistics, but it does not apply to non-affine statistics.

To see why no mechanism can weakly elicit non-affine statistics, we refer to the following example. If the joint distribution F puts probability one on a single state ω , then values v_i are independent across agents i . In that case, the no-information benchmark has a mean-square error of

$$\text{Var}[E[\sigma(v_i) \mid \omega]] = \text{Var}[E[\sigma(v_i)]] = 0$$

In other words, when the state ω is fixed with probability one, even agents that acquire no information can infer any statistic of the realized distribution $F(\omega)$.

Although this example may seem like an easy case for the researcher, the intuition for Theorem 1 tells us that it is not. Recall from Section 3.1 that the only mechanisms that appear capable of identifying non-affine statistics involve asking the agents for information about the instance I . Such mechanisms inevitably feature equilibria where agents mislead the researcher, and these equilibria can lead to a positive mean-square error that exceeds the no-information benchmark. That is, these mechanisms open the door to outcomes that are worse than if the researcher had simply given up on the idea of incentivizing agents to acquire information.

At this point, we have fully characterized the set of elicitable statistics, under both the strong and weak forms of elicibility. These results are largely negative, but suggest that it may be possible to learn about the population's average value. Next, we explore an important class of applications where eliciting the population's average value is precisely what is needed: social choice.

5 Application to Social Choice

We now turn to the problem of social choice, where a planner must choose a single alternative on behalf of a population. Specifically, we seek social choice mechanisms that work well even when processing information is costly.

It turns out that the methods we developed to elicit preferences can be used to design social choice mechanisms that are robust to costly information processing. This new setting raises new theoretical challenges, but it also broadens the set of potential applications. For example:

1. *Democratic elections and referendums, especially at the local level.*

Empirical studies suggest that many voters are poorly-informed (e.g., Delli Carpini and Keeter 1996) and may vote differently if they were better informed (e.g., Bartels 1996). If voters are not processing the information available about a candidate or ballot measure, they could easily make misguided decisions.

2. *Corporate governance.*

Shareholders routinely elect board members and vote on issues that affect their interests (e.g., whether to approve a merger). Although large shareholders have strong financial incentives to stay well-informed, smaller shareholders and retail investors may not (e.g., Kastiel and Nili 2016, Brav et al. 2023).

3. Collective decision-making in organizations.

For example, a committee may make decisions on behalf of the faculty at a university. Alternatively, a workforce may hold votes on whether to unionize, who to elect as a union official, etc. As in the previous applications, individuals may have little incentive to do their due diligence before voting.

Before presenting our result (Theorem 3), we describe a model of social choice (Section 5.1) and the challenges that good mechanisms must overcome (Section 5.2). Then we state our result (Section 5.3), propose a mechanism (Section 5.4), and explain why it works (Section 5.5).

5.1 Model

We convert our model of preference elicitation into a model of social choice.

The researcher is now a planner. The n agents are now voters participating in a committee (or citizen's assembly, mini-public, jury, etc.) that take actions on behalf of a larger population. Alternatives $x \in \mathcal{X}$ indicate whether a measure passes ($x = 1$) or not ($x = 0$). The mechanism must assign the same alternative x to all agents.

Assumption 4. *Restrict attention to mechanisms (\mathbf{x}, \mathbf{t}) where, for all message profiles m ,*

$$\mathbf{x}_1(m) = \mathbf{x}_2(m) = \dots = \mathbf{x}_n(m)$$

In a slight abuse of notation, let $\mathbf{x}(m)$ refer to the one alternative chosen for all agents.

We maintain the assumption that the planner can recruit as many agents as he likes and spend as much money as he likes. This strikes us a reasonable starting point for large populations, like the residents of a large municipality or the employees of a major corporation. Since the population is large, even a large sample of agents will only be a small fraction of the population, and the cost of transfers to n agents will tend to be small relative to the potential gains for everyone else.

The planner seeks a sequence of mechanisms that *converges to efficiency*. We measure efficiency by comparing the expected *welfare* of a given mechanism with *first-best welfare*. Welfare is the population's average value from an alternative x , i.e.,

$$E[v_{ix} \mid \omega]$$

Expected welfare for instance I , mechanism (\mathbf{x}, \mathbf{m}) , and strategy profile (\mathbf{s}, \mathbf{m}) is

$$W(I, \mathbf{s}, \mathbf{m}, \mathbf{x}, \mathbf{t}) = \mathbb{E}[v_{ix(m)} \mid \omega]$$

with messages $m_i = \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i))$. Finally, first-best welfare is the highest expected welfare that the planner could achieve, if he knew the instance I and the hidden state ω . Formally, it is

$$\text{OPT}(I) = \mathbb{E}\left[\max_x \mathbb{E}[v_{ix} \mid \omega]\right]$$

We ignore the cost of information acquisition in all of these definitions, but this is essentially without loss if the population is large. After all, only the committee members will be asked to acquire information. Their information costs will tend to be small relative to the potential gains for the broader population.

Definition 13. Fix an instance I . A mechanism (\mathbf{x}, \mathbf{t}) is strongly α -efficient if expected welfare is within α of first-best for every equilibrium (\mathbf{s}, \mathbf{m}) , i.e.,

$$W(I, \mathbf{s}, \mathbf{m}, \mathbf{x}, \mathbf{t}) \geq \text{OPT}(I) - \alpha \quad (7)$$

It is weakly α -efficient if inequality (7) holds for some equilibrium (\mathbf{s}, \mathbf{m}) . Let the smallest constants α where (\mathbf{x}, \mathbf{t}) are strongly and weakly α -efficient be called $\alpha^S(I, \mathbf{x}, \mathbf{t})$ and $\alpha^W(I, \mathbf{x}, \mathbf{t})$, respectively.

The definition of convergence to efficiency mirrors the definition of weak elicitation. That is, we ignore inefficient equilibria as long as (i) they are not less efficient than a no-information benchmark and (ii) there is an efficient equilibrium. The no-information benchmark, in this case, means that expected welfare is at least as high as if the ex-ante optimal alternative had been chosen.

Definition 14. A sequence $(\mathbf{x}^n, \mathbf{t}^n)$ of mechanisms converges to efficiency if, for every instance $I \in \mathcal{J}$,

$$\lim_{n \rightarrow \infty} \bar{\alpha}^W(I, \mathbf{x}^n, \mathbf{t}^n) = 0 \quad \text{and} \quad \limsup_{n \rightarrow \infty} \bar{\alpha}^S(\mathbf{x}^n, \mathbf{t}^n) \leq \text{OPT}(I) - \max_x \mathbb{E}[v_{ix}]$$

Neither of these requirements is trivial. In particular, standard voting procedures do not necessarily satisfy the no-information benchmark (see e.g., Ali et al. 2025,

Section 1 for an excellent example). Next, we explain what challenges arise.

5.2 Challenges

When trying to solve the planner's problem, we encounter challenges that do not arise in models of social choice where agents know their own values. To explain these challenges, we use the *majority-rule mechanism* as a foil.

Definition 15. *The majority-rule mechanism (\mathbf{x}, \mathbf{t}) asks each agent i to report her preferred alternative $\hat{x}_i \in \{0, 1\}$. Then it sets*

$$\mathbf{t}(\hat{x}_1, \dots, \hat{x}_n) = 0 \quad \text{and} \quad \mathbf{x}(\hat{x}_1, \dots, \hat{x}_n) = \mathbf{1} \left(\frac{1}{n} \sum_{i=1}^n \hat{x}_i \geq \frac{1}{2} \right)$$

In addition to the standard problems with majority rule, there are three additional problems that arise when information acquisition is costly.

1. *Majority rule only rewards information acquisition in the event that an agent is pivotal. This is unlikely (e.g., Mulligan and Hunter 2003).*

Agent i is pivotal if changing her report \hat{x}_i would change the alternative selected by the mechanism. If she is not pivotal, then information has no instrumental value. The probability that an agent is pivotal tends to vanish quickly as the sample size n grows, so there is little expected gain from acquiring information. This is typically true even if n is relatively small.

2. *Majority rule does not account for the positive externalities of information acquisition, where one agent's informed decision can benefit everyone else.*

For example, suppose that agents have common values and that agent i knows she will be pivotal. Acquiring information that identifies the optimal alternative would not only improve her well-being, but also the well-being of everyone else. Since agent i does not internalize this benefit to others, she may acquire less information than the planner would like.

3. *When agents are imperfectly informed, majority rule does not incentivize them to truthfully report the alternative they prefer given the information they have. This can distort the outcome of the vote (e.g., Austen-Smith and Banks 1996).*

Since agent i 's report only matters if she is pivotal, it is optimal for her to condition her beliefs on being pivotal before making a report. For example, suppose that agent i believes a ballot measure is likely to be both good for her and very popular among the other agents. If she is truthful, she would vote in favor of the measure. However, if she conditions on her being pivotal, then she focuses on a contingency where the measure was not as popular as expected. This may lead her to revise her beliefs about her own value v_i downwards.

We stress that these problems are not unique to majority rule. They apply to many other mechanisms – like the Vickrey-Clarke-Groves (VCG) mechanism, or quadratic voting – that do not account for voters' costs of processing information.

Of these challenges, the third one is the most daunting. The first two deal with incentives to acquire information, which we already studied in the preference elicitation context. The third challenge deals with incentives to be truthful when voters have not acquired the revealing signal. We cannot rely on BDM to incentivize truth-telling, because it relies on the ability to allocate different alternatives to different agents. Furthermore, the intuition for Theorems 1-2 suggests that it is not always possible to incentivize voters to become fully informed.

Fortunately, it turns out that we can address all three challenges at once, using a simple adaptation of our method for eliciting the population's average value.

5.3 Statement of Theorem 3

We can now state the main result of our application to social choice.

Theorem 3. *There is a sequence of mechanisms that converge to efficiency.*

The mechanism we use to prove Theorem 3 builds on the Vickrey-Clarke-Groves (VCG) mechanism. Unfortunately, this mechanism has been criticized as impractical (see e.g., Rothkopf 2007). It also involves a form of vote buying, which is a nonstarter in many applications. For these reasons, we now restrict attention to a special case of the model where we can motivate a more realistic mechanism.

Assumption 5. *Restrict attention to instances I with a binary value space $\mathcal{V} = \{-1, 1\}$ and a binary signal space $\mathcal{S} = \{s_i^R, s_i^U\}$, where s_i^R is a revealing signal and s_i^U is an uninformative signal that sets $s_i^U(\cdot) = 0$. This overrides Assumption 1.*

Assumption 5 allows us to build on the majority-rule mechanism. Majority rule is widely-used in practice, but fails to account for the intensity of voters’ preferences. For example, a minority that cares deeply about an issue may be overruled by a majority that is nearly indifferent. This means that majority rule is generally inefficient, even if voters know their values v_i . Assumption 5 ensures that voters’ preferences are equally intense, since their values are $v_i \in \{-1, 1\}$.⁵

We maintain Assumption 5 for the rest of the main text and Appendix A. We leave the general case of Theorem 3 to the Supplemental Appendix.

5.4 Mechanism

Assumption 5 allows us to motivate the *majority-rule-with-betting mechanism*. Majority-rule-with-betting has voters cast their votes according to the majority rule and asks them, before they cast their votes, to predict the results.

To be more precise, majority-rule-with-betting has voters cast their votes according to the majority rule with high probability. For technical reasons, there is also a small probability $\delta > 0$ that a randomly-chosen agent dictates the outcome.⁶

To define this mechanism, we need to update our notation from Section 4.1. Let the *vote share* from agent i ’s perspective be

$$\tilde{n}_i = \frac{1}{n} \sum_{j \neq i} \hat{x}_j$$

Let agent i ’s belief over the vote share be

$$b_i \in \mathcal{B} = \Delta([0, 1])$$

A scoring rule maps a reported belief $\hat{b}_i \in \mathcal{B}$ and a vote share $\tilde{n}_i \in [0, 1]$ to a score, i.e.,

$$\text{SR} : \mathcal{B} \times [0, 1] \rightarrow \mathbb{R}$$

⁵By restricting attention to all-or-nothing information acquisition, it also avoids complications that arise when voters are partially-informed. For example, when voters are partially-informed, a minority of voters who are confident that the measure is bad for them may be overruled by a majority of voters who believe it is slightly more likely that the measure is good for them.

⁶This ensures that voters are pivotal with positive probability and therefore avoids situations where voters must condition on zero probability events when casting their votes.

Definition 16. The majority-rule-with-betting mechanism (\mathbf{x}, \mathbf{t}) is parameterized by a proper scoring rule SR, scaling parameter λ , and probability δ . Agent i sends a message

$$m_i = (\hat{x}_i, \hat{b}_i) \in \{0, 1\} \times \mathcal{B}$$

that consists of a reported alternative \hat{x}_i and a reported belief \hat{b}_i . With probability $1 - \delta$, the planner selects the alternative that the majority reports, i.e.,

$$x(\hat{x}, \hat{b}) = \mathbf{1} \left(\frac{1}{n} \sum_{i=1}^n \hat{x}_i \geq \frac{1}{2} \right)$$

With probability δ , the planner randomly chooses a voter $i \sim \text{UNIFORM}(1, \dots, n)$ and selects their reported alternative, i.e.,

$$x(\hat{x}, \hat{b}) = \hat{x}_i$$

Finally, each agent i is paid for her prediction according to the scoring rule, i.e.,

$$\mathbf{t}_i(\hat{v}, \hat{b}) = \lambda \cdot \text{SR}(\hat{b}_i, \hat{n}_i)$$

This mechanism has two desirable features. First, it does not allow anyone to buy votes. Transfers are used only to reward accurate predictions. Second, the mechanism counts all votes equally. In particular, voter i 's ability to predict the vote share does not affect how much weight is attached to her vote.

Majority-rule-with-betting may not appear that different from the status quo in many democracies, where voters are free to participate in political betting markets. Still, there are two differences. First, we do not ask voters to bet against each other. Doing so would give rise to the Grossman and Stiglitz (1980) paradox and undermine most voters' incentive to acquire information. Second, we rely on *sortition* – that is, random sampling of the population. Voters are members of a citizen's assembly (or a committee, etc.) tasked with making a decision on behalf of the population. Sortition has been used in various forms by democracies dating back to ancient Athens. In the United States, for example, sortition is used to form jury pools.

5.5 Intuition for Theorem 3

The rest of this section provides a high-level intuition for Theorem 3.

Recall the three challenges from Section 5.2. The first two challenges concern incentives to acquire information. The reasons why betting incentivizes information acquisition here are essentially the same as the reasons why betting incentivizes information acquisition in Theorem 2, which we discuss in Section 4.2

The third challenge, which concerns incentives to be truthful, is new. Fortunately, it turns out that the kind of information that voters need to acquire in order to be truthful is closely-related to the kind of information that is relevant to predicting the vote share. This means that a convenient side effect of betting on the vote share is that it restores truth-telling incentives.

To see this, recall why majority rule may not incentivize truth-telling on its own. Agent i knows that her vote \hat{x}_i is pivotal only if the vote share \tilde{v}_i is exactly 50%. Therefore, it is in her best interest to vote in favor of the measure whenever her expected value v_i conditional on being pivotal is positive, i.e.,

$$E[v_i \mid s_i, s_i(\omega, \theta_i, v_i), \tilde{v}_i = 50\%] \geq 0 \quad (8)$$

A truthful agent would not condition on pivotality and vote in favor whenever

$$E[v_i \mid s_i, s_i(\omega, \theta_i, v_i)] \geq 0 \quad (9)$$

What determines whether agent i will be truthful? It is whether agent i perceives the vote share \tilde{v}_i to be correlated with her value v_i even after conditioning on her information. If not, then conditioning on pivotality does not affect agent i 's beliefs about her value. As a result, conditions (8) and (9) will be the same.

The reason why high-stakes betting incentivizes agent i to acquire information until she is truthful is because it incentivizes her to acquire any information that is correlated with the vote share \tilde{v}_i . On the one hand, if v_i is correlated with the vote share \tilde{v}_i , then agent i benefits from acquiring her revealing signal. Once she acquires the revealing signal, she prefers to vote truthfully. On the other hand, if v_i is not correlated with the vote share \tilde{v}_i , then agent i also prefers to vote truthfully.

We leave further discussion to the proof outline in Appendix A.4.

6 Related Literature

We build on and contribute to several research areas in economics, computer science, political science, and marketing.

Revealed Preferences with Limited Attention. There is a sizable literature that studies the implications of limited or costly attention for revealed preferences (e.g., Masatlioglu et al. 2012; Manzini and Mariotti 2014; Cattaneo et al. 2020), demand estimation (e.g., Barseghyan et al. 2021a; Barseghyan et al. 2021b), and preference elicitation (e.g., Cao and Zhang 2021; Mamadehussene and Sguera 2023).

We contribute to this work by exploring the extent to which costly attention is a fundamental barrier to learning about preferences. These barriers hold regardless of how much money a researcher can spend, or what kind of experiment he sets up. Theoretically, these barriers extend to (i) limited attention, a special case of costly attention where attention is exogenous, and (ii) demand estimation, a special case of preference elicitation where the mechanism is exogenous.

We also propose a modification of the BDM mechanism, which can also apply to other methods of eliciting willingness to pay. This modification may strengthen incentives for information acquisition. Empirical work suggests that BDM may not be effective without such incentives (Mamadehussene and Sguera 2023).

Full Surplus Extraction. The literature on full surplus extraction asks when it is possible to incentivize agents to reveal their private types at no expected cost to the designer (e.g., Crémer and McLean 1988). Roughly, the designer asks each agent to place a bet on the reports of other agents. If types are correlated across agents, and these bets are designed carefully, the designer can infer agents' types from their bets. Although this literature typically focuses on incentives for truthful reporting, Bikhchandani (2010) studies the impact of information acquisition.

Although both involve betting, the mechanisms driving our positive results are simpler and more robust than full surplus extraction mechanisms. BDM-with-betting is detail-free and uses an off-the-shelf proper scoring rule to reward accuracy. Full surplus extraction involves fine-tuned transfer rules that are very sensitive to the details of the prior distribution (or a finite set of prior distributions; see Fu et al. 2021). Moreover, BDM-with-betting does not *require* agents' types to be correlated, whereas.

full surplus extraction only works if types are correlated in particular ways.

More fundamentally, our work differs from full surplus extraction in that we ignore the monetary costs of eliciting information. We try to understand the limits of what a researcher can learn, independently of how well-funded he is. In contrast, full surplus extraction insists on zero costs. It is not surprising that this more ambitious goal requires stronger assumptions and more complicated mechanisms.

Proper Scoring Rules. The literature on proper scoring rules assumes that the researcher can make transfers that depend on the realized state. In that case, there is a wide range of scoring rules that can incentivize agents to truthfully report their beliefs (see e.g., McCarthy 1956). When the researcher has access to unbounded transfers, as in our model, he can also incentivize information acquisition by rescaling these scoring rules. More recent work asks how to incentivize information acquisition when transfers are bounded (e.g., Li et al. 2022).

We do not assume that the researcher can make state-dependent transfers. In market research, even if the researcher eventually observes long-run demand, survey participants may not be willing to wait years before being paid. In social choice, where the state represents the effectiveness of a given policy, society may never reach a consensus around the effectiveness of that policy. Nonetheless, proper scoring rules play an important role in our positive results; we use them to incentivize agents to acquire information that predicts summary statistics of other agents' reports.

We draw particular inspiration from Lambert et al. (2008). Under the assumption that the researcher can make transfers that depend on the state, they ask what kinds of statistics can be elicited by a direct mechanism. We formulate our problem similarly, albeit in a different setting and without restricting attention to direct mechanisms.

Crowdsourcing. A sizable literature in both economics and computer science studies crowdsourcing mechanisms. These mechanisms encourage many agents to report information about a common state of nature. Perhaps the most prominent examples are prediction markets (e.g., Hanson 2003), which assume that the designer can make transfers that depend on the realized state. When the designer never directly observes the hidden state, mechanisms like the Bayesian truth serum (e.g., Prelec 2004) and peer prediction (e.g., Miller et al. 2005) incentivize truthful reporting.

Of particular relevance is work on peer prediction. Peer prediction mechanisms

tend to reward agents for making reports that are correlated with the reports of other participants. These rewards incentivize both truthful reporting and information acquisition, and the designer uses the reports to infer the hidden state. Much of this work shares our emphasis on costly information processing (e.g., Miller et al. 2005; Dasgupta and Ghosh 2013) and on robustness to the designer’s lack of knowledge (e.g., Shnayder et al. 2016; Kong and Schoenebeck 2019).

Although there are many similarities, peer prediction differs from preference elicitation in two fundamental ways.

1. *There is no product to allocate, or decision to make.*

In peer prediction, the designer can solicit messages m and set transfers t . He typically does not allocate a payoff-relevant alternative x . In our setting, the allocation is essential. We cannot study revealed preference without a product to consume, or social choice without a measure that impacts the population.

2. *Participants are learning about the same thing.*

In peer prediction, a designer encourages many participants to learn about the same object – the hidden state. For example, in peer grading, several students review the same assignment. In our setting, the designer would like each participant to learn about their own preferences. So, unless preferences are homogeneous, every participant is learning about a different object.

In light of these differences, neither negative results nor positive results necessarily translate from one problem to another. On the one hand, the fact that our designer can make allocations makes elicitation easier. On the other hand, the fact that each of our participants are learning about different things makes elicitation harder.

In addition to these fundamental differences, there are two notable modeling differences. First, we deal with multiple equilibria differently. Peer prediction mechanisms typically have what we call “unreasonably bad” equilibria, which violate weak elicitation, even after restricting attention to participant-optimal equilibria (e.g., Dasgupta and Ghosh 2013). They also have uninformative equilibria, which violate strong elicitation, and these equilibria sometimes arise in practice (Gao et al. 2014). Second, we only make mild assumptions on how participants acquire information. Most results on peer prediction make strong assumptions on how participants acquire information. Relaxing those assumptions can change the results (Gao et al. 2019).

It is also worth clarifying how our BDM-with-betting mechanism differs from peer prediction mechanisms. BDM-with-betting asks each participant to (i) report her willingness to pay for a product, and (ii) to predict the average reported willingness to pay of other participants. Here, (i) incentivizes truthful reporting and (ii) incentivizes information acquisition. In particular, truthful reporting is optimal regardless of what the other participants do. A typical peer prediction mechanism asks each participant to report her signal realization; she earns money if her reports are correlated with the reports of other participants. Here, truth-telling incentives are more complicated; participants only want to be truthful if other participants are also truthful.

Finally, two innovative papers use crowdsourcing to accomplish goals very much like our own: designing mechanisms that are robust to the possibility that agents do not know their own preferences. Schoenebeck and Tao (2021) study social choice. Pakzad-Hurson (2022) studies a more general setting. Interestingly, these papers do not need to assume that voters can learn their values in order to achieve efficiency. Unfortunately, we cannot adapt their mechanisms to our setting, for three reasons. First, they study particular information structures. Second, they assume agents learn about a common state. Third, they either focus on favorable equilibria or need stronger assumptions to rule out unfavorable equilibria.

Mechanism Design with Information Acquisition. We build on the literature on mechanism design with information acquisition. Much of this work focuses on different applications, like auctions (e.g., Persico 2000; Bergemann et al. 2009), matching (e.g., Immorlica et al. 2020), and settings with verifiable information (e.g., Ben-Porath et al. 2024). We are especially motivated by two negative results. First, Bergemann and Välimäki (2002) find that efficient mechanisms may not exist when information is costly and preferences are correlated. Second, Pei and Strulovici (2024) find that planners with sufficiently-imprecise knowledge of preferences and the costs of information can only implement trivial social choice rules.

Our work both applies and contributes to this literature. Specifically, it further develops applications to preference elicitation and social choice. It also proposes a model of unstructured information acquisition and insists on mechanisms that are robust to the details of how agents acquire information. We draw heavily from prior work in mechanism design that insists on robustness to unknown production or information acquisition technologies (e.g., Carroll 2015; Carroll 2019). Relative to this

work, we focus on a particular known technology: the revealing signal that allows an agent to learn her own value, but not necessarily anything else. We also build on work on unstructured information acquisition in games (e.g., Denti and Ravid 2024), which relies on very natural assumptions that happen to be different from ours.

Interestingly, some prior work highlights how incentives to learn about others can actually undermine efficiency (Pernoud and Gleyze 2024) and strategic simplicity (Gleyze and Pernoud 2023). In contrast, we obtain efficient social choice only after explicitly incentivizing agents to learn about each other. Moreover, this simplifies the agents' problem insofar as it restores incentives to report truthfully.

Voting and Social Choice. The literature on voting, and social choice more generally, has studied voter information both empirically and theoretically.

On the empirical side, the literature establishes several stylized facts that motivate our interest in social choice with costly information processing.

1. *Many voters are poorly informed.*

There are many poorly-informed voters, and better-informed voters are not demographically or ideologically representative of the electorate (e.g., Palfrey and Poole 1987; Delli Carpini and Keeter 1996; Angelucci and Prat 2024).

2. *Voters respond to information.*

In experiments, giving publicly-available information to voters can affect how they vote (e.g., Lau and Redlawsk 1997; Fowler and Margolis 2014), and people with different ideologies respond to information in similar ways (e.g., Coppock 2023). In practice, well-informed voters vote differently than poorly-informed ones, controlling for observable characteristics (e.g., Bartels 1996; Gilens 2001).

3. *Voter respond to incentives to acquire information.*

There is some evidence that voters trade-off the costs of processing information with the instrumental value of that information (e.g., Larcinese 2009; Shineman 2018). Moreover, in surveys, participants' political knowledge responds to monetary incentives (e.g., Prior and Lupia 2008).

4. *Increasing voter information has benefits.*

More informed electorates are better at identifying low-quality officials and removing them from office (see e.g., Pande 2011 and references therein).

On the theoretical side, we build on a line of work on combating uninformed voting (e.g., Persico 2004; Gerardi and Yariv 2008; Gershkov and Szentes 2009; Cai 2009; Tyson 2016). Relative to this work, we consider a much more powerful planner and obtain efficiency results under much weaker assumptions. Specifically, our planner has unlimited access to transfers. In return, we allow general forms of heterogeneity, correlation, and information acquisition. We also do not require the planner to have any prior knowledge. In a sense, Assumption 5 bridges our approach with that of prior work. It motivates a more practical mechanism – i.e., one that does not involve vote buying – in a more stylized model.

We also build on research that studies voting behavior when information is costly (e.g., Martinelli 2006; Feddersen and Sandroni 2006) or imperfect (e.g., Feddersen and Pesendorfer 1996; Jaehoon Kim and Fey 2007; Battaglini et al. 2010; Bhattacharya 2013; Ali et al. 2025). Some of this work predicts a swing voter’s curse, where poorly-informed voters abstain or appear to vote against their own interests, which is empirically supported by Battaglini et al. (2010). Relative to this work, we identify mechanisms that alleviate the swing voter’s curse in all equilibria, and encourage voter informedness in favorable equilibria.

Finally, we are motivated by research on sortition, as well as minipublics and citizen’s assemblies (e.g., Dahl 1989; Fishkin 2011). As in our model, sortition involves random sampling of voters drawn from a broader electorate. This practice raises a number of important political, economic, and computational questions that other work explores (e.g., Flanigan et al. 2021; Bardhi and Bobkova 2023).

7 Conclusion

Revealed preference plays a foundational role in economics, by linking observed choices with unobserved preferences. This link is disrupted when people find it costly to process information about the goods they consume, the services they receive, and the policies that affect them. We ask whether it is possible to repair that link. That is, we ask whether observed choices, in any possible mechanism and from arbitrarily many individuals, can identify the distribution of preferences in a population.

We find that it is not possible to fully restore the link between choices and preferences when processing information is costly (Theorem 1). It is possible to partially restore the link (Theorem 2), and that may be enough for some applications

(Theorem 3). For other applications, it may be possible to obtain informative bounds on statistics that are relevant to those applications. But it is clear that choices cannot reveal preferences at the same level of generality and precision as classical theory suggests (e.g., Becker et al. 1964; Clarke 1971).

Readers may object to our emphasis on informed preferences, rather than uninformed preferences. After all, as long as people are uninformed, it is their uninformed preferences that govern their behavior. But revealed preference is not only used to predict behavior, it is also used to measure welfare. When structural models are used to guide policy decisions, or theoretical models are used to motivate particular markets, it would be a mistake to equate welfare with people’s uninformed preferences. Even if the goal is to predict behavior, informed preferences matter for many applications (e.g., estimating long-run demand for a product).

Naturally, our results have weaknesses that may limit their practical relevance. For example, the researcher in our model can recruit an unlimited number of individuals and spend an unlimited amount of money. This strengthens our negative results, but weakens our positive results. In addition, we rely on Bayes-Nash equilibrium, a solution concept that is quite strong and often leads to unrealistic conclusions. Finally, our results are entirely theoretical. We do not provide empirical evidence that mechanisms like BDM-with-betting work in practice.

Nonetheless, we believe that further work that addresses these weaknesses could have broad practical relevance. In particular, there are many institutions where costly information processing can lead to market failures (e.g., elections), and many institutions that exist in part to correct such market failures (e.g., reputation and recommender systems). Understanding and addressing these market failures may ultimately require methods for eliciting informed preferences.

A Proof Outlines

A.1 Proof of Theorem 1

We split the proof of Theorem 1 into two cases, based on whether the σ -statistic is affine or not. The advantage to handling these cases separately is that we can obtain the negative result in Theorem 2 without much additional effort (see Corollary 1).

A.1.1 Non-Affine Statistics

Suppose there is a non-affine σ -statistic that is strongly elicitable, and let $(\mathbf{x}^n, \mathbf{t}^n)$ be a sequence of mechanisms that strongly identifies it. The following claim will be useful.

Claim 1. *For every $\epsilon > 0$, there exist constants $a, b \in \mathbb{R}$ where $a \leq b \leq a + \epsilon$ and*

$$\sigma\left(\frac{a+b}{2}\right) \neq \frac{\sigma(a) + \sigma(b)}{2} \quad (10)$$

We construct two joint distributions F and F' that the planner must be able to distinguish if the σ -statistic is identified. There are three steps to this construction.

1. Consider two basic instances I and I' where $\Omega = \{1, 2\}$ and $\mathcal{V} = \{a, b, (a+b)/2\}$. With probability one, the state is $\omega = 1$. Let the cost of the good signal given the good technology be $c(\theta_i^G, s_i^G) = \tilde{c} > 0$.
2. Construct a joint distribution F for basic instance I as follows. Let agent i have the bad technology with probability $p \in (0, 1)$. Let values be drawn according to

$$v_i \sim a + (b - a) \cdot \text{BERNOULLI}(0.5)$$

Note that the expected value v_i and σ -statistic are

$$E_F[v_i] = \frac{a+b}{2} \quad \text{and} \quad E_F[\sigma(v_i)] = \frac{\sigma(a) + \sigma(b)}{2}$$

3. Construct another joint distribution F' for basic instance I' as follows. Agents have the good technology with probability one. Let values be drawn according to

$$\Pr_{F'}\left[v_i = \frac{a+b}{2}\right] = p \quad \text{and} \quad \Pr[v_i = a] = \Pr[v_i = b] = \frac{1-p}{2}$$

Note that the expected value v_i and σ -statistic are

$$E_{F'}[v_i] = \frac{a+b}{2} \quad \text{and} \quad E_{F'}[\sigma(v_i)] = (1-p) \cdot \frac{\sigma(a) + \sigma(b)}{2} + p \cdot \sigma\left(\frac{a+b}{2}\right)$$

Critically, the σ -statistic takes on a different value when the joint distribution is F relative to when it is F' . It follows that a planner who can identify the σ -statistic can

also distinguish between joint distributions F and F' . Accordingly, our goal is to show that the planner cannot distinguish between joint distributions F and F' .

Lemma 1. *Set $\epsilon > 0$ sufficiently small. For any mechanism (\mathbf{x}, \mathbf{t}) , there exists a pair of strategy profiles (\mathbf{s}, \mathbf{m}) and $(\mathbf{s}', \mathbf{m}')$ where:*

1. *Given instance I , (\mathbf{s}, \mathbf{m}) is an equilibrium.*
2. *Given instance I' , $(\mathbf{s}', \mathbf{m}')$ is an equilibrium.*
3. *Consider message profiles m and m' , defined as*

$$\begin{aligned} m_i &= \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i)) \\ m'_i &= \mathbf{m}'_i(\mathbf{s}'_i, \mathbf{s}'_i(\omega, \theta_i, v_i)) \end{aligned}$$

The distribution of m given joint distribution F is the same as the distribution of m' given joint distribution F' .

Lemma 1 ensures that, with suitable equilibrium selection, the data has the same distribution regardless of whether the joint distribution is F or F' . In turn, the estimate Y is identically-distributed in both cases. This leads to a contradiction, as follows. Strong identification requires that

$$Y \rightarrow_p E_F[\sigma(v_i) \mid \omega] \quad \text{and} \quad Y \rightarrow_p E_{F'}[\sigma(v_i) \mid \omega] \quad (11)$$

However, by construction, the σ -statistic takes on a different value for distributions F and F' . That is,

$$E_F[\sigma(v_i) \mid \omega] \neq E_{F'}[\sigma(v_i) \mid \omega] \quad (12)$$

Together, conditions (11) and (12) imply that the estimate Y converges to two different quantities. This is a contradiction.

It also follows from this construction that non-affine statistics are not weakly elicitable. This is the second part of Theorem 2.

Corollary 1. *If σ is non-affine, then the σ -statistic is not weakly elicitable.*

Proof. The no-information benchmark for instances I and I' are

$$\text{Var}_F[E[\sigma(v_i) \mid \omega]] = \text{Var}_{F'}[E[\sigma(v_i) \mid \omega]] = 0$$

This follows from the fact that $\omega = 1$ with probability one. Therefore, satisfying the no-information benchmark means that the estimate must converge to the σ -statistic in every equilibrium, for both instance I and I' . We have shown that is not the case. \square

A.1.2 Affine Statistics

Suppose there is a non-constant affine statistic that is strongly elicitable, and let $(\mathbf{x}^n, \mathbf{t}^n)$ be a sequence of mechanisms that strongly identifies it. Since the statistic is non-constant, there must be some values $a, b \in \mathbb{R}$ such that $\sigma(a) \neq \sigma(b)$.

Consider a basic instance I'' with state space $\Omega = \{a, b\}$ and value space

$$\mathcal{V} = \{a, b, (a + b)/2\}$$

Let $p \in (0, 1)$ be the probability that agent i has the good technology. Let agents with the good technology have value $v_i = \text{UNIFORM}\{a, b\}$. Let agents with the bad technology have value $v_i = \omega$. Let the cost of the good signal given the bad technology be $\tilde{c} > \bar{v}$. Note that the σ -statistic has value

$$p \cdot \sigma\left(\frac{a + b}{2}\right) + (1 - p) \cdot \sigma(\omega)$$

Moreover, it varies with ω , since $\sigma(a) \neq \sigma(b)$.

Recall the basic instance I' from earlier in this proof outline. Redefine the parameters a, b, \tilde{c}, p to match those of the basic instance I'' . Note that the value of the σ -statistic does not vary with ω , since $\omega = 1$ with probability one.

Claim 2. *There exists an equilibrium $(\mathbf{s}'', \mathbf{m}'')$ of the mechanism for instance I'' where the value of the σ -statistic does not vary with ω .*

It follows from this claim that the σ -statistic is not strongly identified.

A.2 Special Cases of Theorem 2

We work through Theorem 2 in three examples, before outlining the general proof in Section A.3. These examples are intended to build intuition for how and when BDM-with-betting elicits (or weakly elicits) the average.

In these examples, we maintain three conventions. First, as in Section 2, we assume there are two alternatives: an outside option with value zero and a product

with value v_i to agent i . Second, we temporarily drop the requirement that the value space \mathcal{V} must be finite. Third, we always focus on the limit as both the sample size and the betting stakes grow (i.e., $n \rightarrow \infty$ and $\lambda_n \rightarrow \infty$).

A.2.1 Elicitation without an Informed Minority

The first example illustrates that there are natural instances, which violate the informed minority condition, where BDM-with-betting has a unique equilibrium. This stands in contrast to the example in Section 4.2, which had multiple equilibria.

Example 1. Fix a instance where the state represents product quality, i.e.,

$$\omega \sim \text{BERNOULLI}(0.5)$$

Each agent i 's value v_i is positively correlated with the state, i.e.,

$$v_i \sim N(\omega, 1)$$

Each agent i has a technology that represents a noisy assessment of the state, i.e.,

$$\theta_i \sim \text{BERNOULLI}\left(\frac{1 + \omega}{3}\right)$$

Each agent i has access to a revealing signal s_i^R that costs \bar{c} , where

$$s^R(\omega, \theta_i, v_i) = (v_i, \theta_i)$$

and an uninformative signal s_i^U that costs nothing and only reveals the technology, i.e.,

$$s_i^U(\omega, \theta_i, v_i) = \theta_i$$

The unique equilibrium that survives asymptotically sets

$$\forall i = 1, \dots, n : \quad \mathbf{s}_i(\theta_i, c_i) = s_i \quad \text{and} \quad \mathbf{m}_i(s_i(\omega, \theta_i, v_i), \theta_i, c_i) = v_i$$

That is, all agents i acquire the revealing signal and report their values v_i .

This equilibrium is unique because acquiring the revealing signal always helps agents predict the average reported value. To see this, consider two extreme cases.

First, if agents $j \neq i$ acquire the informative signal, the average reported value is

$$\tilde{v}_i = \frac{1}{n-1} \sum_{j \neq i} \hat{v}_j = \frac{1}{n-1} \sum_{j \neq i} v_j \rightarrow_p \omega$$

Second, if agents $j \neq i$ acquire the uninformative signal, the average reported value is

$$\tilde{v}_i = \frac{1}{n-1} \sum_{j \neq i} \hat{v}_j = \frac{1}{n-1} \sum_{j \neq i} E[v_j | \theta_j] \rightarrow_p \frac{4 + \omega}{9}$$

In both cases, the more that agent i learns about the state ω , the more accurately she is able to predict the average reported value \tilde{v}_i . Moreover, she can learn more about the state ω by acquiring her revealing signal; this reveals her value v_i , which is correlated with ω . As long as the scaling parameter λ_n is sufficiently large compared to the cost \bar{c} , agent i is better off acquiring the revealing signal.

In this example, the average reported value identifies the population's average value because every agent i learns her value v_i . Next, we consider an example where the same holds even though no agent learns her value v_i perfectly.

A.2.2 Elicitation with Partial Information Acquisition

The second example illustrates why BDM-with-betting remains effective if agents have access to multiple signals. In particular, we consider what happens if agents react to betting incentives by learning about others' values rather than their own.

Example 2. *The instance is the same as in Example 1, except that agents have access to two additional signals. First, there is a product quality signal s_i^Q that costs $\bar{c}/2$ and reveals the state ω , i.e.,*

$$s_i^Q(\omega, \theta_i, v_i) = (\omega, \theta_i)$$

Second, to satisfy Assumption 2, there is a combined signal s_i^C that costs $3\bar{c}/2$, i.e.,

$$s_i^C(\omega, \theta_i, v_i) = (\omega, \theta_i, v_i)$$

The unique equilibrium that survives asymptotically sets

$$\forall i = 1, \dots, n : \quad \mathbf{s}_i = s_i^Q \quad \text{and} \quad \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i)) = \omega$$

Here, agents i do not learn their own values v_i . They prefer the quality signal s_i^Q to the revealing signal s_i^R and the combined signal s_i^C . All three signals are equally good at predicting the vote margin, but the quality signal costs the least. Nonetheless, the average reported value equals the population's average value, since

$$\frac{1}{n} \sum_{j=1}^n \hat{v}_j = \frac{1}{n} \sum_{j=1}^n \mathbb{E}[v_j \mid \theta_j, \omega] = \omega = \mathbb{E}[v_i \mid \omega]$$

Examples 1 and 2 both have a symmetric pure-strategy equilibrium. Next, we consider an example where that is not the case.

A.2.3 Elicitation with Mixed-Strategy Equilibria

The third example illustrates how unstable best response dynamics can lead to mixed-strategy equilibria. The best response function is unstable in the sense that, as other agents j acquire more information, agent i wants to acquire *less* information. This is contrary to the intuition from the example in Section 4.2, where the incentive to acquire information is strongest when other agents also acquire information.

Example 3. *The instance is the same as in Example 1, with two changes. First, agent i 's value v_i is no longer correlated with the state, i.e.,*

$$v_{i1} \sim \text{UNIFORM}[-2, 2 + 2\gamma]$$

where $\gamma \in \mathbb{R}$ is a constant. Second, agent i 's technology θ_i provides information about both her value v_i and the state ω , i.e.,

$$\theta_i = \mathbf{1}(v_i > 1 - 2\omega)$$

Here, the state ω no longer represents quality. Instead, it controls whether the agents' uninformative signal is biased in favor of the product ($\omega = 1$) or not ($\omega = 0$). Each agent i either receives “good news” ($\theta_i = 1$), or “bad news” ($\theta_i = 0$).

In any symmetric pure strategy profile, each agent i prefers to become informed if and only if other agents are *not* informed. To see this, consider two cases.

1. If agents $j \neq i$ acquire the revealing signal s_j^R , the average reported value is

$$\tilde{v}_i = \frac{1}{n-1} \sum_{j \neq i} \hat{v}_j = \frac{1}{n-1} \sum_{j \neq i} v_j \rightarrow_p \gamma$$

Since the parameter γ is common knowledge, agent i does not need to acquire the revealing signal s_i^R in order to predict the vote margin. She prefers to deviate to the uninformative signal s_i^U , which costs less.

2. If agents $j \neq i$ acquire the uninformative signal s_j^U , the average reported value is

$$\tilde{v}_i = \frac{1}{n-1} \sum_{j \neq i} \hat{v}_j = \frac{1}{n-1} \sum_{j \neq i} E[v_j | \theta_j] \rightarrow_p \begin{cases} \frac{\gamma^2+2\gamma-2}{\gamma+4} & \omega = 0 \\ \frac{\gamma^2+6\gamma+2}{\gamma+4} & \omega = 1 \end{cases}$$

When λ_n is sufficiently large, agent i prefers to acquire the revealing signal, which is correlated with ω and therefore helps her predict the vote margin.

However, there is an equilibrium in mixed strategies. In that equilibrium, agents acquire the informative signal with high probability and the uninformative signal with low probability. Intuitively, the minority of uninformed agents is large enough to affect the average reported value (but just barely). The remaining agents are willing to acquire information to predict what this uninformed minority will report (but again, just barely). The size of the uninformed minority needed to incentivize information acquisition is decreasing as the betting stakes grow ($\lambda_n \rightarrow \infty$). Asymptotically, the uninformed minority vanishes as a fraction of the sample size, and the average reported value converges to the population's average value.

In this example, characterizing the mixed-strategy equilibrium is straightforward because there are only two signals, and they are common to all agents. When many signals are available, strategies are complicated and high-dimensional, which makes constructing mixed-strategy equilibria quite difficult. We overcome this difficulty with a non-constructive approach to proving Theorem 2.

A.3 Proof of Theorem 2

We outline the proof of Theorem 2. We focus on the positive result, since the negative result follows immediately from Corollary 1. After describing the sequence of mechanisms in more detail, we introduce key notation and sufficient conditions for weakly identifying the population's average value. Lemmas 2-4 verify those sufficient conditions. Corollary 2 handles the no-information benchmark.

Let each mechanism $(\mathbf{x}^n, \mathbf{t}^n)$ be a BDM-with-betting mechanism where the proper scoring rule is a simple extension of the continuous ranked probability score. More precisely, the scoring rule is

$$\text{SR}(b_i, \tilde{v}_i) = \bar{v}k - \sum_{x=1}^k \int_{[0, \bar{v}]} (\Pr_{b_i}[\tilde{v}_{ix} \leq z] - \mathbf{1}(\tilde{v}_{ix} \leq z))^2 dz$$

Note that we apply the continuous ranked probability score separately, for each alternative x , to the marginal distribution of the average reported value \tilde{v}_{ix} . The overall score is the sum of the scores associated with each alternative x .

A.3.1 Notation and Sufficient Conditions

We begin by introducing key notation and finding sufficient conditions for weakly identifying the population's average value.

Let agent i 's *error* be the difference between her true and expected values, i.e.,

$$\epsilon_i = v_i - \mathbb{E}[v_i \mid s_i, s_i(\omega, \theta_i, v_i)]$$

Since the BDM mechanism incentivizes truth-telling, her reported value is

$$\hat{v}_i = \mathbb{E}[v_i \mid s_i, s_i(\omega, \theta_i, v_i)] = v_i - \epsilon_i$$

The *average error* is the difference between the *average value* and the average reported value (across all agents). That is,

$$\underbrace{\frac{1}{n} \sum_{i=1}^n \epsilon_i}_{\text{average error}} = \underbrace{\frac{1}{n} \sum_{i=1}^n v_i}_{\text{average value}} - \underbrace{\frac{1}{n} \sum_{i=1}^n \hat{v}_i}_{\text{average reported value}}$$

If the average error is small, the alternative will be nearly optimal.

In the limit as the sample size n grows, the average error, average value, and average reported value converge in probability:

$$\frac{1}{n} \sum_{i=1}^n \epsilon_i \rightarrow_p E[\epsilon_i | \omega] \quad \frac{1}{n} \sum_{i=1}^n v_i \rightarrow_p E[v_i | \omega] \quad \frac{1}{n} \sum_{i=1}^n \hat{v}_i \rightarrow_p E[\hat{v}_i | \omega] \quad (13)$$

These averages converge in probability to expectations $E[\cdot | \omega]$ conditional on the realized state ω . This follows from the fact that the variables $\epsilon_i, v_i, \hat{v}_i$ are conditionally i.i.d., which allows us to invoke the law of large numbers after conditioning on ω .

For convenience, we often refer to the probability limits (13) as the average error, average value, and average reported value, respectively.

To show that the alternative is nearly optimal, it is sufficient to show that the average error vanishes, i.e.,

$$E[\epsilon_i | \omega] \rightarrow_p 0 \quad (14)$$

which ensures that the average reported value converges to the average value. But, as we saw in Example 2, it is not critical that the individual errors vanish. As long as the errors ϵ_i are uncorrelated across participants i , the average error will vanish. This is why BDM-with-betting works even though it does not directly incentivize participants to learn about their own values: it really only needs to incentivize participants i, j to acquire information up to the point where their errors ϵ_i, ϵ_j are uncorrelated.

To show that the average error vanishes (14), it suffices to show that

$$\text{Cov}[\epsilon_i, E[\hat{v}_i | \omega]] \rightarrow 0 \quad \text{and} \quad \text{Cov}[\epsilon_i, E[v_j | \omega]] \rightarrow 0 \quad (15)$$

Intuitively, if the error ϵ_i is neither correlated with the average reported value $E[\hat{v}_i | \omega]$, nor with the average value $E[v_j | \omega]$, then it cannot be correlated with the average error

$$E[\epsilon_j | \omega] = E[v_j | \omega] - E[\hat{v}_i | \omega]$$

If ϵ_i is not correlated with the average error, then it must not be correlated with the errors ϵ_j of other participants j . As we just discussed, this ensures that the average error vanishes.

A.3.2 Verifying Sufficient Conditions

We verify the sufficient conditions for eliciting the average, in three lemmas. First, we show that errors ϵ_i cannot be too correlated with the average reported value $E[\tilde{v}_i | \omega]$.

Lemma 2. *In any equilibrium of the mechanism $(\mathbf{x}_n, \mathbf{t}_n)$, the covariance between the error $\epsilon_{i,n}$ and conditional expectation $E[\tilde{v}_{i,n} | \omega]$ is bounded, i.e.,*

$$|\text{Cov}[\epsilon_{i,n}, E[\tilde{v}_{i,n} | \omega]]| = O\left(\frac{1}{\sqrt{\lambda_n}}\right)$$

Lemma 2 follows from the fact that participant i wants to acquire any signal that helps her better predict the average reported value $E[\hat{v}_j | \omega]$. Suppose that, for the sake of contradiction, her error ϵ_i is correlated with the average reported value. Then, if there were a signal that revealed her error ϵ_i , she would want to acquire it. However, this signal exists. By Assumptions 2 and 3, i can combine her revealing signal (which reveals v_i) with the signals that she has already acquired (which determine \hat{v}_i). This reveals her error $\epsilon_i = v_i - \hat{v}_i$ and costs at most \bar{c} more than the signals she already acquired. This combination is a profitable deviation whenever the scaling parameter λ_n is large relative to the cost \bar{c} . That, in turn, contradicts the premise that ϵ_i is correlated with the average reported value in equilibrium.

As a corollary of Lemma 2, we find that our BDM-with-betting mechanism satisfies the no-information benchmark. While this benchmark may seem trivial, simpler mechanisms like BDM may violate it (consider e.g., Example 3).

Corollary 2. *The mechanism $(\mathbf{x}_n, \mathbf{t}_n)$ satisfies the no-information benchmark.*

Next, we must show that ϵ_i is not too correlated with the average value $E[v_j | \omega]$. However, this is not true in every equilibrium. We must show that there exists an equilibrium in which ϵ_i is not too correlated with the average value.

We take an indirect approach. For every instance, we define an *auxiliary instance* where, by construction, ϵ_i is not correlated with the average value. Using Lemma 2, we show that the average error vanishes in every equilibrium of the auxiliary instance (Lemma 3). Then we show, when n is large, every equilibrium of the auxiliary instance is also an equilibrium of the original instance (Lemma 4).

Essentially, the auxiliary instance forces agents to only acquire signals that are maximally predictive of the conditional expectation $E[v_j | \omega]$.

Definition 17. For any given instance $I = (\Omega, \mathcal{V}, \Theta, \mathcal{S}, F, c)$, we define an auxiliary instance $\tilde{I} = (\Omega, \mathcal{V}, \Theta, \mathcal{S}, F, \tilde{c})$. There are three steps to this construction.

1. Let $\bar{s}_i \in \mathcal{S}$ be the signal that combines all other signals $s_i \in \mathcal{S}$ that have finite cost $c(s_i, \theta_i) < \infty$ for all technologies θ_i in the support of F .⁷
2. We can evaluate a given signal s_i by how well it predicts average value $E[v_j | \omega]$. Let predictiveness be the maximum expected score when agent i acquires s_i , i.e.,

$$\rho(s_i) = E \left[\max_{b_i} E \left[\text{SR}^{\text{CRPS}}(b_i, E[v_j | \omega]) \mid s_i(\omega, \theta_i, v_i) \right] \right] \quad (16)$$

3. For any given technology $\theta_i \in \Theta$, let every signal s_i that is less predictive than the combined signal \bar{s}_i have infinite cost according to the auxiliary cost function \tilde{c} . Let other signals cost the same as in the original cost function c . More precisely,

$$\tilde{c}(\theta_i, s_i) = \begin{cases} \infty & \rho(\theta_i, s_i) < \rho(\theta_i, \bar{s}_i) \\ c(\theta_i, s_i) & \text{otherwise} \end{cases} \quad (17)$$

In Lemma 3, we show that the average error vanishes in every equilibrium of the auxiliary instance.

Lemma 3. Fix the auxiliary instance. For any sequence of equilibria $(\mathbf{s}^n, \mathbf{m}^n)$ of the mechanisms $(\mathbf{x}^n, \mathbf{t}^n)$, the average reported value converges in probability to the population's average value.

The proof of Lemma 3 follows the logic of condition (15). Lemma 2 already that the error ϵ_i cannot be too correlated with the average reported value $E[\hat{v}_j | \omega]$. All that remains, by condition (15), is to show that the error ϵ_i cannot be correlated with the average value $E[v_j | \omega]$. This is because the only signals s_i with finite cost in the auxiliary instance are those that are maximally predictive of the average value. Suppose, for the sake of contradiction, that ϵ_i is correlated with the average value. Then there exists another signal – which combines s_i with the revealing signal, and reveals ϵ_i – that predicts the average value better than s_i . This contradicts the premise that s_i is maximally predictive.

⁷If signal s_i is itself a combined signal (Assumption 2), let \bar{s}_i include all base signals that s_i combines.

At this point, we have shown that the average error vanishes for all auxiliary instances. Lemma 4 extends this result to general instances.

Lemma 4. *There exists a constant N such that, for any sample size $n \geq N$, every equilibrium of the auxiliary instance is also an equilibrium of the original instance.*

To prove Lemma 4, it is enough to show that, in the candidate equilibrium of the original instance, no participant i will deviate to a signal that is unavailable in the auxiliary instance. That is, we want to show that participant i wants to acquire only signals that are maximally predictive of the average value $E[v_j | \omega]$. We know, in the limit, that she will acquire only signals that are maximally predictive of the average reported value $E[\hat{v}_j | \omega]$. Moreover, by Lemma 3, in any equilibrium of the auxiliary instance, the average reported value converges to the average value (as the average error vanishes). Therefore, she will only acquire signals that are maximally predictive of the average value.

This completes the proof of Theorem 1. The proofs of Lemmas 2-4 are somewhat involved, so we leave them to the Supplemental Appendix.

A.4 Proof of Theorem 3 (Special Case)

We outline the proof of Theorem 3 when Assumption 5 holds. We construct a sequence of majority-rule-with-betting mechanisms $(\mathbf{x}^n, \mathbf{t}^n)$ with proper scoring rule SR, scaling parameter λ_n and probability parameter δ_n . Then we show that it provides incentives for truthful reporting and efficient information acquisition.

First, we define the proper scoring rule. Without loss of generality, let the sample size n be odd. Let $Q_{i,n}$ be the event that agent i is pivotal; that is, either i is dictator or

$$(n-1) \cdot \tilde{n}_{i,n} = \sum_{j \neq i} \hat{x}_{i,n} = \frac{n-1}{2}$$

Let $q_{i,n}$ be the probability of this conditional on i 's information. Let SR be the sum of a quadratic scoring rule and a continuous ranked probability score. That is,

$$\text{SR}(\hat{b}_i, \tilde{v}_{i,n}) = \text{SR}_n^Q(\hat{b}_{i,n}, \tilde{v}_{i,n}) + \text{SR}^{CRPS}(\hat{b}_{i,n}, \tilde{v}_{i,n})$$

where SR^{CRPS} is defined as in Appendix A.3 (replacing $\tilde{v}_{i,n}$ with $\tilde{n}_{i,n}$). To define SR_n^Q , let

$$\text{SR}_n^Q(\hat{b}_i, \tilde{v}_{i,n}) = 2 \cdot \mathbf{1}(Q_{i,n}) \cdot \hat{q}_{i,n} + 2 \cdot \mathbf{1}(\neg Q_{i,n}) \cdot (1 - \hat{q}_{i,n}) - \hat{q}_{i,n}^2 - (1 - \hat{q}_{i,n})^2$$

Each of these scoring rules is convenient in different parts of the proof. Since each one is proper, the sum SR is also proper.

We do not attempt to optimize tuning parameters. Generally, it is important that $\delta_n \rightarrow 0$ and $\lambda_n \rightarrow \infty$. When δ_n vanishes quickly, agents are unlikely to be pivotal, and λ_n must grow even more quickly in order to ensure near-truthfulness. For the sake of concreteness, set

$$\delta_n = n^{-1/2} \quad \text{and} \quad \lambda_n = n^7$$

The proof of Lemma 5 describes trade-offs between these parameters in more detail.

Next, we show that these majority-rule-with-betting mechanisms are *nearly truthful*. This means that each agent's reported alternative \hat{x}_i probably approximately maximizes her expected value conditional on the information she acquires.

Definition 18. *The majority-rule-with-betting mechanism $(\mathbf{x}^n, \mathbf{t}^n)$ is (ϕ_1, ϕ_2) -truthful if the following holds for every instance $I \in \mathcal{I}$ and equilibrium $(\mathbf{s}, \hat{\mathbf{x}}, \hat{\mathbf{b}})$. Let agent i report alternative $\hat{x}_i = \hat{\mathbf{x}}_i(\mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i))$. Then*

$$\Pr \left[\max_{i=1, \dots, n} \left| \max_{x \in \mathcal{X}} \mathbb{E}[v_{ix} \mid \mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i)] - \mathbb{E}[v_{i\hat{x}_i} \mid \mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i)] \right| > \phi_1 \right] \leq \phi_2$$

Lemma 5 formalizes the high-level intuition in Section 5.5.

Lemma 5. *The mechanism $(\mathbf{x}^n, \mathbf{t}^n)$ is (ϕ_{1n}, ϕ_{2n}) -truthful, where $\phi_{1n} \rightarrow 0$ and $\phi_{2n} \rightarrow 0$.*

We also prove a result analogous to Lemma 2. This says that agent i 's error ϵ_i is uncorrelated with the vote share $\tilde{n}_{i,n}$.

Lemma 6. *In any strategy profile of the mechanism $(\mathbf{x}_n, \mathbf{t}_n)$ where agent i does not have a profitable deviation,*

$$|\text{Cov}[\epsilon_{i,n}, \mathbb{E}[\tilde{n}_{i,n} \mid \omega]]| = O(\lambda_n^{-1/2})$$

So far, we have not relied on Assumption 5. In particular, even if that assumption fails, majority-rule-with-betting would alleviate the swing voter's curse.

However, we do need Assumption 5 to show that efficient equilibria exist. Let $\mu = \mathbb{E}[v_i]$. Without loss of generality, let $\mu \neq 0$. There are two cases to consider.

1. Suppose that

$$E[(v_i - \mu) \cdot \Pr[v_j = 1 \mid \omega]] > 0$$

Then there exists some number N such that, for any sample size $n \geq N$, there is an equilibrium where every agent i acquires the revealing signal and reports $\hat{x}_i = \mathbf{1}(v_i = 1)$. To see this, note that the only deviation for agent i that might be profitable involves acquiring the uninformative signal. Then $\epsilon_i = v_i - \mu$, and

$$E[\epsilon_i \cdot E[\tilde{n}_{i,n} \mid \omega]] = E[\epsilon_i \cdot \Pr[v_j = 1 \mid \omega]] = E[(v_i - \mu) \cdot \Pr[v_j = 1 \mid \omega]] > 0$$

Since $\lambda_n \rightarrow \infty$, this eventually contradicts Lemma 6.

Finally, observe that the vote share converges to $\Pr[v_i = 1 \mid \omega]$, which exceeds 50% if and only if $E[v_i \mid \omega] \geq 0$. In the limit, since $\delta_n \rightarrow 0$, alternative $x = 1$ is chosen if and only if it maximizes welfare.

2. Suppose that

$$E[(v_i - \mu) \cdot \Pr[v_j = 0 \mid \omega]] = 0 \tag{18}$$

Observe that

$$\begin{aligned} E[v_i \mid \omega] &= \Pr[v_i = 1 \mid \omega] - \Pr[v_i = -1 \mid \omega] \\ &= \Pr[v_i = 1 \mid \omega] - 1 + \Pr[v_i = 1 \mid \omega] \\ &= 2\Pr[v_i = 1 \mid \omega] - 1 \end{aligned}$$

Plugging this back into inequality (18) gives us

$$E\left[(v_i - \mu) \cdot \frac{E[v_i \mid \omega] + 1}{2}\right] = 0$$

After simplifying this expression, we obtain

$$\text{Var}[E[v_i \mid \omega]] = 0$$

Therefore, $E[v_i \mid \omega] = \mu$ with probability one. Furthermore, the outcome $x = 1$ is efficient if $\mu > 0$ and inefficient if $\mu < 0$.

Now, consider a strategy profile where every agent i acquires the uninformative signal and optimally reports $\hat{x}_i = \mathbf{1}(\mu = 1)$. Since the vote share $\tilde{n}_{i,n}$ does not

depend on the state ω , acquiring the revealing signal is not a profitable deviation. Therefore, this is an equilibrium. In this equilibrium, the chosen alternative is $x = \mathbf{1}(\mu > 0)$. As we just argued, this is efficient.

In both cases, an efficient equilibrium exists.

All that remains is to verify that, in all equilibria, welfare is at least as high as if the planner had chosen the ex-ante optimal alternative. Let n be large. Continue to assume, without loss of generality, that $\mu \neq 0$. In a given equilibrium, let the random set N include all agents that acquire the revealing signal. All agents $i \in N$ report $\hat{x}_i = \mathbf{1}(v_i = 1)$. If this group of agents prevails in the vote, then the ex-post optimal alternative is chosen with high probability. It follows from Lemma 5 that, with high probability, every agent $i \notin N$ reports $\hat{x}_i = \mathbf{1}(\mu > 0)$. If this group of agents prevails in the vote, then the ex-ante optimal alternative is chosen. Either way, the chosen alternative is at least as good as the ex-ante optimal alternative.

References

- Ali, S. N., Mihm, M., & Siga, L. (2025). The political economy of zero-sum thinking. *Econometrica*, 93(1), 41–70.
- Angelucci, C., & Prat, A. (2024). Is journalistic truth dead? measuring how informed voters are about political news. *American Economic Review*, 114(4), 887–925.
- Arrow, K. J. (1950). A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4), 328–346.
- Austen-Smith, D., & Banks, J. S. (1996). Information aggregation, rationality, and the condorcet jury theorem. *The American Political Science Review*, 90(1), 34–45.
- Bardhi, A., & Bobkova, N. (2023). Local evidence and diversity in minipublics. *Journal of Political Economy*, 131(9), 2451–2508.
- Barseghyan, L., Coughlin, M., Molinari, F., & Teitelbaum, J. C. (2021a). Heterogeneous choice sets and preferences. *Econometrica*, 89(5), 2015–2048.
- Barseghyan, L., Molinari, F., & Thirkettle, M. (2021b). Discrete choice under risk with limited consideration. *American Economic Review*, 111(6), 1972–2006.
- Bartels, L. M. (1996). Uninformed votes: Information effects in presidential elections. *American Journal of Political Science*, 40(1), 194–230.

- Battaglini, M., Morton, R. B., & Palfrey, T. R. (2010). The swing voter's curse in the laboratory. *The Review of Economic Studies*, 77(1), 61–89.
- Becker, G. M., Degroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3), 226–232.
- Ben-Porath, E., Dekel, E., & Lipman, B. L. (2024). *Sequential mechanisms for evidence acquisition*.
- Bergemann, D., Shi, X., & Välimäki, J. (2009). Information Acquisition in Interdependent Value Auctions. *Journal of the European Economic Association*, 7(1), 61–89.
- Bergemann, D., & Välimäki, J. (2002). Information acquisition and efficient mechanism design. *Econometrica*, 70(3), 1007–1033.
- Bhattacharya, S. (2013). Preference monotonicity and information aggregation in elections. *Econometrica*, 81(3), 1229–1247.
- Bikhchandani, S. (2010). Information acquisition and full surplus extraction. *Journal of Economic Theory*, 145(6), 2282–2308.
- Brav, A., Malenko, A., & Malenko, N. (2023). Corporate governance implications of the growth in indexing.
- Cai, H. (2009). Costly participation and heterogeneous preferences in informational committees. *The RAND Journal of Economics*, 40(1), 173–189.
- Cao, X., & Zhang, J. (2021). Preference learning and demand forecast. *Marketing Science*, 40(1), 62–79.
- Carroll, G. (2015). Robustness and linear contracts. *American Economic Review*, 105(2), 536–63.
- Carroll, G. (2019). Robust incentives for information acquisition. *Journal of Economic Theory*, 181, 382–420.
- Cattaneo, M. D., Ma, X., Masatlioglu, Y., & Suleymanov, E. (2020). A random attention model. *Journal of Political Economy*, 128(7), 2796–2836.
- Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11, 17–33.
- Coppock, A. (2023). *Persuasion in parallel: How information changes minds about politics*. University of Chicago Press.
- Crémer, J., & McLean, R. P. (1985). Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2), 345–361.

- Crémer, J., & McLean, R. P. (1988). Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions. *Econometrica*, 56(6), 1247–1257.
- Dahl, R. A. (1989). *Democracy and its critics*. Yale University Press.
- Dasgupta, A., & Ghosh, A. (2013). Crowdsourced judgement elicitation with endogenous proficiency. *Proceedings of the 22nd International Conference on World Wide Web*, 319–330.
- Delli Carpini, M. X., & Keeter, S. (1996). *What americans know about politics and why it matters*. Yale University Press.
- Denti, T., & Ravid, D. (2024). Robust predictions in games with rational inattention. *Proceedings of the 25th ACM Conference on Economics and Computation*, 573.
- Devanur, N., Hartline, J., Karlin, A., & Nguyen, T. (2011). Prior-independent multi-parameter mechanism design. In N. Chen, E. Elkind, & E. Koutsoupias (Eds.), *Internet and network economics* (pp. 122–133). Springer Berlin Heidelberg.
- Feddersen, T. J., & Pesendorfer, W. (1996). The swing voter’s curse. *The American Economic Review*, 86(3), 408–424.
- Feddersen, T. J., & Sandroni, A. (2006). Ethical voters and costly information acquisition. *Quarterly Journal of Political Science*, 1(3), 287–311.
- Fishkin, J. S. (2011). *When the people speak: Deliberative democracy and public consultation*. Oxford University Press.
- Flanigan, B., Gözl, P., Gupta, A., Hennig, B., & Procaccia, A. D. (2021). Fair algorithms for selecting citizens’ assemblies. *Nature*, 596(7873), 548–552.
- Fowler, A., & Margolis, M. (2014). The political consequences of uninformed voters. *Electoral Studies*, 34, 100–110.
- Fu, H., Haghpanah, N., Hartline, J., & Kleinberg, R. (2021). Full surplus extraction from samples. *Journal of Economic Theory*, 193, 105230.
- Gao, X. A., Mao, A., Chen, Y., & Adams, R. P. (2014). Trick or treat: Putting peer prediction to the test. *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, 507–524.
- Gao, X. A., Wright, J. R., & Leyton-Brown, K. (2019). Incentivizing evaluation with peer prediction and limited access to ground truth. *Artificial Intelligence*, 275, 618–638.
- Gerardi, D., & Yariv, L. (2008). Information acquisition in committees. *Games and Economic Behavior*, 62(2), 436–459.

- Gershkov, A., & Szentes, B. (2009). Optimal voting schemes with costly information acquisition. *Journal of Economic Theory*, 144(1), 36–68.
- Gilens, M. (2001). Political ignorance and collective policy preferences. *American Political Science Review*, 95(2), 379–396.
- Gleyze, S., & Pernoud, A. (2023). Informationally simple incentives. *Journal of Political Economy*, 131(3), 802–837.
- Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review*, 70(3), 393–408.
- Groves, T., & Ledyard, J. (1977). Optimal allocation of public goods: A solution to the “free rider” problem. *Econometrica*, 45(4), 783–809.
- Hanson, R. (2003). Combinatorial information market design. *Information Systems Frontiers*, 5(1), 107–119.
- Immorlica, N., Leshno, J. D., Lo, I. Y., & Lucier, B. J. (2020). *Information acquisition in matching markets: The role of price discovery*.
- Kastiel, K., & Nili, Y. (2016). In search of the absent shareholders: A new solution to retail investors’ apathy. *Delaware Journal of Corporate Law*, 41, 55.
- Kim, J. [Jaehoon], & Fey, M. (2007). The swing voter’s curse with adversarial preferences. *Journal of Economic Theory*, 135(1), 236–252.
- Kong, Y., & Schoenebeck, G. (2019). An information theoretic framework for designing information elicitation mechanisms that reward truth-telling. *ACM Trans. Econ. Comput.*, 7(1).
- Lambert, N. S., Pennock, D. M., & Shoham, Y. (2008). Eliciting properties of probability distributions. *Proceedings of the 9th ACM Conference on Electronic Commerce*, 129–138.
- Larcinese, V. (2009). Information acquisition, ideology and turnout: Theory and evidence from Britain. *Journal of Theoretical Politics*, 21(2), 237–276.
- Lau, R. R., & Redlawsk, D. P. (1997). Voting correctly. *The American Political Science Review*, 91(3), 585–598.
- Li, Y., Hartline, J. D., Shan, L., & Wu, Y. (2022). Optimization of scoring rules. *Proceedings of the 23rd ACM Conference on Economics and Computation*, 988–989.
- Mamadehussene, S., & Sguera, F. (2023). On the reliability of the bdm mechanism. *Management Science*, 69(2), 1166–1179.

- Manzini, P., & Mariotti, M. (2014). Stochastic choice and consideration sets. *Econometrica*, 82(3), 1153–1176.
- Martinelli, C. (2006). Would rational voters acquire costly information? *Journal of Economic Theory*, 129(1), 225–251.
- Masatlioglu, Y., Nakajima, D., & Ozbay, E. Y. (2012). Revealed attention. *American Economic Review*, 102(5), 2183–2205.
- McCarthy, J. (1956). Measures of the value of information. *Proceedings of the National Academy of Sciences*, 42(9), 654–655.
- Miller, N., Resnick, P., & Zeckhauser, R. (2005). Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9), 1359–1373.
- Mulligan, C. B., & Hunter, C. G. (2003). The empirical frequency of a pivotal vote. *Public Choice*, 116(1/2), 31–54.
- Pakzad-Hurson, B. (2022). Crowdsourcing and optimal market design.
- Palfrey, T. R., & Poole, K. T. (1987). The relationship between information, ideology, and voting behavior. *American Journal of Political Science*, 31(3), 511–530.
- Pande, R. (2011). Can informed voters enforce better governance? experiments in low-income democracies. *Annual Review of Economics*, 3(1), 215–237.
- Pei, H., & Strulovici, B. (2024). Robust implementation with costly information. *The Review of Economic Studies*, 92(1), 476–505.
- Pernoud, A., & Gleyze, S. (2024). *How competition shapes information in auctions*.
- Persico, N. (2000). Information acquisition in auctions. *Econometrica*, 68(1), 135–148.
- Persico, N. (2004). Committee Design with Endogenous Information. *The Review of Economic Studies*, 71(1), 165–191.
- Prelec, D. (2004). A bayesian truth serum for subjective data. *Science*, 306(5695), 462–466.
- Prior, M., & Lupia, A. (2008). Money, time, and political knowledge: Distinguishing quick recall and political learning skills. *American Journal of Political Science*, 52(1), 169–183.
- Rothkopf, M. H. (2007). Thirteen reasons why the vickrey-clarke-groves process is not practical. *Oper. Res.*, 55(2), 191–197.
- Schoenebeck, G., & Tao, B. (2021). Wisdom of the crowd voting: Truthful aggregation of voter information and preferences. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, & J. W. Vaughan (Eds.), *Advances in neural information processing systems* (pp. 1872–1883, Vol. 34). Curran Associates, Inc.

- Shineman, V. A. (2018). If you mobilize them, they will become informed: Experimental evidence that information acquisition is endogenous to costs and incentives to participate. *British Journal of Political Science*, 48(1), 189–211.
- Shnayder, V., Agarwal, A., Frongillo, R., & Parkes, D. C. (2016). Informed truthfulness in multi-task peer prediction. *Proceedings of the 2016 ACM Conference on Economics and Computation*, 179–196.
- Tyson, S. A. (2016). Information acquisition, strategic voting, and improving the quality of democratic choice. *The Journal of Politics*, 78(4), 1016–1031.

B Supplemental Appendix

We maintain the following notation. In a strategy profile (s^n, m^n) where agent i acquires signal s_i , let

$$\Pr_{i,n}[\cdot] = \Pr[\cdot \mid s_i, s_i(\omega, \theta_i, v_i)]$$

That is, $\Pr_{i,n}[\cdot]$ is the probability conditioned on agent i 's information. Define the expected value $E_{i,n}[\cdot]$, variance $\text{Var}_{i,n}[\cdot]$, and covariance $\text{Cov}_{i,n}[\cdot]$ similarly.

B.1 Proof of Theorem 3 (General Case)

We outline a proof of Theorem 3 that does not rely on Assumption 5. The high-level intuition from Section 5.5 still applies, but now we must formalize that intuition in a significantly more general setting and using a new mechanism.

B.1.1 VCG-with-Betting Mechanism

We define a new mechanism called *VCG-with-betting*. This mechanism is analogous to BDM-with-betting and majority-rule-with-betting. The difference is that, in the second stage, agents report their willingness to pay to the VCG mechanism. As with majority-rule-with-betting, we also need to add some noise to the mechanism to avoid issues that arise when agents condition on low-probability events.

The VCG-with-betting mechanism has four tuning parameters. First, there is proper scoring rule SR, which evaluates the accuracy of agent i 's reported beliefs \hat{b}_i over the average reported value \tilde{v}_i . Second, there is the scaling parameter λ that controls how large the betting stakes are. Third, there is a probability δ_n that some agent i is chosen as dictator. Fourth, there is a randomized *bias term* $v_{0,n} \in \mathcal{V}$. This biases the mechanism towards alternatives x where v_{0x} is large. Let the average reported value \tilde{v}_i include the bias term, i.e.,

$$\tilde{v}_i = \frac{1}{n} \left(v_{0,n} + \sum_{j \neq i} \hat{v}_j \right)$$

The agents do not know the realization of the bias term.

Definition 19. *The VCG-with-betting mechanism (\mathbf{x}, \mathbf{t}) is parameterized by a proper scoring rule SR, scaling parameter λ , probability δ , and randomized bias term $v_{0,n}$. Each*

agent i sends a message

$$m_i = (\hat{v}_i, \hat{b}_i) \in [v_L, v_H] \times \mathcal{B}$$

that consists of a reported value \hat{v}_i and a reported belief \hat{b}_i . There are two cases.

1. With probability $1 - \delta$, there is no dictator ($D = 0$). The planner selects the alternative that maximizes the average reported value plus the bias term, i.e.,

$$\mathbf{x}(\hat{v}, \hat{b}) \in \arg \max_{x \in \mathcal{X}} \left(v_{0x} + \sum_{i=1}^n \hat{v}_{ix} \right)$$

Each agent i is paid according to the VCG mechanism's transfer rule and the proper scoring rule, i.e.,

$$\mathbf{t}_i(\hat{v}, \hat{b}) = v_{0x(\hat{v}, \hat{b})} + \sum_{j \neq i} \hat{v}_{jx(\hat{v}, \hat{b})} - \max_{x \in \mathcal{X}} \left(v_{0x} + \sum_{j \neq i} \hat{v}_{jx} \right) + \lambda \cdot \text{SR}(\hat{b}_i, \tilde{v}_i)$$

2. With probability δ_n , the planner selects a random agent as dictator, i.e.,

$$D \sim \text{UNIFORM}(1, \dots, n)$$

This agent is paid according to a BDM mechanism with random price $p \sim \text{UNIFORM}([v_L, v_H])$ and the proper scoring rule, i.e.,

$$\mathbf{t}_D(\hat{v}, \hat{b}) = \lambda \cdot \text{SR}(\hat{b}_D, \tilde{v}_D) - p \cdot \mathbf{1}(\hat{v}_D \geq p)$$

The other agents $j \neq D$ are according to the proper scoring rule, i.e.,

$$\mathbf{t}_j(\hat{v}, \hat{b}) = \lambda \cdot \text{SR}(\hat{b}_j, \tilde{v}_j)$$

Finally, the allocation is

$$\mathbf{x}_i(\hat{v}, \hat{b}) = \mathbf{1}(\hat{v}_D \geq p)$$

B.1.2 Sequence of VCG-with-Betting Mechanisms

We construct a sequence $(\mathbf{x}^n, \mathbf{t}^n)$ of VCG-with-betting mechanisms.

First, we introduce useful notation. Let $\tilde{p}_{i,n}$ be the effective price of switching

from alternative 0 to 1, i.e.,

$$\tilde{p}_{i,n} = \mathbf{1}(D = i) \cdot p + \mathbf{1}(D = 0) \cdot (-n\tilde{v}_{i,n}) + \mathbf{1}(D \notin \{0, i\}) \cdot \infty$$

Let $Q_{i,n}$ indicate the event that agent i is potentially pivotal, i.e.,

$$v_L \leq \tilde{p}_{i,n} \leq v_H$$

Let $q_{i,n}$ be the probability that agent i is potentially pivotal, i.e.,

$$q_{i,n} = \Pr_{i,n}[Q_{i,n}]$$

Let $\hat{q}_{i,n}$ be the probability according to reported belief $\hat{b}_{i,n}$.

Second, we specify the scoring rule SR. Let SR be the sum of the quadratic scoring rule and two continuous ranked probability scores. That is,

$$\text{SR}(\hat{b}_i, \tilde{v}_{i,n}) = \text{SR}_n^Q(\hat{b}_{i,n}, \tilde{v}_{i,n}) + \text{SR}^{CRPS}(\hat{b}_{i,n}, \tilde{v}_{i,n}) + \text{SR}^{CRPS}(\hat{b}_{i,n}, \mathbf{1}(\tilde{v}_{i,n} \geq 0))$$

where SR^{CRPS} is defined as in Appendix A.3, and SR^Q is defined as in Appendix A.4 (replacing $\tilde{n}_{i,n}$ with $\tilde{v}_{i,n}$). Each of these scoring rules is convenient in different parts of the proof. Since each one is proper, the sum SR is also proper.

Third, we specify the distribution of the bias term. Let v_{0x} follow a Laplace distribution with scaling parameter $\beta_n > 0$, i.e.,

$$v_{0x} \sim \text{LAPLACE}(0, \beta_n) \quad \text{i.i.d. across } x \in \mathcal{X}$$

The Laplace distribution ensures that agent i 's expected value conditional on the vote margin $\tilde{v}_{i,n}$ is a smooth function of $\tilde{v}_{i,n}$.

We do not attempt to optimize these tuning parameters. Generally, it is important that $\delta_n \rightarrow 0$, $\beta_n \rightarrow \infty$, and $\beta_n/n \rightarrow 0$. It is also important that $\lambda_n \rightarrow \infty$. When δ_n vanishes quickly, agents are unlikely to be pivotal, and λ_n must grow even more quickly in order to ensure near-truthfulness. For the sake of concreteness, set

$$\beta_n = 2e(v_H - v_L) \cdot n^{1/2}, \quad \delta_n = n^{-1/2}, \quad \lambda_n = n^7$$

The proof of Lemma 7 describes trade-offs between these parameters in more detail.

B.1.3 Convergence to Efficiency

We begin by formalizing what it means for VCG-with-betting to be nearly truthful. Essentially, each agent's reported values should be probably approximately equal to her expected values (conditional on any information she has acquired).

Definition 20. *The VCG-with-betting mechanism $(\mathbf{x}^n, \mathbf{t}^n)$ is (ϕ_1, ϕ_2) -truthful if the following holds for every instance $I \in \mathcal{I}$ and equilibrium $(\mathbf{s}, \hat{\mathbf{v}}, \hat{\mathbf{b}})$. Let agent i acquire signal $s_i = \mathbf{s}_i(\theta_i, c_i)$ and report values $\hat{v}_i = \hat{\mathbf{v}}_i(s_i, s_i(\omega, \theta_i, v_i))$. Then*

$$\Pr \left[\max_{i=1, \dots, n} |\hat{v}_i - \mathbb{E}[v_i \mid s_i, s_i(\omega, \theta_i, v_i)]| > \phi_1 \right] \leq \phi_2$$

Next, we verify that our VCG-with-betting mechanism is nearly-truthful. This is the key step in proving Theorem 3.

Lemma 7. *The mechanism $(\mathbf{x}^n, \mathbf{t}^n)$ is (ϕ_{1n}, ϕ_{2n}) -truthful, where $\phi_{1n} \rightarrow 0$ and $\phi_{2n} \rightarrow 0$.*

The next three lemmas are analogs Lemmas 2-4. The main differences are that, rather than require truthful reporting, they only rely on nearly-truthful reporting.

Lemma 8. *In any equilibrium of the mechanism $(\mathbf{x}_n, \mathbf{t}_n)$,*

$$|\text{Cov}[\epsilon_{i,n}, \mathbb{E}[\tilde{v}_{i,n} \mid \omega]]| = O\left(\frac{1}{\sqrt{\lambda_n}}\right) \quad \text{and} \quad |\text{Cov}[\epsilon_{i,n}, \mathbf{1}(\tilde{v}_{i,n} \geq 0)]| = O\left(\frac{1}{\sqrt{\lambda_n}}\right)$$

Lemma 9. *Fix the auxiliary instance. For any sequence of equilibria $(\mathbf{s}^n, \mathbf{m}^n)$,*

$$\frac{1}{n+1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \right) \rightarrow_p \mathbb{E}[v_j \mid \omega] \quad (19)$$

Lemma 10. *There exists a constant N such that, for any sample size $n \geq N$, every equilibrium of the auxiliary instance is also an equilibrium of the original instance.*

Lemmas 9 and 10 guarantee an equilibrium $(\mathbf{s}, \hat{\mathbf{v}}^n, \hat{\mathbf{b}}^n)$ where condition (19) holds. It follows from the argmax continuous mapping theorem (Jeankyung Kim and Pollard 1990) that the alternative $\mathbf{x}^n(\hat{\mathbf{v}}^n, \hat{\mathbf{b}}^n)$ converges into the set of welfare-maximizing alternatives.

At this point, we have shown that there are good equilibria. Finally, Lemma 11 rules out unreasonably bad equilibria. It relies critically on Lemmas 7 and 8.

Lemma 11. *The mechanism $(\mathbf{x}^n, \mathbf{t}^n)$ performs at least as well as the ex-ante optimal alternative, in the limit as $n \rightarrow \infty$ and for every equilibrium. Formally, for every $I \in \mathcal{J}$,*

$$\limsup_{n \rightarrow \infty} \bar{\alpha}^S(\mathbf{x}^n, \mathbf{t}^n) \leq \text{OPT}(I) - \max_x \mathbb{E}[v_{ix}]$$

B.2 Proof of Claim 1

First, suppose that this condition fails and σ is continuous. Consider a small interval $[c, d] \subseteq \mathcal{V}$, where “small” means that $d - c \leq \epsilon$. By definition, since we cannot find constants $a, b \in [c, d]$ such that inequality (10) holds, σ is both midpoint convex and midpoint concave on the interval $[c, d]$. By Jensen (1905), it follows that σ is both convex and concave – in other words, affine – on the interval $[c, d]$. If σ is affine on every small subinterval $[c, d]$, it must (i) be differentiable and (ii) have a constant derivative on the entire domain \mathcal{V} . Therefore, it is affine.

Second, suppose that σ is discontinuous at $d \in \mathbb{R}$. By Definition 6, σ has a finite discontinuity set. So, there are intervals $[d', d]$ and $(d, d'']$ over which σ is continuous. Let $a \in [d', d]$ and $b \in [d, d'']$. As $d', d'' \rightarrow d$, either (i) $\lim \sigma(d') \neq \sigma(d)$, (ii) $\lim \sigma(d'') \neq \sigma(d)$, or both. If (i), let $a = d'$ and $b = d$. If (ii), let $b = d''$ and $a = d$. As $d', d'' \rightarrow d$, eventually $b - a \leq \epsilon$ and inequality (10) holds.

B.3 Proof of Lemma 1

We first discuss strategy profile (\mathbf{s}, \mathbf{m}) , and then turn to strategy profile $(\mathbf{s}', \mathbf{m}')$.

Let (\mathbf{s}, \mathbf{m}) be some equilibrium when the instance is I . Let $s_i = \mathbf{s}_i$ and

$$m_i = \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i))$$

I claim that the vector $(v_i, \theta_i, s_i, m_i)$ is independent across agents i . This follows from three observations. First, the values v_i and technologies θ_i are independent across agents i after conditioning on the state ω . Second, the state is $\omega = 1$ with probability one, so values v_i and technologies θ_i are unconditionally independent across agents i . Third, s_i and m_i are random functions of (ω, θ_i, v_i) . The arguments are independent across agents (since $\omega = 1$ with probability one) and the mixing in mixed strategies is independent across agents.

Consider the distribution of acquired signals s_i . We claim that agent i learns her

value v_i with probability $1 - p$. We begin by making some observations. The good signal is only useful to agent i insofar as it affects her choice of which message m_i to send. With probability $1 - p$, agent i has the good technology and learns her value v_i regardless of which signal she acquires. Conditional on having the bad technology, the cost of the good signal is $\tilde{c} > 0$ and the agent will only acquire it if she benefits from the information that it provides. In this case, the state is fixed with probability one, so the information is her value v_i . However, this information does not affect her beliefs about other agents' messages, since v_i and m_{-i} are independent.

Suppose that agent i acquires the good signal. Her expected cost of doing so is $p\tilde{c}$. We want to find a profitable deviation. Consider her expected utility

$$U_i(\mathbf{s}, \mathbf{m}, \mathbf{x}, \mathbf{t}) = E[u_i(v_i, \mathbf{x}_i(m), \mathbf{t}_i(m), s_i, \theta_i)]$$

Recall that her value is $v_i \in \{a, b\}$. Suppose agent i sends message m_i^a if she learns $v_i = a$ and m_i^b if she learns $v_i = b$. Then

$$\begin{aligned} & E[u_i(a, \mathbf{x}_i(m_i^b, m_{-i}), \mathbf{t}_i(m_i^b, m_{-i}), s_i, \theta_i, c_i)] \\ & \geq E[u_i(b, \mathbf{x}_i(m_i^b, m_{-i}), \mathbf{t}_i(m_i^b, m_{-i}), s_i, \theta_i, c_i)] - \epsilon \end{aligned} \quad (20)$$

$$\geq E[u_i(b, \mathbf{x}_i(m_i^a, m_{-i}), \mathbf{t}_i(m_i^a, m_{-i}), s_i, \theta_i, c_i)] - \epsilon \quad (21)$$

$$\geq E[u_i(a, \mathbf{x}_i(m_i^a, m_{-i}), \mathbf{t}_i(m_i^a, m_{-i}), s_i, \theta_i, c_i)] - 2\epsilon \quad (22)$$

where inequalities (20) and (22) hold because the distance between a and b is at most ϵ , and inequality (21) holds because agent i prefers to send message m_i^b when $v_i = b$. Now, consider the deviation where agent i does not acquire the revealing signal and always sends message m_i^b . She reduces her costs by $p\tilde{c}$ and, by inequalities (20)-(22), reduces her gains by at most 2ϵ . This is a profitable deviation as long as $\epsilon < p\tilde{c}/2$.

Next, we construct a strategy profile $(\mathbf{s}', \mathbf{m}')$ and show that it is an equilibrium for instance I' . Let agent i acquire the bad signal. Since she has the good technology with probability one, she learns her value v_i . Let her send messages as follows.

1. If $v_i = (a + b)/2$, send the same (random) message m_i that an agent with the bad technology would have sent in instance I .
2. If type $v_i = a$, send the same (random) message m_i that an agent with the good technology and value $v_i = a$ would have sent in instance I .

3. If type $v_i = b$, send the same (random) message m_i that an agent with the good technology and value $v_i = b$ would have sent in instance I .

That is, the three types of agents that arise in joint distribution F' mimic the three types of agents that arise in joint distribution F .

To see that $(\mathbf{s}', \mathbf{m}')$ is an equilibrium, first consider information acquisition. When the joint distribution is F' , agent v_i knows her value perfectly, and she knows the state $\omega = 1$ with probability one. Therefore, there is no informational benefit to acquiring the good signal. The agent will always weakly prefer to acquire the bad signal, as prescribed.

We must also show that agent i does not want to deviate from message rule \mathbf{m}'_i . We rely on the following claim, which we prove later.

Claim 3. *The distribution of message profiles m' given joint distribution F' , where*

$$m'_i = \mathbf{m}'_i(\mathbf{s}'_i, \mathbf{s}'_i(\omega, \theta_i, v_i))$$

is the same as the distribution of message profiles m given joint distribution F .

Suppose that agent i sends message \tilde{m}_i . Having acquired the bad signal, her conditional expected utility is

$$E_{F'}[v_i \cdot \mathbf{x}_i(\tilde{m}_i, m'_{-i}) + \mathbf{t}_i(\tilde{m}_i, m'_{-i}) \mid v_i]$$

There are three cases to consider, corresponding to the three possible values v_i .

1. Suppose $v_i = (a + b)/2$. Then agent i 's expected utility is

$$\begin{aligned} & E_{F'} \left[v_i \cdot \mathbf{x}_i(\tilde{m}_i, m'_{-i}) + \mathbf{t}_i(\tilde{m}_i, m'_{-i}) \mid v_i = \frac{a + b}{2} \right] \\ &= E_{F'} \left[\frac{a + b}{2} \cdot \mathbf{x}_i(\tilde{m}_i, m'_{-i}) + \mathbf{t}_i(\tilde{m}_i, m'_{-i}) \right] && \text{(since } \theta_i \perp m'_{-i} \text{)} \\ &= E_F \left[\frac{a + b}{2} \cdot \mathbf{x}_i(\tilde{m}_i, m_{-i}) + \mathbf{t}_i(\tilde{m}_i, m_{-i}) \right] && \text{(by Claim 3)} \\ &= E_F [E_F[v_i] \cdot \mathbf{x}_i(\tilde{m}_i, m_{-i}) + \mathbf{t}_i(\tilde{m}_i, m_{-i})] && \text{(by inspection of } F \text{)} \\ &= E_F [v_i \cdot \mathbf{x}_i(\tilde{m}_i, m_{-i}) + \mathbf{t}_i(\tilde{m}_i, m_{-i})] && \text{(by LIE)} \end{aligned}$$

This objective is identical to the expected utility of the types in instance I with the bad technology. Since (\mathbf{s}, \mathbf{m}) is an equilibrium when the joint distribution is

F , those types must (almost surely) choose a message \tilde{m}_i that maximizes this objective. It follows that agent i must also choose messages \tilde{m}_i that maximize this objective, since she mimics those types.

2. Suppose $v_i = a$. Then agent i 's expected utility is

$$\begin{aligned}
& E_{F'}[a \cdot \mathbf{x}_i(\tilde{m}_i, m'_{-i}) + \mathbf{t}_i(\tilde{m}_i, m'_{-i}) \mid v_i = a] \\
&= E_{F'}[a \cdot \mathbf{x}_i(\tilde{m}_i, m'_{-i}) + \mathbf{t}_i(\tilde{m}_i, m'_{-i})] && \text{(since } v_i \perp m'_{-i}\text{)} \\
&= E_F[a \cdot \mathbf{x}_i(\tilde{m}_i, m_{-i}) + \mathbf{t}_i(\tilde{m}_i, m_{-i})] && \text{(by Claim 3)} \\
&= E_F[v_i \cdot \mathbf{x}_i(\tilde{m}_i, m_{-i}) + \mathbf{t}_i(\tilde{m}_i, m_{-i}) \mid v_i = a] && \text{(since } v_i \perp m_{-i}\text{)}
\end{aligned}$$

This objective is identical to the expected utility of the type in instance I with the good technology and value $v_i = a$. That type conditions on $v_i = a$ because she has the good technology. Since agent i mimics that type, we use the same argument as in the previous case to conclude that agent i is optimizing.

3. Suppose $v_i = b$. This argument is identical to the previous case, except that we replace $v_i = a$ with $v_i = b$.

It follows that message m'_i maximizes agent i 's expected utility.

Having shown that $(\mathbf{s}', \mathbf{m}')$ is an equilibrium when the instance is I' , the proof is almost complete. The final part of Lemma 1 states that the distribution of message profile m given joint distribution F is the same as the distribution of m' given joint distribution F' . This is precisely Claim 3.

B.4 Proof of Claim 2

It follows from the proof of Lemma 1 that instance I' has an equilibrium where all agents acquire the bad signal. Let agents with the good technology and value $v_i = a$ in instance I'' mimic agents with value $v_i = a$ in instance I' . Let agents with the good technology and value $v_i = b$ in instance I'' mimic agents with value $v_i = b$ in instance I' . Let agents with the bad technology in instance I'' mimic agents with value $v_i = (a + b)/2$ in instance I' . If all agents acquire the bad signal, we can verify (as in Lemma 1) that these messages are best responses for each type of agent in instance I'' . Furthermore, note that these messages do not vary with the state ω (because $\omega = 1$ with probability one in instance I'). It follows from the same argument as in Lemma 1

that, for \tilde{c} sufficiently large, agent i 's best response is to acquire the bad signal. Therefore, we have constructed an equilibrium for instance I'' where messages do not vary with the state ω .

B.5 Proof of Claim 3

We begin by showing that the messages m_i and m'_i have identical marginal distributions. This follows from three observations.

1. Suppose the joint distribution is F . With probability p , agent i has the bad technology. Her message is

$$m_i = \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i^{\theta}(\omega, \theta_i^B, v_i)) = \mathbf{m}_i(s_i^B, \theta_i^B)$$

where the second equality follows from the facts that (i) as previously discussed, agent i acquires the bad signal, and (ii) $s_i^B(\omega, \theta_i^B, v_i) = \theta_i^B$.

Suppose the joint distribution is F' . With probability p , agent i 's value is $v_i = (a + b)/2$. By construction, her message is the same as above.

2. Suppose the joint distribution is F . With probability $(1 - p)/2$, agent i has the good technology and learns that her value is $v_i = a$. Her message is

$$m_i = \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i^{\theta_i}(\omega, \theta_i^G, v_i)) = \mathbf{m}_i(s_i^B, (\omega, \theta_i^G, v_i)) = \mathbf{m}_i(s_i^B, (1, \theta_i^G, a))$$

Suppose the joint distribution is F' . With probability $(1 - p)/2$, agent i 's value is $v_i = a$. By construction, her message is the same as above.

3. Repeat the last argument, replacing $v_i = a$ with $v_i = b$.

Furthermore, the message m'_i is independent of m'_{-i} , by the same argument that showed m_i is independent of m_{-i} . Therefore, the distribution of message profiles m' given F' is the same as the distribution of message profiles m given F .

B.6 Proof of Lemmas 2, 6, and 8

Let $Z_{i,n}$ be a random variable that represents $\tilde{v}_{i,n}$ in Lemma 2, $\tilde{n}_{i,n}$ in Lemma 6, and $\tilde{v}_{i,n}$ and $\mathbf{1}(\tilde{v}_{i,n} \geq 0)$ in Lemma 8. We assume that $Z_{i,n}$ takes on values in the interval

$[z_L, z_H]$, and that $\epsilon_{i,n}$ and $Z_{i,n}$ are independent conditional on the state ω . We begin by introducing notation.

1. Let $\bar{c} < \infty$ be an upper bound on the cost of combined signal \bar{s}_i (Definition 17).
2. Define an analog $\tilde{\rho}_n$ to predictiveness ρ that asks how well a given signal s_i predicts $E[Z_{i,n} | \omega]$. Formally

$$\tilde{\rho}_n(s_i) = \lambda_n E \left[\max_{b_i} E_{i,n} \left[\text{SR}^{\text{CRPS}}(b_i, Z_{i,n}) \right] \right]$$

This proof consists of two claims.

Claim 4. Fix a strategy profile where agent i acquires a signal s_i where $c(\theta_i, s_i) < \infty$ for all technologies θ_i in the support of the joint distribution F . If the covariance between her error $\epsilon_{i,n}$ and the random variable $Z_{i,n}$ is positive, she can increase her expected score by acquiring the combined signal \bar{s} . That is,

$$\tilde{\rho}_n(\bar{s}_i) \geq \tilde{\rho}_n(s_i) + \frac{\lambda_n}{(v_H - v_L)^3} |\text{Cov}[\epsilon_{i,n}, Z_{i,n}]|^2$$

Proof. There are five steps to this proof. First, we derive a more useful expression for agent i 's maximum expected score given signal s_i .

$$\begin{aligned} \tilde{\rho}_n(s_i) &= E \left[\max_{b_i} E_{i,n} \left[\text{SR}^{\text{CRPS}}(b_i, Z_{i,n}) \right] \right] && \text{(definition of } \tilde{\rho}_n) \\ &= E \left[\max_{b_i} E_{i,n} \left[-\lambda_n \int_{[z_L, z_H]} (\Pr_{b_i}[Z_{i,n} \leq z] - \mathbf{1}(Z_{i,n} \leq z))^2 dz \right] \right] && \text{(defn. of SR)} \\ &= E \left[E_{i,n} \left[-\lambda_n \int_{[z_L, z_H]} (\Pr_{i,n}[Z_{i,n} \leq z] - \mathbf{1}(Z_{i,n} \leq z))^2 dz \right] \right] && \text{(SR is proper)} \\ &= -\lambda_n \int_{[z_L, z_H]} E \left[E_{i,n} \left[(\Pr_{i,n}[Z_{i,n} \leq z] - \mathbf{1}(Z_{i,n} \leq z))^2 \right] \right] dz && \text{(linearity)} \\ &= -\lambda_n \int_{[z_L, z_H]} E \left[E_{i,n} \left[(E_{i,n}[\mathbf{1}(Z_{i,n} \leq z)] - \mathbf{1}(Z_{i,n} \leq z))^2 \right] \right] dz && \text{(Pr}[E] = E[\mathbf{1}_E]) \\ &= -\lambda_n \int_{[z_L, z_H]} E \left[\text{Var}_{i,n}[\mathbf{1}(Z_{i,n} \leq z)] \right] dz && \text{(definition of variance)} \end{aligned}$$

Second, we derive a lower bound for agent i 's maximum expected score if she were to deviate to the combined signal \bar{s}_i . Before proceeding, we reiterate two points.

1. The combined signal \bar{s}_i includes the signal s_i , because s_i has finite cost for all technologies θ_i in the support of the joint distribution F .
2. The combined signal \bar{s}_i includes a revealing signal s'_i , which has finite cost for all technologies θ_i by Assumption 3.

It follows from these two observations that if agent i acquires \bar{s}_i , she will know the value v_i and the signal realization $s_i(\omega, \theta_i, v_i)$. As a result, agent i can infer $\epsilon_{i,n}$ – the error in her report when she only acquires signal s_i . With that, we can proceed.

$$\begin{aligned}
\tilde{\rho}_n(\bar{s}_i) &= \mathbb{E} \left[\max_{b_i} \mathbb{E} \left[\text{SR}^{\text{CRPS}}(b_i, Z_{i,n}) \mid \bar{s}_i, \bar{s}_i(\omega, \theta_i, v_i) \right] \right] && \text{(definition of } \tilde{\rho}_n) \\
&\geq \mathbb{E} \left[\max_{b_i} \mathbb{E} \left[\text{SR}^{\text{CRPS}}(b_i, Z_{i,n}) \mid s_i, s_i(\omega, \theta_i, v_i), \epsilon_{i,n} \right] \right] && \text{(throwing out info.)} \\
&= \mathbb{E} \left[\max_{b_i} \mathbb{E}_{i,n} \left[\text{SR}^{\text{CRPS}}(b_i, Z_{i,n}) \mid \epsilon_{i,n} \right] \right] && \text{(simplifying notation)} \\
&= -\lambda_n \int_{[z_L, z_H]} \mathbb{E} \left[\text{Var}_{i,n} [\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}] \right] dz && \text{(similar to earlier derivation)} \\
&= -\lambda_n \int_{[z_L, z_H]} \mathbb{E} \left[\mathbb{E}_{i,n} \left[\text{Var}_{i,n} [\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}] \right] \right] dz && \text{(LIE)}
\end{aligned}$$

Third, we obtain a lower bound on agent i 's expected gain in score that agent i if she were to deviate from signal s_i to the combined signal \bar{s}_i . Combining the first two parts of this proof,

$$\begin{aligned}
\tilde{\rho}_n(\bar{s}_i) - \tilde{\rho}_n(s_i) &\geq \lambda_n \int_{[z_L, z_H]} \mathbb{E} \left[\text{Var}_{i,n} [\mathbf{1}(Z_{i,n} \leq z)] - \mathbb{E}_{i,n} \left[\text{Var}_{i,n} [\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}] \right] \right] dz && (23) \\
&= \lambda_n \int_{[z_L, z_H]} \mathbb{E} \left[\text{Var}_{i,n} \left[\mathbb{E}_{i,n} [\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}] \right] \right] dz && \text{(law of total variance)}
\end{aligned}$$

Fourth, for any constant z , we use the conditional covariance between $\epsilon_{i,n}$ and $\mathbf{1}(Z_{i,n} \leq z)$ to bound the conditional variance of $\mathbb{E}_{i,n} [\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}]$.

$$\text{Cov}_{i,n} [\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)]$$

$$\begin{aligned}
&= \mathbb{E}_{i,n}[\epsilon_{i,n} (\mathbf{1}(Z_{i,n} \leq z) - \mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z)])] && \text{(since } \mathbb{E}_{i,n}[\epsilon_{i,n}] = 0) \\
&= \mathbb{E}_{i,n}[\mathbb{E}_{i,n}[\epsilon_{i,n} (\mathbf{1}(Z_{i,n} \leq z) - \mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z)]) \mid \epsilon_{i,n}]] && \text{(LIE)} \\
&= \mathbb{E}_{i,n}[\epsilon_{i,n} \mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z) - \mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z)] \mid \epsilon_{i,n}]] && \text{(linearity)} \\
&\leq \sqrt{\mathbb{E}_{i,n}[\epsilon_{i,n}^2 \mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z) - \mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z)] \mid \epsilon_{i,n}]^2]} && \text{(Jensen's inequality)} \\
&\leq \sqrt{\mathbb{E}_{i,n}[(v_H - v_L)^2 \mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z) - \mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z)] \mid \epsilon_{i,n}]^2]} && \text{(since } |\epsilon_{i,n}| \leq v_H - v_L \text{ and } \mathbb{E}_{i,n}[\cdot]^2 \geq 0) \\
&= (v_H - v_L) \sqrt{\mathbb{E}_{i,n}[(\mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}] - \mathbb{E}_{i,n}[\mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}]])^2]} && \text{(linearity)} \\
&= (v_H - v_L) \sqrt{\text{Var}_{i,n}[\mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}]]} && \text{(definition of variance)}
\end{aligned}$$

Rearranging this inequality gives us

$$\text{Var}_{i,n}[\mathbb{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}]] \geq \frac{1}{(v_H - v_L)^2} \text{Cov}_{i,n}[\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)]^2 \quad (24)$$

Fifth and finally, we relax the lower bound in step 3, in terms of quantities that appear in the claim statement. Combining inequalities (23) and (24), we obtain

$$\begin{aligned}
&\tilde{\rho}_n(\bar{s}_i) - \tilde{\rho}_n(s_i) && (25) \\
&\geq \frac{\lambda_n}{(v_H - v_L)^2} \int_{[z_L, z_H]} \mathbb{E}[\text{Cov}_{i,n}[\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)]^2] dz && \text{(combine (23) and (24))} \\
&= \frac{\lambda_n}{(v_H - v_L)^2} \mathbb{E} \left[\int_{[z_L, z_H]} \text{Cov}_{i,n}[\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)]^2 dz \right] && \text{(linearity)}
\end{aligned}$$

Next, focus on the term inside the expectation.

$$\begin{aligned}
\int_{[z_L, z_H]} \text{Cov}_{i,n}[\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)]^2 dz &= \|\text{Cov}_{i,n}[\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)]\|_2^2 && \text{(definition of } L^2 \text{ norm)} \\
&\geq \frac{1}{(v_H - v_L)} \|\text{Cov}_{i,n}[\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)]\|_1^2 && \text{(Hölder's inequality)} \\
&= \frac{1}{(v_H - v_L)} \left(\int_{[z_L, z_H]} |\text{Cov}_{i,n}[\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)]| dz \right)^2 && \text{(defn. of } L^1 \text{ norm)}
\end{aligned}$$

$$\geq \frac{1}{(v_H - v_L)} \left(\int_{[z_L, z_H]} \text{Cov}_{i,n}[\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)] dz \right)^2$$

(since $|X| \geq X$)

Now, focus on the term inside the square.

$$\begin{aligned} & \int_{[z_L, z_H]} \text{Cov}_{i,n}[\epsilon_{i,n}, \mathbf{1}(Z_{i,n} \leq z)] dz \\ &= \int_{[z_L, z_H]} \text{E}_{i,n}[\epsilon_{i,n} \mathbf{1}(Z_{i,n} \leq z)] dz && \text{(since } \text{E}_{i,n}[\epsilon_{i,n}] = 0\text{)} \\ &= \int_{[z_L, z_H]} \text{E}_{i,n}[\text{E}_{i,n}[\epsilon_{i,n} \mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}]] dz && \text{(LIE)} \\ &= \int_{[z_L, z_H]} \text{E}_{i,n}[\epsilon_{i,n} \text{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}]] dz && \text{(linearity)} \\ &= \text{E}_{i,n} \left[\epsilon_{i,n} \int_{[z_L, z_H]} \text{E}_{i,n}[\mathbf{1}(Z_{i,n} \leq z) \mid \epsilon_{i,n}] dz \right] && \text{(linearity)} \\ &= \text{E}_{i,n} \left[\epsilon_{i,n} \int_{[z_L, z_H]} \text{Pr}_{i,n}[Z_{i,n} \leq z \mid \epsilon_{i,n}] dz \right] && \text{(since } \text{Pr}[E] = \text{E}[\mathbf{1}_E]\text{)} \\ &= (v_H - v_L) \text{E}_{i,n}[\epsilon_{i,n}] - \text{E}_{i,n} \left[\epsilon_{i,n} \int_{[z_L, z_H]} (1 - \text{Pr}_{i,n}[Z_{i,n} \leq z \mid \epsilon_{i,n}]) dz \right] && \text{(linearity)} \\ &= -\text{E}_{i,n} \left[\epsilon_{i,n} \int_{[z_L, z_H]} (1 - \text{Pr}_{i,n}[Z_{i,n} \leq z \mid \epsilon_{i,n}]) dz \right] && \text{(since } \text{E}_{i,n}[\epsilon_{i,n}] = 0\text{)} \\ &= -\text{E}_{i,n}[\epsilon_{i,n} \text{E}_{i,n}[Z_{i,n} \mid \epsilon_{i,n}]] && \text{(definition of } \text{E}[\cdot]\text{)} \\ &= -\text{E}_{i,n}[\text{E}_{i,n}[\epsilon_{i,n} Z_{i,n} \mid \epsilon_{i,n}]] && \text{(linearity)} \\ &= -\text{E}_{i,n}[\epsilon_{i,n} Z_{i,n}] && \text{(LIE)} \\ &= -\text{Cov}_{i,n}[\epsilon_{i,n}, Z_{i,n}] && \text{(since } \text{E}_{i,n}[\epsilon_{i,n}] = 0\text{)} \end{aligned}$$

Combine this sequence of inequalities, starting with (25), to obtain

$$\begin{aligned} \tilde{\rho}_n(\bar{s}_i) - \tilde{\rho}_n(s_i) &\geq \frac{\lambda_n}{(v_H - v_L)^2} \text{E} \left[\frac{1}{(v_H - v_L)} (-\text{Cov}_{i,n}[\epsilon_{i,n}, Z_{i,n}])^2 \right] \\ &= \frac{\lambda_n}{(v_H - v_L)^3} \text{E}[\text{Cov}_{i,n}[\epsilon_{i,n}, Z_{i,n}]^2] && \text{(linearity)} \\ &= \frac{\lambda_n}{(v_H - v_L)^3} \text{Cov}[\epsilon_{i,n}, Z_{i,n}]^2 && \text{(LIE and since } \text{E}_{i,n}[\epsilon_{i,n}] = 0\text{)} \end{aligned}$$

(26)

This, rearranged, gives us the statement of the claim. \square

Having established conditions under which the agent could increase her expected score by acquiring the combined signal \bar{s} , we can limit the extent to which those conditions hold in an equilibrium. More precisely, we argue that the covariance between the error $\epsilon_{i,n}$ and the random variable $Z_{i,n}$ cannot be too large in equilibrium.

Claim 5. *In any equilibrium, the covariance between the error $\epsilon_{i,n}$ and the conditional expectation $E[Z_{i,n} | \omega]$ is bounded. That is,*

$$|\text{Cov}[\epsilon_{i,n}, Z_{i,n}]| = |\text{Cov}[\epsilon_{i,n}, E[Z_{i,n} | \omega]]| \leq \sqrt{\frac{(v_H - v_L)^3 \bar{c}}{\lambda_n}}$$

Proof. First, we study the covariance between agent i 's error $\epsilon_{i,n}$ and $Z_{i,n}$. There are two cases to consider.

1. Suppose that agent i does not acquire the combined signal \bar{s}_i . Then she could deviate from the signal s_i that she did acquire to \bar{s}_i . This will change her expected transfers associated with the scoring rule SR^{CRPS} from $\tilde{\rho}_n(s_i)$ to $\tilde{\rho}_n(\bar{s})$. This may also increase her costs, but the increase will be at most \bar{c} . Putting everything together, the fact that the deviation cannot be profitable in equilibrium implies

$$\tilde{\rho}_n(\bar{s}_i) - \bar{c} \leq \tilde{\rho}_n(s_i)$$

Combining this inequality with Claim 4 yields⁸

$$\bar{c} \geq \frac{\lambda_n}{(v_H - v_L)^3} |\text{Cov}[\epsilon_{i,n}, Z_{i,n}]|^2 \quad (27)$$

2. Suppose that agent i acquires the combined signal \bar{s}_i . Then, almost surely, she knows v_i and reports $\hat{v}_{i,n} = v_i$. Conditional on agent i 's information, the error $\epsilon_{i,n}$ is almost surely constant and so its covariance with any other random variable is zero. Therefore, inequality (27) still holds, albeit vacuously.

⁸Claim 4 requires the signal s_i acquired by agent i to have finite cost for all technologies in the support of the joint distribution F , i.e., $c(\theta_i, s_i) < \infty$. But this holds in every equilibrium; otherwise it would be profitable to deviate to the combined signal \bar{s}_i , which has finite expected cost.

Second, we verify that the covariance between ϵ_i and $Z_{i,n}$ is equal to the covariance between ϵ_i and the conditional expectation $E[Z_{i,n} | \omega]$.

$$\begin{aligned}
\text{Cov}[\epsilon_{ix}, Z_{i,n}] &= E[\epsilon_{ix} Z_{i,n}] && \text{(since } E[\epsilon_{ix}] = 0\text{)} \\
&= E[E[\epsilon_{ix} Z_{i,n} | \omega]] && \text{(LIE)} \\
&= E[E[\epsilon_{ix} | \omega] E[Z_{i,n} | \omega]] && \text{(cond. independence)} \\
&= E[E[\epsilon_{ix} E[Z_{i,n} | \omega] | \omega]] && \text{(linearity)} \\
&= E[\epsilon_{ix} E[Z_{i,n} | \omega]] && \text{(LIE)} \\
&= \text{Cov}[\epsilon_{ix}, E[Z_{i,n} | \omega]] && \text{(since } E[\epsilon_{ix}] = 0\text{)}
\end{aligned}$$

Combining this identity with inequality (27) completes the proof. \square

B.7 Proof of Corollary 2

Fix an instance $I \in \mathcal{J}$, an equilibrium $(\mathbf{s}^n, \mathbf{m}^n)$, and an alternative $x \in \mathcal{X}$. Let agent i send message $m_i = \mathbf{m}_i(\mathbf{s}_i, \mathbf{s}_i(\omega, \theta_i, v_i))$. The researcher's mean-square error is

$$\begin{aligned}
&E\left[(f_n(m_1, \dots, m_n) - E[\sigma(v_i) | \omega])^2\right] && (28) \\
&= E\left[\left(\frac{1}{n} \sum_{i=1}^n \hat{v}_{ix} - E[v_{ix} | \omega]\right)^2\right] && \text{(defn. of } f_n \text{ and } \sigma) \\
&= E[E[v_{ix} | \omega]^2] + E\left[\left(\frac{1}{n} \sum_{i=1}^n \hat{v}_{ix}\right)^2\right] - 2E\left[E[v_{ix} | \omega] \cdot \frac{1}{n} \sum_{i=1}^n \hat{v}_{ix}\right] && \text{(linearity of } E[\cdot])
\end{aligned}$$

Focus on the third term, i.e.,

$$\begin{aligned}
&-2E\left[E[v_{ix} | \omega] \cdot \frac{1}{n} \sum_{i=1}^n \hat{v}_{ix}\right] && (29) \\
&= -2E\left[E[\hat{v}_{ix} | \omega] \cdot \frac{1}{n} \sum_{i=1}^n \hat{v}_{ix}\right] + 2E\left[E[\epsilon_{ix} | \omega] \cdot \frac{1}{n} \sum_{i=1}^n \hat{v}_{ix}\right] && \text{(defn. of } \epsilon_{ix}) \\
&= -2E[E[\hat{v}_{ix} | \omega] \cdot \hat{v}_{ix}] + 2E[E[\epsilon_{ix} | \omega] \cdot \hat{v}_{ix}] && \text{(linearity of } E[\cdot]) \\
&= -2E[E[E[\hat{v}_{ix} | \omega] \cdot \hat{v}_{ix} | \omega]] + 2E[E[E[\epsilon_{ix} | \omega] \cdot \hat{v}_{ix} | \omega]] && \text{(LIE)} \\
&= -2E[E[\hat{v}_{ix} | \omega] \cdot E[\hat{v}_{ix} | \omega]] + \frac{2}{n} \sum_{i=1}^n E[E[\epsilon_{ix} | \omega] \cdot E[\hat{v}_{ix} | \omega]] && \text{(linearity of } E[\cdot])
\end{aligned}$$

$$\begin{aligned}
&= -2\mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + 2\mathbb{E}[\mathbb{E}[\epsilon_{ix} \cdot \mathbb{E}[\hat{v}_{ix} \mid \omega] \mid \omega]] && \text{(linearity of } \mathbb{E}[\cdot] \text{)} \\
&= -2\mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + 2\mathbb{E}[\epsilon_{ix} \cdot \mathbb{E}[\hat{v}_{ix} \mid \omega]] && \text{(LIE)} \\
&= -2\mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + 2\text{Cov}[\epsilon_{ix} \cdot \mathbb{E}[\hat{v}_{ix} \mid \omega]] && \text{(since } \mathbb{E}[\epsilon_{ix}] = 0 \text{)} \\
&= -2\mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + o(1) && \text{(by Lemma 2)}
\end{aligned}$$

Returning to the earlier expression (28), we find that

$$\begin{aligned}
(28) &= \mathbb{E}[\mathbb{E}[v_{ix} \mid \omega]^2] + \mathbb{E}\left[\left(\frac{1}{n} \sum_{i=1}^n \hat{v}_{ix}\right)^2\right] - 2\mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + o(1) && (30) \\
&= \mathbb{E}[\mathbb{E}[v_{ix} \mid \omega]^2] + \mathbb{E}\left[\mathbb{E}\left[\left(\frac{1}{n} \sum_{i=1}^n \hat{v}_{ix}\right)^2 \mid \omega\right]\right] - 2\mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + o(1) && \text{(LIE)}
\end{aligned}$$

Since \hat{v}_{ix} depends on the equilibrium $(\mathbf{s}^n, \mathbf{m}^n)$, its distribution varies with n . The weak law of large numbers for triangular arrays ensures that $\frac{1}{n} \sum_{i=1}^n \hat{v}_{ix} \rightarrow_p \mathbb{E}[\hat{v}_{ix} \mid \omega]$.

Combining this with the continuous mapping theorem, we find that

$$\begin{aligned}
(28) &= \mathbb{E}[\mathbb{E}[v_{ix} \mid \omega]^2] + \mathbb{E}[\mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2 + o_p(1) \mid \omega]] - 2\mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + o(1) && (31) \\
&= \mathbb{E}[\mathbb{E}[v_{ix} \mid \omega]^2] + \mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] - 2\mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + o(1) && \text{(LIE)} \\
&= \mathbb{E}[\mathbb{E}[v_{ix} \mid \omega]^2] - \mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + o(1) && (32) \\
&= \mathbb{E}[\mathbb{E}[v_{ix} \mid \omega]^2] - \mathbb{E}[v_{ix}]^2 + \mathbb{E}[\hat{v}_{ix}]^2 - \mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + o(1) && \text{(since } \mathbb{E}[\epsilon_{ix}] = 0 \text{)} \\
&= \mathbb{E}[\mathbb{E}[v_{ix} \mid \omega]^2] - \mathbb{E}[\mathbb{E}[v_{ix} \mid \omega]]^2 + \mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]]^2 - \mathbb{E}[\mathbb{E}[\hat{v}_{ix} \mid \omega]^2] + o(1) && \text{(LIE)} \\
&= \text{Var}[\mathbb{E}[v_{ix} \mid \omega]] - \text{Var}[\mathbb{E}[\hat{v}_{ix} \mid \omega]] + o(1) && \text{(defn. of } \text{Var}[\cdot] \text{)} \\
&\leq \text{Var}[\mathbb{E}[v_{ix} \mid \omega]] + o(1) && \text{(since } \text{Var}[\cdot] \geq 0 \text{)}
\end{aligned}$$

Therefore, the no-information benchmark is satisfied.

B.8 Proof of Lemmas 3 and 9

This proof consists of three claims. It uses notation from the proof of Lemmas 2 and 8.

Let agent i acquire signal $s_i = \mathbf{s}_i(\theta_i, c_i)$ and report values $\hat{v}_i = \hat{\mathbf{v}}_i(s_i, s_i(\omega, \theta_i, v_i))$. We assume that

$$\Pr\left[\max_{i=1, \dots, n} |\hat{v}_i - \mathbb{E}[v_i \mid s_i, s_i(\omega, \theta_i, v_i)]| > \phi_{1,n}\right] \leq \phi_{2,n} \quad (33)$$

We define

$$\tilde{v}_{i,n} = \frac{1}{n'} \left(v_{0,n} + \sum_{j \neq i} \hat{v}_{j,n} \right)$$

For Lemma 3, $\phi_{1,n} = \phi_{2,n} = 0$, $n' = n - 1$, $v_{0,n} = 0$, and $\delta_n = 0$. For Lemma 9, define these parameters as in Appendix B.1.

We begin by showing that the cross-covariance between the error $\epsilon_{i,n}$ and the conditional expectation $E[v_j | \omega]$ is zero. This follows by construction of the auxiliary instance and an argument similar to Lemmas 2 and 8.

Claim 6. *Consider a strategy profile in the auxiliary instance where agent i acquires a signal s_i with finite expected cost, i.e., $E[\tilde{c}(\theta_i, s_i)] < \infty$. Then the covariance between the error $\epsilon_{i,n}$ and the conditional expected values $E[v_j | \omega]$ is zero, i.e.,*

$$\left\| \text{Cov}[\epsilon_{i,n}, E[v_j | \omega]] \right\| = 0$$

Proof. Note that the proof of Claim 4 holds verbatim if we replace $\tilde{\rho}_n$ with $\lambda_n \rho$ and $\tilde{v}_{i,n}$ with $E[v_j | \omega]$. After dividing by λ_n , the statement of that claim becomes

$$\rho(\bar{s}_i) \geq \rho(s_i) + \frac{1}{(v_H - v_L)^3} \left\| \text{Cov}[\epsilon_i, E[v_j | \omega]] \right\|^2$$

However, since the signal s_i has finite expected cost according to auxiliary cost function \tilde{c} , it follows from the definition of \tilde{c} in (17) that

$$\rho(\bar{s}_i) \geq \rho(s_i)$$

Combining these two inequalities yields,

$$0 = \frac{1}{(v_H - v_L)^3} \left\| \text{Cov}[\epsilon_i, E[v_j | \omega]] \right\|^2 = \left\| \text{Cov}[\epsilon_i, E[v_j | \omega]] \right\|$$

which implies the statement of the claim. □

With Claims 5 and 6 in hand, we can argue that the variance of the conditional expectation of the error cannot be too large.

Claim 7. *In any equilibrium of the auxiliary instance, the variance of the conditional*

expected error is bounded, i.e.,

$$\text{Var} \left[\mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n \epsilon_{j,n} \mid \omega \right] \right] = o(1) \quad \text{and} \quad \text{Var} \left[\mathbb{E} \left[\frac{1}{n-1} \sum_{j \neq i} \epsilon_{j,n} \mid \omega \right] \right] = o(1)$$

Proof. Observe that, with probability $1 - \phi_{2,n}$,

$$\begin{aligned} & \text{Var} \left[\mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n \epsilon_{j,n} \mid \omega \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n \epsilon_{j,n} \mid \omega \right]^2 \right] && \text{(since } \mathbb{E}[\epsilon_{i,n}] = 0 \text{)} \\ &= \mathbb{E} \left[\mathbb{E} \left[\mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n \epsilon_{j,n} \mid \omega \right] \cdot \frac{1}{n} \sum_{j=1}^n \epsilon_{j,n} \mid \omega \right] \right] && \text{(linearity)} \\ &= \mathbb{E} \left[\mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n \epsilon_{j,n} \mid \omega \right] \cdot \frac{1}{n} \sum_{j=1}^n \epsilon_{j,n} \right] && \text{(LIE)} \\ &= \frac{1}{n^2} \sum_{j=1}^n \mathbb{E} \left[\left(\mathbb{E}[\epsilon_{j,n} \mid \omega] + \sum_{i \neq j} \mathbb{E}[\epsilon_{i,n} \mid \omega] \right) \cdot \epsilon_{i,n} \right] && \text{(linearity)} \\ &= \frac{1}{n^2} \sum_{j=1}^n \mathbb{E} \left[\left(\mathbb{E}[\epsilon_{j,n} \mid \omega] + \sum_{i \neq j} \mathbb{E}[v_i \mid \omega] - \sum_{i \neq j} \mathbb{E}[\mathbb{E}_{i,n}[v_i] \mid \omega] \right) \cdot \epsilon_{j,n} \right] && \text{(defn. of } \epsilon_{i,n} \text{)} \\ &\leq \frac{1}{n^2} \sum_{j=1}^n \mathbb{E} \left[\left(\mathbb{E}[\epsilon_{j,n} \mid \omega] + \sum_{i \neq j} \mathbb{E}[v_i \mid \omega] - \sum_{i \neq j} \mathbb{E}[\hat{v}_{i,n} \mid \omega] \right) \cdot \epsilon_{j,n} \right] \\ &\quad + \frac{n(n-1)(v_H - v_L)}{n^2} \cdot \phi_{1,n} && \text{(condition (33))} \\ &= \frac{1}{n^2} \sum_{j=1}^n \mathbb{E} \left[\left(\mathbb{E}[\epsilon_{j,n} \mid \omega] - \sum_{i \neq j} \mathbb{E}[\hat{v}_{i,n} \mid \omega] \right) \cdot \epsilon_{j,n} \right] + \frac{n(n-1)(v_H - v_L)}{n^2} \cdot \phi_{1,n} && \text{(Claim 6)} \\ &\leq \frac{(v_H - v_L)^2}{n} - \frac{1}{n^2} \sum_{j=1}^n \mathbb{E} \left[\epsilon_{j,n} \sum_{i \neq j} \mathbb{E}[\hat{v}_{i,n} \mid \omega] \right] + \frac{n(n-1)(v_H - v_L)}{n^2} \cdot \phi_{1,n} \\ & && \text{(since } \epsilon_{i,n} \in [v_L, v_H] \text{)} \\ &= \frac{(v_H - v_L)^2}{n} - \frac{n'}{n^2} \sum_{j=1}^n \mathbb{E}[\epsilon_{j,n} \tilde{v}_{j,n}] + \frac{n(n-1)(v_H - v_L)}{n^2} \cdot \phi_{1,n} \\ & && \text{(defn. of } \tilde{v}_{i,n} \text{ and since } v_{0,n} \perp \epsilon_{j,n} \text{)} \end{aligned}$$

$$\begin{aligned}
&= \frac{(v_H - v_L)^2}{n} + \frac{n'}{n} \sqrt{\frac{(v_H - v_L)^3 \bar{c}}{\lambda_n}} + \frac{n(n-1)(v_H - v_L)}{n^2} \cdot \phi_{1,n} && \text{(Claim 5)} \\
&\leq \frac{(v_H - v_L)^2}{n} + \sqrt{\frac{(v_H - v_L)^3 \bar{c}}{\lambda_n}} + (v_H - v_L) \phi_{1,n} && \text{(since } n' \leq n \text{ and } n-1 \leq n)
\end{aligned}$$

Note that the probability $\phi_{2,n}$ event where this may not hold can affect the variance by at most $(v_H - v_L)^2 \cdot \phi_{2,n}$. Finally, the second part of the claim follows from the fact that adding a term of size $O(n^{-1})$ has a vanishing impact on the variance. \square

We use Claim 7 to argue that the average error concentrates around zero.

Claim 8. *For agent i , the average reported value $\tilde{v}_{i,n}$ concentrates around the conditional expectation $E[v_j \mid \omega]$. That is,*

$$\forall t > 0 : \quad \lim_{n \rightarrow \infty} \Pr[|\tilde{v}_{i,n} - E[v_j \mid \omega]| \geq t] = 0$$

Proof. The proof has three parts. First, we show that a particular sum of random variables has expected value $\mu_n = 0$. This sum will become relevant later on.

$$\begin{aligned}
\mu_n &:= E \left[\frac{1}{n-1} \sum_{j \neq i} (v_j - E[v_j \mid \omega] - \epsilon_{j,n}) \right] \\
&= \frac{1}{n-1} \sum_{j \neq i} (E[v_j] - E[E[v_j \mid \omega]]) && \text{(linearity, and since } E[\epsilon_{j,n}] = 0) \\
&= \frac{1}{n-1} \sum_{j \neq i} (E[v_j] - E[v_j]) && \text{(LIE)} \\
&= 0
\end{aligned}$$

Second, we bound the variance σ_n^2 of this same sum of random variables.

$$\begin{aligned}
\sigma_n^2 &:= \text{Var} \left[\frac{1}{n-1} \sum_{j \neq i} (v_j - E[v_j \mid \omega] - \epsilon_{j,n}) \right] \\
&= E \left[\left(\frac{1}{n-1} \sum_{j \neq i} (v_j - E[v_j \mid \omega] - \epsilon_{j,n}) \right)^2 \right] && \text{(since } \mu_n = 0)
\end{aligned}$$

$$\begin{aligned}
&\leq 2\mathbb{E}\left[\left(\frac{1}{n-1}\sum_{j\neq i}(v_j - \mathbb{E}[v_j | \omega])\right)^2\right] + 2\mathbb{E}\left[\left(\frac{1}{n-1}\sum_{j\neq i}\epsilon_{j,n}\right)^2\right] \quad (\text{Cauchy-Schwarz}) \\
&= 2\mathbb{E}\left[\mathbb{E}\left[\left(\frac{1}{n-1}\sum_{j\neq i}(v_j - \mathbb{E}[v_j | \omega])\right)^2 \mid \omega\right]\right] + 2\mathbb{E}\left[\mathbb{E}\left[\left(\frac{1}{n-1}\sum_{j\neq i}\epsilon_{j,n}\right)^2 \mid \omega\right]\right] \\
&\hspace{20em} (\text{LIE})
\end{aligned}$$

We will consider both of these terms in turn. The first term is

$$\begin{aligned}
\mathbb{E}\left[\left(\frac{1}{n-1}\sum_{j\neq i}(v_j - \mathbb{E}[v_j | \omega])\right)^2 \mid \omega\right] &= \text{Var}\left[\frac{1}{n-1}\sum_{j\neq i}v_j \mid \omega\right] \quad (\text{defn. of variance}) \\
&= \frac{1}{n-1}\text{Var}[v_j \mid \omega] \quad (\text{cond. indep.}) \\
&\leq \frac{(v_H - v_L)^2}{n-1} \quad (\text{since } v_j \in [v_L, v_H])
\end{aligned}$$

Now, consider the second term.

$$\begin{aligned}
&\mathbb{E}\left[\left(\frac{1}{n-1}\sum_{j\neq i}\epsilon_{j,n}\right)^2 \mid \omega\right] \hspace{10em} (34) \\
&= \text{Var}\left[\frac{1}{n-1}\sum_{j\neq i}\epsilon_{j,n} \mid \omega\right] + \left(\mathbb{E}\left[\frac{1}{n-1}\sum_{j\neq i}\epsilon_{j,n} \mid \omega\right]\right)^2 \quad (\text{defn. of Var}[\cdot]) \\
&= \frac{1}{(n-1)^2}\sum_{j\neq i}\text{Var}[\epsilon_{j,n} \mid \omega] + \left(\mathbb{E}\left[\frac{1}{n-1}\sum_{j\neq i}\epsilon_{j,n} \mid \omega\right]\right)^2 \quad (\text{cond. indep.}) \\
&\leq \frac{(v_H - v_L)^2}{n-1} + \left(\mathbb{E}\left[\frac{1}{n-1}\sum_{j\neq i}\epsilon_{j,n} \mid \omega\right]\right)^2 \quad (\text{since } \epsilon_{j,n} \in [v_L, v_H])
\end{aligned}$$

We plug these new expressions for the two terms back into our inequality.

$$\sigma_n^2 \leq \frac{4(v_H - v_L)^2}{n-1} + 2\mathbb{E}\left[\left(\mathbb{E}\left[\frac{1}{n-1}\sum_{j\neq i}\epsilon_{j,n} \mid \omega\right]\right)^2\right] \quad (35)$$

$$\begin{aligned}
&= \frac{4(v_H - v_L)^2}{n-1} + 2\text{Var}\left[\mathbb{E}\left[\frac{1}{n-1}\sum_{j \neq i} \epsilon_{j,n} \mid \omega\right]\right] + 2\left(\mathbb{E}\left[\mathbb{E}\left[\frac{1}{n-1}\sum_{j \neq i} \epsilon_{j,n} \mid \omega\right]\right]\right)^2 \\
&\hspace{15em} \text{(defn. of variance)} \\
&= \frac{4(v_H - v_L)^2}{n-1} + 2\text{Var}[\mathbb{E}[\epsilon_{j,n} \mid \omega]] \hspace{10em} \text{(LIE and since } \mathbb{E}[\epsilon_{j,n}] = 0) \\
&\leq \frac{4(v_H - v_L)^2}{n-1} + o(1) \hspace{15em} \text{(Claim 7)}
\end{aligned}$$

Third and finally, we show that $\tilde{v}_{i,n}$ concentrates around $\mathbb{E}[v_j \mid \omega]$. Given $t > 0$,

$$\begin{aligned}
&\Pr\left[|\tilde{v}_{i,n} - \mathbb{E}[v_j \mid \omega]| \geq t\right] \\
&= \Pr\left[\left|\frac{1}{n'}\left(v_{0,n} + \sum_{j \neq i} \hat{v}_{j,n}\right) - \mathbb{E}[v_j \mid \omega]\right| \geq t\right] \hspace{5em} \text{(defn. of } \tilde{v}_{i,n}) \\
&\leq \Pr\left[\left|\frac{v_{0,n}}{n}\right| + \left|\frac{n' - n + 1}{n'(n-1)}\left(\sum_{j \neq i} \hat{v}_{j,n}\right)\right| + \left|\frac{1}{n-1}\left(\sum_{j \neq i} \hat{v}_{j,n}\right) - \mathbb{E}[v_j \mid \omega]\right| \geq t\right] \\
&\hspace{15em} \text{(triangle inequality)} \\
&\leq \Pr\left[\left|\frac{v_{0,n}}{n}\right| \geq \frac{t}{3}\right] + \Pr\left[\left|\frac{n' - n + 1}{n'(n-1)}\left(\sum_{j \neq i} \hat{v}_{j,n}\right)\right| \geq \frac{t}{3}\right] \hspace{5em} (36) \\
&\quad + \Pr\left[\left|\frac{1}{n-1}\left(\sum_{j \neq i} \hat{v}_{j,n}\right) - \mathbb{E}[v_j \mid \omega]\right| \geq \frac{t}{3}\right] \hspace{10em} \text{(union bound)}
\end{aligned}$$

The first two terms of line (36) are vanishing, since $\beta_n = o(n)$ and $n' = O(n)$. The third term is vanishing as well, since

$$\begin{aligned}
&\Pr\left[\left|\frac{1}{n-1}\sum_{j \neq i} (\hat{v}_j - \mathbb{E}[v_j \mid \omega])\right| \geq \frac{t}{3}\right] \\
&\leq \delta_n + \phi_{1,n} + \phi_{2,n} + \Pr\left[\left|\frac{1}{n-1}\sum_{j \neq i} (\mathbb{E}_{i,n}[v_j] - \mathbb{E}[v_j \mid \omega])\right| \geq \frac{t}{3}\right] \\
&\hspace{15em} \text{(defn. of } \delta_n, \phi_{1,n}, \phi_{2,n}) \\
&= \delta_n + \phi_{1,n} + \phi_{2,n} + \Pr\left[\left|\frac{1}{n-1}\sum_{j \neq i} (v_j - \epsilon_{j,n} - \mathbb{E}[v_j \mid \omega])\right| \geq \frac{t}{3}\right] \hspace{5em} \text{(definition of } \epsilon_{i,n})
\end{aligned}$$

$$\begin{aligned}
&= \delta_n + \phi_{1,n} + \phi_{2,n} + \Pr \left[\left| \frac{1}{n-1} \sum_{j \neq i} (v_j - \epsilon_{j,n} - \mathbb{E}[v_j | \omega]) - \mu_n \right| \geq \frac{t}{3} \right] \quad (\text{since } \mu_n = 0) \\
&\leq \delta_n + \phi_{1,n} + \phi_{2,n} + \frac{9\sigma_n^2}{t^2} \quad (\text{Chebyshev's inequality}) \\
&\leq \delta_n + \phi_{1,n} + \phi_{2,n} + \frac{9}{t^2} \left(\frac{4(v_H - v_L)^2}{n-1} + o(1) \right) \quad (\text{inequality (35)})
\end{aligned}$$

Therefore, $\tilde{v}_{i,n}$ converges in probability to $\mathbb{E}[v_j | \omega]$. \square

Finally, let the betting stakes $\lambda_n \rightarrow \infty$ as the sample size $n \rightarrow \infty$. Claim 8 implies

$$\tilde{v}_{i,n} \rightarrow_p \mathbb{E}[v_j | \omega]$$

It follows from this and the continuous mapping theorem that

$$\frac{1}{n'+1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \right) = \frac{\hat{v}_{i,n}}{n'+1} + \left(\frac{n'}{n'+1} \right) \tilde{v}_{i,n} \rightarrow_p \mathbb{E}[v_j | \omega]$$

This completes the proof.

B.9 Proof of Lemmas 4 and 10

This proof has four parts. It uses notation from Appendices B.6 and B.8.

First, consider any signal s_i where, for any technology θ_i , its cost in the auxiliary instance is different from its cost in the original instance. Equivalently, consider any signal s_i where

$$\rho(\bar{s}_i) - \rho(s_i) > 0$$

Since \mathcal{S} is a finite set, we can define a gap $\delta > 0$ where

$$\delta := \min_{s_i \text{ s.t. } \rho(s_i) < \rho(\bar{s}_i)} \rho(\bar{s}_i) - \rho(s_i) \quad (37)$$

This will be useful later in the proof.

Second, we consider a sequence of moments and show that it converges. For any random variable Z , define a moment that depends on additional parameter $z \in \mathbb{R}$, i.e.,

$$f(Z, z) = \mathbb{E} \left[(\Pr_i[Z \leq z] - \mathbf{1}(Z \leq z))^2 \right]$$

Let Z_n be a sequence of random variables where $Z_n \rightarrow_p Z$. To find $\lim_{n \rightarrow \infty} f(Z_n, \cdot)$, we first consider the limits of two random variables that appear in the definition of f .

1. The random variable $\Pr_i[Z_n \leq z]$ converges in probability to $\Pr_i[Z \leq z]$ for almost all z . To see this, fix any constant $t > 0$ and define the probability

$$p_n^+ := \Pr[\Pr_i[Z \leq z] - \Pr_i[Z_n \leq z] \geq t]$$

I begin by showing that $p_n^+ \rightarrow 0$. Observe that

$$\begin{aligned} \Pr[Z \leq z] - \Pr[Z_n \leq z] &= \mathbb{E}[\Pr_i[Z \leq z] - \Pr_i[Z_n \leq z]] && \text{(LIE)} \\ &\geq p_n^+ \cdot t + (1 - p_n^+) \cdot O(t) && \text{(LIE and defn. of } p_n^+) \end{aligned}$$

However, since $Z_n \rightarrow_p Z$, we know that $Z_n \rightarrow_d Z$. By the Portmanteau theorem,

$$\Pr[Z_n \leq z] \rightarrow \Pr[Z \leq z] \quad \text{for almost all } z \in \mathbb{R}^9$$

Combining this with the previous inequality implies

$$p_n^+ \cdot t + (1 - p_n^+) \cdot O(t) \rightarrow 0$$

Since t does not depend on n , this can only be true when $p_n^+ \rightarrow 0$. By symmetry, we can apply this argument again to argue that

$$p_n^- := \Pr[\Pr_i[Z_n \leq z] - \Pr_i[Z \leq z] \geq t] \rightarrow 0$$

By the definition of convergence in probability, the fact that $p_n^+ \rightarrow 0$ and $p_n^- \rightarrow 0$ implies that $\Pr_i[Z_n \leq z] \rightarrow_p \Pr_i[Z \leq z]$.

2. The random variable $\mathbf{1}(Z_n \leq z)$ converges in probability to $\mathbf{1}(Z \leq z)$ for almost all z . This follows from the continuous mapping theorem. The distribution of Z can have at most finitely many atoms and the indicator function has only one discontinuity point, at z . Provided that z is not one of those atoms, the probability that Z matches a discontinuity point is zero. Therefore, we can

⁹The distribution of Z can have at most finitely many atoms. As long as z is not one of those atoms, the set of points $\{z' \in \mathbb{R} \mid z' \leq z\}$ is a continuity set for that distribution, and we can apply the Portmanteau theorem.

invoke the continuous mapping theorem for all but finitely many values of z .

It follows from these two results and the continuous mapping theorem that

$$(\Pr_i[Z_n \leq z] - \mathbf{1}(Z_n \leq z))^2 \rightarrow_p (\Pr_i[Z \leq z] - \mathbf{1}(Z \leq z))^2$$

for almost all $z \in \mathbb{R}$. Since convergence in probability implies convergence in distribution, we can apply the Portmanteau theorem to show that

$$f(Z_n, \cdot) \rightarrow f(Z, \cdot) \quad \text{a.e.}$$

Set $Z_n = \tilde{v}_{i,n}$ and $Z = E[v_j | \omega]$, where $Z_n \rightarrow_p Z$ by Claim 8. Then

$$f(\tilde{v}_{i,n}, \cdot) \rightarrow f(E[v_j | \omega], \cdot) \quad \text{a.e.}$$

We will use this fact momentarily, when applying the bounded convergence theorem.

Third, we characterize the limiting behavior of $\tilde{\rho}_n(s_i)$ as $n \rightarrow \infty$.

$$\begin{aligned} \lim_{n \rightarrow \infty} \lambda_n^{-1} \tilde{\rho}_n(s_i) &= - \sum_{x=1}^k \lim_{n \rightarrow \infty} \int_0^{\bar{v}} f(\tilde{v}_{i,n}, z) dz && \text{(definition of } \tilde{\rho}_n) \\ &= - \sum_{x=1}^k \int_0^{\bar{v}} f(E[v_j | \omega], z) dz && \text{(bounded convergence theorem)} \\ &= \rho(s_i) && \text{(definition of } \rho) \end{aligned}$$

It follows that there exists a threshold $N(s_i)$ such that, for any $n \geq N(s_i)$,

$$|\rho(s_i) - \lambda_n^{-1} \tilde{\rho}_n(s_i)| < \frac{\delta}{4} \quad (38)$$

Let $N = \max_{s_i \in \mathcal{S}} N(s_i)$. This quantity exists since \mathcal{S} is a finite set.

Fourth and finally, we show that agent i does not have a profitable deviation when $n \geq N$. Assume that all other agents continue to play according to an equilibrium of the auxiliary instance. Let s_i be the signal that she acquires in the equilibrium of the auxiliary instance. If agent i continues to acquire the same signal s_i , then her best response \hat{v}_i does not change. Suppose that agent i acquires a different signal s'_i . We only need to consider deviations to signals s'_i that were not available in

the auxiliary instance, which implies

$$\rho(\bar{s}_i) - \rho(s'_i) \geq \delta \quad (39)$$

where $\delta > 0$ was defined in equation (37). Since s_i was chosen in equilibrium of the auxiliary instance, it must have finite cost. By definition of \tilde{c} in (17), this implies

$$\rho(\bar{s}_i) = \rho(s_i) \quad (40)$$

Combining the last few inequalities, we find

$$\tilde{\rho}_n(s_i) - \tilde{\rho}_n(s'_i) > \lambda_n \rho(s_i) - \rho(s'_i) - \frac{\delta \lambda_n}{2} \quad (\text{equation (38)})$$

$$= \lambda_n \rho(\bar{s}_i) - \rho(s'_i) - \frac{\delta \lambda_n}{2} \quad (\text{equation (40)})$$

$$\geq \frac{\delta \lambda_n}{2} \quad (\text{inequality (39)})$$

This is the expected loss in transfers from the scoring rule SR^{CRPS} , which is a lower bound for the expected loss from the scoring rule SR. The maximum gain from deviating is at most the cost \bar{c} of the combined signal, plus the maximum payoff $v_H - v_L$ from the BDM mechanism (Lemma 4) or the VCG mechanism (Lemma 10). Therefore, a sufficient condition for the deviation to not be profitable is

$$\frac{\delta \lambda_n}{2} \geq \bar{c} + v_H - v_L$$

Since $\lambda_n \rightarrow \infty$, there is some threshold N' such that this condition holds for all $n \geq N'$. Then the agent has no profitable deviations whenever $n \geq \max\{N, N'\}$.

B.10 Proof of Lemma 5

Let $\bar{c}(\theta_i)$ be the maximum cost, given technology θ_i and across all base signals $s_i \in \mathcal{S}$, of the signal s'_i that combines s_i with a revealing signal. Assumptions 2 and 3 ensure that this is finite. Let $\bar{c} = \max_{\theta_i \in \Theta} \bar{c}(\theta_i)$.

Agent i reports the alternative \hat{x}_i that maximizes her expected value conditional on being pivotal. Claim 9 says that agent i conditioning on being pivotal does not meaningfully affect her expected value v_i .

Claim 9. Fix a constant $t > 0$. Then

$$\Pr[|E_{i,n}[v_i | Q_{i,n}] - E_{i,n}[v_i]| \geq t] \leq \frac{1}{t} \cdot \sqrt{\frac{|\mathcal{V}| \cdot |v_H - v_L| \cdot \bar{c}}{(\delta_n/n)^2 \cdot \lambda_n}}$$

Suppose x_i is the truthful report. Then, with high probability,

$$\begin{aligned} E_{i,n}[v_{ix_i}] &\geq E_{i,n}[v_{i\hat{x}_i}] && \text{(since } x_i \text{ is truthful)} \\ &\geq E_{i,n}[v_{i\hat{x}_i} | Q_{i,n}] - t && \text{(Claim 9)} \\ &\geq E_{i,n}[v_{ix_i} | Q_{i,n}] - t && \text{(since } \hat{x}_i \text{ is optimal)} \\ &\geq E_{i,n}[v_{ix_i}] - 2t && \text{(Claim 9)} \end{aligned}$$

This and the union bound ensures that

$$\Pr\left[\max_{i=1,\dots,n} \left| \max_{x \in \mathcal{X}} E_{i,n}[v_{ix}] - E_{i,n}[v_{i\hat{x}_i}] \right| \geq 2t\right] \leq \frac{n}{t} \cdot \sqrt{\frac{|\mathcal{V}| \cdot |v_H - v_L| \cdot \bar{c}}{(\delta_n/n)^2 \cdot \lambda_n}}$$

Setting t appropriately completes the proof.

B.11 Proof of Lemma 7

Define \bar{c} as in Lemma 5. Claim 10 says that agent i conditioning on being potentially pivotal does not meaningfully affect her expected value v_i .

Claim 10. Fix a constant $t > 0$. Then

$$\Pr[|E_{i,n}[v_i | Q_{i,n}] - E_{i,n}[v_i]| \geq t] \leq \frac{1}{t} \cdot \sqrt{\frac{|\mathcal{V}| \cdot |v_H - v_L| \cdot \bar{c}}{(\delta_n/n)^2 \cdot \lambda_n}}$$

Claim 11 is a useful fact that we use to prove Claims 12 and 13. It refers to G_n , which is the cumulative distribution function of the bias term $v_{0,n}$.

Claim 11. Fix constants a, b where $0 \leq b - a \leq \beta_n$. Then

$$\frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \left(\frac{b-a}{\beta_n} - e \left(\frac{b-a}{\beta_n} \right)^2 \right) \right) \leq G_n(b) - G_n(a) \leq \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \left(\frac{b-a}{\beta_n} + e \left(\frac{b-a}{\beta_n} \right)^2 \right) \right)$$

Claims 12 and 13 say that agent i 's expected value is roughly the same regardless of whether (i) she conditions on being potentially pivotal or (ii) she conditions on the

price $\tilde{p}_{i,n}$ being below some threshold z . Claim 12 conditions on agent i not being dictator, and Claim 13 drops this requirement. They refer to the following sequence:

$$\psi_n = \left(\frac{\beta_n + e(v_H - v_L)}{\beta_n - e(v_H - v_L)} \right)^2 - \left(\frac{\beta_n - e(v_H - v_L)}{\beta_n + e(v_H - v_L)} \right)^2$$

Claim 12. Fix constants z, z' where $v_L \leq z < z' \leq v_H$. Then

$$|\mathbb{E}_{i,n}[v_i \mid v_L \leq -n\tilde{v}_{i,n} \leq v_H] - \mathbb{E}_{i,n}[v_i \mid z \leq -n\tilde{v}_{i,n} \leq z']| \leq \psi_n$$

Claim 13. Fix constants z, z' where $v_L \leq z < z' \leq v_H$. Then

$$|\mathbb{E}_{i,n}[v_i \mid Q_{i,n}] - \mathbb{E}_{i,n}[v_i \mid z \leq \tilde{p}_{i,n} \leq z']| \leq \psi_n + O\left(\frac{1}{\beta_n} + \frac{\beta_n \delta_n}{n}\right)$$

Claim 14 says that agent i 's reported value will not be too far from her expected value. It relies heavily on Claims 10 and 11.

Claim 14. Fix a constant $t > 0$. Then

$$\Pr\left[|\hat{v}_{i,n} - \mathbb{E}_{i,n}[v_i]| \geq 2t + 2\psi_n + O\left(\frac{1}{\beta_n} + \frac{\beta_n \delta_n}{n}\right)\right] \leq \frac{1}{t} \sqrt{\frac{|\mathcal{V}| \cdot |v_H - v_L| \cdot \bar{c}}{(\delta_n/n)^2 \cdot \lambda_n}}$$

Let $t_n = n^{-1/2}$. Apply the union bound to Claim 14 to show that

$$\Pr\left[\max_{i=1}^n |\hat{v}_{i,n} - \mathbb{E}_{i,n}[v_i]| \geq t_n + \psi_n + O\left(\frac{1}{\beta_n} + \frac{\beta_n \delta_n}{n}\right)\right] \leq \frac{n}{t_n} \sqrt{\frac{|\mathcal{V}| \cdot |v_H - v_L| \cdot \bar{c}}{(\delta_n/n)^2 \cdot \lambda_n}}$$

We set the truthfulness parameters as follows. First, let

$$\phi_{1,n} = t_n + \psi_n + O\left(\frac{1}{\beta_n} + \frac{\beta_n \delta_n}{n}\right) = n^{-1/2} + O(n^{-1/2}) + O(n^{-1/2} + n^{-1}) = O(n^{-1/2})$$

Second, let

$$\phi_{2,n} = \frac{n}{t_n} \sqrt{\frac{|\mathcal{V}| \cdot |v_H - v_L| \cdot \bar{c}}{(\delta_n/n)^2 \cdot \lambda_n}} = n^{3/2} \sqrt{\frac{|\mathcal{V}| \cdot |v_H - v_L| \cdot \bar{c}}{n^{-3} \cdot n^7}} = O(n^{-1/2})$$

This completes the proof.

B.12 Proof of Claims 9 and 10

We want to bound the following term:

$$\begin{aligned}
& \left| \mathbb{E}_{i,n}[v_i \mid Q_{i,n}] - \mathbb{E}_{i,n}[v_i] \right| && (41) \\
& = \left| \sum_{y \in \mathcal{V}} y \cdot \Pr_{i,n}[v_i = y \mid Q_{i,n}] - \sum_{y \in \mathcal{V}} y \cdot \Pr_{i,n}[v_i = y] \right| && (\text{LIE}) \\
& \leq \sum_{y \in \mathcal{V}} y \cdot \left| \Pr_{i,n}[v_i = y \mid Q_{i,n}] - \Pr_{i,n}[v_i = y] \right| && (\text{triangle inequality}) \\
& = \sum_{y \in \mathcal{V}} y \cdot \left| \frac{\Pr_{i,n}[Q_{i,n} \mid v_i = y] \cdot \Pr_{i,n}[v_i = y]}{q_{i,n}} - \Pr_{i,n}[v_i = y] \right| && (\text{Bayes' rule}) \\
& = \sum_{y \in \mathcal{V}} y \cdot \Pr_{i,n}[v_i = y] \cdot \left| \frac{\Pr_{i,n}[Q_{i,n} \mid v_i = y]}{q_{i,n}} - 1 \right| \\
& = \sum_{y \in \mathcal{V}} y \cdot \Pr_{i,n}[v_i = y] \cdot \left| \frac{\Pr_{i,n}[Q_{i,n} \mid v_i = y] - q_{i,n}}{q_{i,n}} \right| \\
& = \frac{1}{q_{i,n}} \sum_{y \in \mathcal{V}} y \cdot \Pr_{i,n}[v_i = y] \cdot \left| \Pr_{i,n}[Q_{i,n} \mid v_i = y] - q_{i,n} \right| \\
& \leq \frac{\sqrt{|\mathcal{V}|}}{q_{i,n}} \sqrt{\sum_{y \in \mathcal{V}} (y \cdot \Pr_{i,n}[v_i = y] \cdot \left| \Pr_{i,n}[Q_{i,n} \mid v_i = y] - q_{i,n} \right|)^2} \\
& \hspace{20em} (\text{since } \|\alpha\|_1 \leq \sqrt{d} \cdot \|\alpha\|_2 \text{ for } \alpha \in \mathbb{R}^d) \\
& \leq \frac{\sqrt{|\mathcal{V}|}}{q_{i,n}} \sqrt{\sum_{y \in \mathcal{V}} \Pr_{i,n}[v_i = y] \cdot (y \cdot \left| \Pr_{i,n}[Q_{i,n} \mid v_i = y] - q_{i,n} \right|)^2} \quad (\text{since } \Pr[\cdot]^2 \leq \Pr[\cdot]) \\
& \leq \frac{\sqrt{|\mathcal{V}|} \cdot |v_H - v_L|}{q_{i,n}} \sqrt{\sum_{y \in \mathcal{V}} \Pr_{i,n}[v_i = y] \cdot (\Pr_{i,n}[Q_{i,n} \mid v_i = y] - q_{i,n})^2} \\
& \hspace{20em} (\text{since } y \leq |v_H - v_L|) \\
& = \frac{\sqrt{|\mathcal{V}|} \cdot |v_H - v_L|}{q_{i,n}} \sqrt{\mathbb{E}_{i,n}[(\Pr_{i,n}[Q_{i,n} \mid v_i] - q_{i,n})^2]} && (\text{LIE})
\end{aligned}$$

Focus on the term inside the square root.

$$\begin{aligned}
& \mathbb{E}_{i,n}[(\Pr_{i,n}[Q_{i,n} \mid v_i] - q_{i,n})^2] \\
& \leq \mathbb{E}_{i,n}[(\Pr_{i,n}[Q_{i,n} \mid v_i] - q_{i,n})^2 + (\Pr_{i,n}[\tilde{v}_{i,n} \notin Q_n \mid v_i] - (1 - q_{i,n}))^2] \quad (\text{since } (\cdot)^2 \geq 0)
\end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{E}_{i,n} \left[\Pr_{i,n}[Q_{i,n} | v_i]^2 + q_{i,n}^2 - 2q_{i,n} \cdot \Pr_{i,n}[Q_{i,n} | v_i] \right] \\
&\quad + \mathbb{E}_{i,n} \left[\Pr_{i,n}[\tilde{v}_{i,n} \notin Q_n | v_i]^2 + (1 - q_{i,n})^2 - 2(1 - q_{i,n}) \cdot \Pr_{i,n}[\tilde{v}_{i,n} \notin Q_n | v_i] \right] \\
&\leq \mathbb{E}_{i,n} \left[\Pr_{i,n}[Q_{i,n} | v_i]^2 + \Pr_{i,n}[\tilde{v}_{i,n} \notin Q_n | v_i]^2 \right] \\
&\quad + q_{i,n}^2 + (1 - q_{i,n})^2 - 2(1 - q_{i,n}) \cdot \Pr_{i,n}[\tilde{v}_{i,n} \notin Q_n] - 2q_{i,n} \cdot \Pr_{i,n}[Q_{i,n}] \quad (\text{LIE}) \\
&\leq \mathbb{E}_{i,n} \left[\Pr_{i,n}[Q_{i,n} | v_i]^2 + \Pr_{i,n}[\tilde{v}_{i,n} \notin Q_n | v_i]^2 \right] - q_{i,n}^2 - (1 - q_{i,n})^2 \quad (\text{defn. of } q_{i,n}) \\
&= \mathbb{E}_{i,n} \left[\max_{b_i} \mathbb{E}_{i,n} \left[\text{SR}^Q(b_i, \tilde{v}_i) | v_i \right] \right] - \max_{b_i} \mathbb{E}_{i,n} \left[\text{SR}^Q(b_i, \tilde{v}_i) \right] \quad (42)
\end{aligned}$$

The last equality follows from Example 1 of Gneiting and Raftery (2007). Next,

$$\begin{aligned}
(42) &\leq (42) + \mathbb{E}_{i,n} \left[\max_{b_i} \mathbb{E}_{i,n} \left[\left(\text{SR}(b_i, \tilde{v}_i) - \text{SR}^Q(b_i, \tilde{v}_i) \right) | v_i \right] \right] \\
&\quad - \max_{b_i} \mathbb{E}_{i,n} \left[\text{SR}(b_i, \tilde{v}_i) - \text{SR}^Q(b_i, \tilde{v}_i) \right] \quad (\mathbb{E}[\max(\cdot)] \geq \max \mathbb{E}[\cdot]) \\
&= \mathbb{E}_{i,n} \left[\max_{b_i} \mathbb{E}_{i,n} [\text{SR}(b_i, \tilde{v}_i) | v_i] \right] - \max_{b_i} \mathbb{E}_{i,n} [\text{SR}(b_i, \tilde{v}_i)] \quad (43)
\end{aligned}$$

The last equality follows from the fact that $\text{SR} - \text{SR}^Q$ is also a proper scoring rule. Recall that the agent can learn v_i by acquiring a revealing signal in addition to her current signal s_i . The cost of this is no greater than \bar{c} . In equilibrium, therefore,

$$\lambda_n \cdot \mathbb{E}[(43)] \leq \bar{c} \quad (44)$$

Combining all of the preceding arguments, we obtain

$$\begin{aligned}
\mathbb{E}[(41)]^2 &\leq \mathbb{E}[(41)^2] \quad (\text{Jensen's inequality}) \\
&\leq \mathbb{E} \left[\frac{|\mathcal{V}| \cdot |v_H - v_L|}{q_{i,n}^2} \cdot \mathbb{E}_{i,n} \left[(\Pr_{i,n}[Q_{i,n} | v_i] - q_{i,n})^2 \right] \right] \quad (\text{inequality (41)}) \\
&\leq \frac{|\mathcal{V}| \cdot |v_H - v_L|}{(\delta_n/n)^2} \cdot \mathbb{E} \left[\mathbb{E}_{i,n} \left[(\Pr_{i,n}[Q_{i,n} | v_i] - q_{i,n})^2 \right] \right] \quad (\text{since } q_{i,n} \geq \delta_n/n) \\
&\leq \frac{|\mathcal{V}| \cdot |v_H - v_L|}{(\delta_n/n)^2} \cdot \mathbb{E}[(42)] \quad (\text{inequality (42)}) \\
&\leq \frac{|\mathcal{V}| \cdot |v_H - v_L| \cdot \bar{c}}{(\delta_n/n)^2 \cdot \lambda_n} \quad (\text{inequality (44)})
\end{aligned}$$

Combining this with Markov's inequality gives us

$$\Pr[(41) \geq t] \leq \frac{\mathbb{E}[(41)]}{t} \leq \frac{1}{t} \cdot \sqrt{\frac{|\mathcal{V}| \cdot |v_H - v_L| \cdot \bar{c}}{(\delta_n/n)^2 \cdot \lambda_n}}$$

This completes the proof.

B.13 Proof of Claim 11

We rely on the fact that, for any constant $|w| \leq 1$,

$$e^w \in \left[1 + w - \frac{e \cdot w^2}{2}, 1 + w + \frac{e \cdot w^2}{2} \right] \quad (45)$$

There are three cases to consider.

1. Suppose $a \geq 0$. Then

$$\begin{aligned} G(b) - G(a) &= \frac{1}{2} \left(e^{-\frac{|a|}{\beta_n}} - e^{-\frac{|b|}{\beta_n}} \right) && \text{(since } v_{0,n} \sim \text{LAPLACE}(\beta_n)) \\ &= \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(e^{-\frac{|a|}{\beta_n} + \frac{|b|}{\beta_n}} - 1 \right) \\ &= \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(e^{\frac{b-a}{\beta_n}} - 1 \right) && \text{(since } b \geq a \geq 0) \\ &\in \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(\frac{b-a}{\beta_n} + \frac{e}{2} \left[-\left(\frac{b-a}{\beta_n} \right)^2, \left(\frac{b-a}{\beta_n} \right)^2 \right] \right) && \text{(bound (45))} \end{aligned}$$

2. Suppose $b < 0$. Then

$$\begin{aligned} G(b) - G(a) &= \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} - e^{-\frac{|a|}{\beta_n}} \right) && \text{(since } v_{0,n} \sim \text{LAPLACE}(\beta_n)) \\ &= \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(1 - e^{\frac{|b|}{\beta_n} - \frac{|a|}{\beta_n}} \right) \\ &= \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(1 - e^{\frac{a-b}{\beta_n}} \right) && \text{(since } a \leq b < 0) \\ &\in \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(\frac{b-a}{\beta_n} + \frac{e}{2} \left[-\left(\frac{b-a}{\beta_n} \right)^2, \left(\frac{b-a}{\beta_n} \right)^2 \right] \right) && \text{(bound (45))} \end{aligned}$$

3. Suppose $b \geq 0, a < 0$. Note that $-\beta_n \leq a - b \leq b + a \leq b < b - a \leq \beta_n$. Then

$$\begin{aligned}
& G(b) - G(a) \\
&= 1 - \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} + e^{-\frac{|a|}{\beta_n}} \right) && \text{(since } v_{0,n} \sim \text{LAPLACE}(\beta_n)) \\
&= \frac{1}{2} \left(2e^0 - e^{-\frac{|b|}{\beta_n}} - e^{-\frac{|a|}{\beta_n}} \right) \\
&= \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(2e^{\frac{|b|}{\beta_n}} - 1 - e^{\frac{|b|}{\beta_n} - \frac{|a|}{\beta_n}} \right) \\
&= \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(2e^{\frac{b}{\beta_n}} - 1 - e^{\frac{b+a}{\beta_n}} \right) && \text{(since } a < 0 \leq b) \\
&\in \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(\frac{b-a}{\beta_n} + e \left[-\left(\frac{b}{\beta_n} \right)^2 - \left(\frac{b+a}{\beta_n} \right)^2, \left(\frac{b}{\beta_n} \right)^2 + \left(\frac{b+a}{\beta_n} \right)^2 \right] \right) \\
&&& \text{(bound (45))} \\
&\subseteq \frac{1}{2} \left(e^{-\frac{|b|}{\beta_n}} \right) \left(\frac{b-a}{\beta_n} + e \left[-\left(\frac{b-a}{\beta_n} \right)^2, \left(\frac{b-a}{\beta_n} \right)^2 \right] \right) && \text{(since } a < 0 \leq b)
\end{aligned}$$

The bound in case 3 is the same as the bound in the statement of the claim, while the bounds in cases 1 and 2 are slightly tighter.

B.14 Proof of Claim 12

For convenience, for event $E \subseteq \mathbb{R}$, let

$$F^{z,z'}(E) = \Pr_{i,n} \left[\sum_{j \neq i} \hat{v}_{j,n} \in E \mid z \leq -v_{0,n} - \sum_{j \neq i} \hat{v}_{j,n} \leq z' \right]$$

Observe that

$$\begin{aligned}
& \mathbb{E}_{i,n} [v_i \mid z \leq -n\tilde{v}_{i,n} \leq z'] && (46) \\
&= \mathbb{E}_{i,n} \left[v_i \mid z \leq -v_{0,n} - \sum_{j \neq i} \hat{v}_{j,n} \leq z' \right] && \text{(defn. of } \tilde{v}_i) \\
&= \int_{\mathbb{R}} \mathbb{E}_{i,n} \left[v_i \mid z \leq -v_{0,n} - \sum_{j \neq i} \hat{v}_{j,n} \leq z', \sum_{j \neq i} \hat{v}_{j,n} = y \right] \cdot F^{z,z'}(dy) && \text{(LIE)}
\end{aligned}$$

$$= \int_{\mathbb{R}} \mathbb{E}_{i,n} \left[v_i \mid \sum_{j \neq i} \hat{v}_{j,n} = y \right] \cdot F^{z,z'}(dy) \quad (\text{since } v_i \perp v_{0,n})$$

Since the distribution $F^{z,z'}$ is not necessarily discrete or continuous, we need to refer to the definition of Lebesgue integration.

Definition 21. A function $h : \mathbb{R} \rightarrow \mathbb{R}$ is simple if there is a sequence $V_1^h, \dots, V_{K_h}^h \subseteq \mathcal{V}$ of disjoint $F^{z,z'}$ -measurable sets, and values $\alpha_1^h, \dots, \alpha_{K_h}^h \in \mathbb{R}_+$, such that

$$h(y) = \sum_{k=1}^{K_h} \alpha_k \cdot \mathbf{1}(y \in V_k)$$

Let \mathcal{H} be the set of simple functions.

By definition of Lebesgue integration,

$$(46) = \sup_{h \in \mathcal{H}} \sum_{k=1}^{K_h} \alpha_k^h \cdot F^{z,z'}(V_k^h) - \sup_{h' \in \mathcal{H}} \sum_{k=1}^{K_{h'}} \alpha_k^{h'} \cdot F^{z,z'}(V_k^{h'}) \quad (47)$$

$$\text{s.t. } h(y) \leq \left(\mathbb{E}_{i,n} \left[v_i \mid \sum_{j \neq i} \hat{v}_{j,n} = y \right] \right)^+ \text{ and } h'(y) \leq - \left(\mathbb{E}_{i,n} \left[v_i \mid \sum_{j \neq i} \hat{v}_{j,n} = y \right] \right)^-$$

Consider a sequence of simple functions $(h_l)_{l=1}^{\infty}$. For convenience, let $K_l = K_{h_l}$ and $V_k^l = V_k^{h_l}$. Define

$$\Delta_l^k := \sup V_k^l - \inf V_k^l \quad \text{and} \quad \Delta_l = \max_{k=1, \dots, K_l} \Delta_l^k$$

We eventually let this sequence converge to one of the suprema in equation (47), so it is without loss of generality to restrict attention to sequences that satisfy $\Delta_l \rightarrow 0$.

Applying Bayes' rule, we find that

$$F^{z,z'}(V_k^l) = \frac{\Pr_{i,n} \left[z \leq -v_{0,n} - \sum_{j \neq i} \hat{v}_{j,n} \leq z' \mid \sum_{j \neq i} \hat{v}_{j,n} \in V_k^l \right] \Pr_{i,n} \left[\sum_{j \neq i} \hat{v}_{j,n} \in V_k^l \right]}{\Pr_{i,n} \left[z \leq -v_{0,n} - \sum_{j \neq i} \hat{v}_{j,n} \leq z' \right]}$$

$$= \frac{\Pr_{i,n} \left[z \leq -v_{0,n} - \sum_{j \neq i} \hat{v}_{j,n} \leq z' \mid \sum_{j \neq i} \hat{v}_{j,n} \in V_k^l \right] \Pr_{i,n} \left[\sum_{j \neq i} \hat{v}_{j,n} \in V_k^l \right]}{\sum_{k'=1}^{K_l} \Pr_{i,n} \left[z \leq -v_{0,n} - \sum_{j \neq i} \hat{v}_{j,n} \leq z' \mid \sum_{j \neq i} \hat{v}_{j,n} \in V_{k'}^l \right] \Pr_{i,n} \left[\sum_{j \neq i} \hat{v}_{j,n} \in V_{k'}^l \right]} \quad (48)$$

where the second line follows from the law of iterated expectations. We focus on the conditional probabilities in (48). Recall that G_n is the cumulative distribution function of $v_{0,n}$. Let l be large enough that $\Delta^l < z' - z$. Then

$$\begin{aligned}
& \Pr_{i,n} \left[z \leq -v_{0,n} - \sum_{j \neq i} \hat{v}_{j,n} \leq z' \mid \sum_{j \neq i} \hat{v}_{j,n} \in V_k^l \right] & (49) \\
&= \Pr_{i,n} \left[z + \sum_{j \neq i} \hat{v}_{j,n} \leq -v_{0,n} \leq z' + \sum_{j \neq i} \hat{v}_{j,n} \mid \sum_{j \neq i} \hat{v}_{j,n} \in V_k^l \right] \\
&\leq \Pr_{i,n} \left[z + \inf V_k^l \leq -v_{0,n} \leq z' + \sup V_k^l \mid \sum_{j \neq i} \hat{v}_{j,n} \in V_k^l \right] \\
&= \Pr[z + \inf V_k^l \leq -v_{0,n} \leq z' + \sup V_k^l] & (\text{independence of } v_{0,n}) \\
&= \Pr[-z' - \sup V_k^l \leq v_{0,n} \leq -z - \inf V_k^l] \\
&= \Pr[v_{0,n} \leq -z - \inf V_k^l] - \Pr[-z' - \sup V_k^l \leq v_{0,n}] \\
&= G_n(-z - \inf V_k^l) - G_n(-z' - \sup V_k^l) & (\text{defn. of } G_n) \\
&\leq \frac{1}{2} \exp\left(-\frac{|-z - \inf V_k^l|}{\beta_n}\right) \left(\frac{z' - z + \Delta_l^k}{\beta_n} + e\left(\frac{z' - z + \Delta_l^k}{\beta_n}\right)^2 \right) & (\text{by Claim 11}) \\
&\leq \frac{1}{2} \exp\left(-\frac{|-z - \inf V_k^l|}{\beta_n}\right) \left(\frac{z' - z + \Delta_l}{\beta_n} + e\left(\frac{z' - z + \Delta_l}{\beta_n}\right)^2 \right) & (\text{defn. of } \Delta_l)
\end{aligned}$$

We can also obtain a lower bound as follows:

$$\begin{aligned}
(49) &\geq \Pr_{i,n} \left[z + \sup V_k^l \leq -v_{0,n} \leq z' + \inf V_k^l \mid \sum_{j \neq i} \hat{v}_{j,n} \in V_k \right] \\
&= \Pr[z + \sup V_k^l \leq -v_{0,n} \leq z' + \inf V_k^l] & (\text{independence of } v_{0,n}) \\
&\geq G_n(-z - \sup V_k^l) - G_n(-z' - \inf V_k^l) & (\text{defn. of } G_n) \\
&\geq \frac{1}{2} \exp\left(-\frac{|-z - \sup V_k^l|}{\beta_n}\right) \left(\frac{z' - z - \Delta_l^k}{\beta_n} - e\left(\frac{z' - z - \Delta_l^k}{\beta_n}\right)^2 \right) & (\text{by Claim 11}) \\
&\geq \frac{1}{2} \exp\left(-\frac{|-z - \sup V_k^l|}{\beta_n}\right) \left(\frac{z' - z - \Delta_l}{\beta_n} - e\left(\frac{z' - z - \Delta_l}{\beta_n}\right)^2 \right) \\
& & (\text{defn. of } \Delta_l \text{ and since } \beta_n \geq e(v_H - v_L))
\end{aligned}$$

Together with equation (48), these bounds imply that

$$\begin{aligned}
& \frac{\exp\left(-\frac{|-z-\sup V_k^l|}{\beta_n}\right)\left(\frac{z'-z-\Delta_l}{\beta_n} - e\left(\frac{z'-z-\Delta_l}{\beta_n}\right)^2\right)\Pr_{i,n}\left[\sum_{j\neq i}\hat{v}_{j,n}\in V_k^l\right]}{\sum_{k'=1}^{K_l}\exp\left(-\frac{|-z-\inf V_{k'}^l|}{\beta_n}\right)\left(\frac{z'-z+\Delta_l}{\beta_n} + e\left(\frac{z'-z+\Delta_l}{\beta_n}\right)^2\right)\Pr_{i,n}\left[\sum_{j\neq i}\hat{v}_{j,n}\in V_{k'}^l\right]} \\
& \leq (48) \\
& \leq \frac{\exp\left(-\frac{|-z-\inf V_k^l|}{\beta_n}\right)\left(\frac{z'-z+\Delta_l}{\beta_n} + e\left(\frac{z'-z+\Delta_l}{\beta_n}\right)^2\right)\Pr_{i,n}\left[\sum_{j\neq i}\hat{v}_{j,n}\in V_k^l\right]}{\sum_{k'=1}^{K_l}\exp\left(-\frac{|-z-\max V_{k'}^l|}{\beta_n}\right)\left(\frac{z'-z-\Delta_l}{\beta_n} - e\left(\frac{z'-z-\Delta_l}{\beta_n}\right)^2\right)\Pr_{i,n}\left[\sum_{j\neq i}\hat{v}_{j,n}\in V_{k'}^l\right]}
\end{aligned} \tag{50}$$

To make our notation more compact, define

$$\eta_{n,l}^+(z'-z) = \frac{\frac{z'-z+\Delta_l}{\beta_n} + e\left(\frac{z'-z+\Delta_l}{\beta_n}\right)^2}{\frac{z'-z-\Delta_l}{\beta_n} - e\left(\frac{z'-z-\Delta_l}{\beta_n}\right)^2} \quad \text{and} \quad \eta_{n,l}^-(z'-z) = \frac{\frac{z'-z-\Delta_l}{\beta_n} - e\left(\frac{z'-z-\Delta_l}{\beta_n}\right)^2}{\frac{z'-z+\Delta_l}{\beta_n} + e\left(\frac{z'-z+\Delta_l}{\beta_n}\right)^2}$$

and their respective limits as $l \rightarrow \infty$,

$$\eta_n^+(z'-z) = \frac{\beta_n + e(z'-z)}{\beta_n - e(z'-z)} \quad \text{and} \quad \eta_n^-(z'-z) = \frac{\beta_n - e(z'-z)}{\beta_n + e(z'-z)}$$

Similarly, define

$$\xi_{i,n,l}(k) = \frac{\exp\left(-\frac{|-z-\inf V_k^l|}{\beta_n}\right)\Pr_{i,n}\left[\sum_{j\neq i}\hat{v}_{j,n}\in V_k^l\right]}{\sum_{k'=1}^{K_l}\exp\left(-\frac{|-z-\max V_{k'}^l|}{\beta_n}\right)\Pr_{i,n}\left[\sum_{j\neq i}\hat{v}_{j,n}\in V_{k'}^l\right]}$$

With this notation, inequality (50) becomes

$$\eta_{n,l}^-(z'-z) \cdot \xi_{i,n,l}(k) \leq F^{z,z'}(V_k^l) \leq \eta_{n,l}^+(z'-z) \cdot \xi_{i,n,l}(k) \tag{51}$$

It follows that

$$\begin{aligned}
& \limsup_{l \rightarrow \infty} \left(\sum_{k=1, \dots, K_l} \alpha_k^l \cdot F^{v_L, v_H}(V_k^l) - \sum_{k=1, \dots, K_l} \alpha_k^l \cdot F^{z, z'}(V_k^l) \right) \tag{52} \\
&= \limsup_{l \rightarrow \infty} \sum_{k=1, \dots, K_l} \alpha_k^l \cdot F^{z, z'}(V_k^l) \left(\frac{F^{v_L, v_H}(V_k^l)}{F^{z, z'}(V_k^l)} - 1 \right) \\
&\leq \limsup_{l \rightarrow \infty} \sum_{k=1, \dots, K_l} \alpha_k^l \cdot F^{z, z'}(V_k^l) \left(\frac{\eta_{n,l}^+(v_H - v_L)}{\eta_{n,l}^-(z' - z)} - 1 \right) \tag{inequality (51)} \\
&= \left(\lim_{l \rightarrow \infty} \frac{\eta_{n,l}^+(v_H - v_L)}{\eta_{n,l}^-(z' - z)} - 1 \right) \left(\limsup_{l \rightarrow \infty} \sum_{k=1, \dots, K_l} \alpha_k^l \cdot F^{z, z'}(V_k^l) \right) \tag{property of lim sup} \\
&= \left(\frac{\eta_n^+(v_H - v_L)}{\eta_n^-(z' - z)} - 1 \right) \left(\limsup_{l \rightarrow \infty} \sum_{k=1, \dots, K_l} \alpha_k^l \cdot F^{z, z'}(V_k^l) \right) \tag{property of lim} \\
&\leq \left(\frac{\eta_n^+(v_H - v_L)}{\eta_n^-(v_H - v_L)} - 1 \right) \left(\limsup_{l \rightarrow \infty} \sum_{k=1, \dots, K_l} \alpha_k^l \cdot F^{z, z'}(V_k^l) \right) \tag{since } \eta_n^-(\cdot) \text{ decreasing}
\end{aligned}$$

Following similar reasoning, we find that

$$(52) \geq \left(\frac{\eta_n^-(v_H - v_L)}{\eta_n^+(v_H - v_L)} - 1 \right) \left(\limsup_{l \rightarrow \infty} \sum_{k=1, \dots, K_l} \alpha_k^l \cdot F^{z, z'}(V_k^l) \right) \tag{53}$$

Finally, we can return to the Lebesgue integral defined in equation (47). Let h_l be the sequence of simple functions that satisfies

$$\lim_{l \rightarrow \infty} \sum_{k=1}^{K_{h_l}} \alpha_k^{h_l} \cdot F^{z, z'}(V_k^{h_l}) = \sup_{h \in \mathcal{H}} \sum_{k=1}^{K_h} \alpha_k^h \cdot F^{z, z'}$$

Let h'_l be the sequence of simple functions that satisfies

$$\sup_{h' \in \mathcal{H}} \sum_{k=1}^{K_{h'}} \alpha_k^{h'} \cdot F^{v_L, v_H}(V_k^{h'}) = \lim_{l \rightarrow \infty} \sum_{k=1}^{K_{h'_l}} \alpha_k^{h'_l} \cdot F^{v_L, v_H}(V_k^{h'_l})$$

Next, observe that

$$\begin{aligned}
& \mathbb{E}_{i,n}[v_i \mid v_L \leq -n\bar{v}_{i,n} \leq v_H] - \mathbb{E}_{i,n}[v_i \mid z \leq -n\bar{v}_{i,n} \leq z'] \tag{54} \\
&= \sup_{h \in \mathcal{H}} \sum_{k=1}^{K_h} \alpha_k^h \cdot F^{v_L, v_H}(V_k^h) - \sup_{h \in \mathcal{H}} \sum_{k=1}^{K_h} \alpha_k^h \cdot F^{z, z'}(V_k^h) \\
&\quad - \sup_{h' \in \mathcal{H}} \sum_{k=1}^{K_{h'}} \alpha_k^{h'} \cdot F^{v_L, v_H}(V_k^{h'}) + \sup_{h' \in \mathcal{H}} \sum_{k=1}^{K_{h'}} \alpha_k^{h'} \cdot F^{z, z'}(V_k^{h'}) \tag{equation (47)} \\
&\geq \limsup_{l \rightarrow \infty} \sum_{k=1}^{K_{h_l}} \alpha_k^{h_l} \cdot F^{v_L, v_H}(V_k^{h_l}) - \lim_{l \rightarrow \infty} \sum_{k=1}^{K_{h_l}} \alpha_k^{h_l} \cdot F^{z, z'}(V_k^{h_l}) \\
&\quad - \lim_{l \rightarrow \infty} \sum_{k=1}^{K_{h'_l}} \alpha_k^{h'_l} \cdot F^{v_L, v_H}(V_k^{h'_l}) + \limsup_{l \rightarrow \infty} \sum_{k=1}^{K_{h'_l}} \alpha_k^{h'_l} \cdot F^{z, z'}(V_k^{h'_l}) \tag{defn. of sup and } h_l, h'_l) \\
&\geq \left(\frac{\eta_n^-(v_H - v_L)}{\eta_n^+(v_H - v_L)} - 1 \right) (v_H - v_L) - \left(\frac{\eta_n^+(v_H - v_L)}{\eta_n^-(v_H - v_L)} - 1 \right) (v_H - v_L) \\
&\tag{inequalities (52) and (53)} \\
&\geq \left(\frac{\eta_n^-(v_H - v_L)}{\eta_n^+(v_H - v_L)} - \frac{\eta_n^+(v_H - v_L)}{\eta_n^-(v_H - v_L)} \right) (v_H - v_L)
\end{aligned}$$

Following similar reasoning, we find that

$$(54) \leq \left(\frac{\eta_n^+(v_H - v_L)}{\eta_n^-(v_H - v_L)} - \frac{\eta_n^-(v_H - v_L)}{\eta_n^+(v_H - v_L)} \right) (v_H - v_L) \tag{55}$$

Combine inequalities (54) and (55), and then simplify, to obtain

$$\begin{aligned}
& \left| \mathbb{E}_{i,n}[v_i \mid z \leq -n\bar{v}_{i,n} \leq v_H] - \mathbb{E}_{i,n}[v_i \mid z \leq -n\bar{v}_{i,n} \leq z'] \right| \\
&\leq \left(\frac{\beta_n + e(v_H - v_L)}{\beta_n - e(v_H - v_L)} \right)^2 - \left(\frac{\beta_n - e(v_H - v_L)}{\beta_n + e(v_H - v_L)} \right)^2
\end{aligned}$$

This completes the proof.

B.15 Proof of Claim 13

Let z, z' be constants where $v_L \leq z < z' \leq v_H$. We begin by bounding the following:

$$\begin{aligned}
& \Pr_{i,n}[D = 0 \mid z \leq \tilde{p}_{i,n} \leq z'] && (56) \\
&= \frac{\Pr_{i,n}[z \leq \tilde{p}_{i,n} \leq z' \mid D = 0] \cdot \Pr_{i,n}[D = 0]}{\Pr_{i,n}[z \leq \tilde{p}_{i,n} \leq z']} && \text{(Bayes' rule)} \\
&= \frac{\Pr_{i,n}[z \leq \tilde{p}_{i,n} \leq z' \mid D = 0] \cdot \Pr[D = 0]}{\Pr_{i,n}[z \leq \tilde{p}_{i,n} \leq z' \mid D = i] \cdot \Pr[D = i] + \Pr_{i,n}[z \leq \tilde{p}_{i,n} \leq z' \mid D = 0] \cdot \Pr[D = 0]} \\
&&& \text{(LIE)} \\
&= \frac{\Pr_{i,n}[z \leq -n\tilde{v}_{i,n} \leq z' \mid D = 0] \cdot (1 - \delta_n)}{\Pr_{i,n}[z \leq p \leq z' \mid D = i] \cdot (\delta_n/n) + \Pr_{i,n}[z \leq -n\tilde{v}_{i,n} \leq z' \mid D = 0] \cdot (1 - \delta_n)} \\
&&& \text{(defn. of } D \text{ and } \tilde{p}_{i,n}\text{)} \\
&= \frac{\Pr_{i,n}[z \leq -n\tilde{v}_{i,n} \leq z'] \cdot (1 - \delta_n)}{\Pr[z \leq p \leq z'] \cdot (\delta_n/n) + \Pr_{i,n}[z \leq -n\tilde{v}_{i,n} \leq z'] \cdot (1 - \delta_n)} && \text{(since } D \perp (p, \tilde{v}_{i,n}\text{))} \\
&= \frac{\mathbb{E}_{i,n}\left[\Pr_{i,n}\left[z \leq -v_{0,n} - \sum_{i \neq j} \hat{v}_{j,n} \leq z', \sum_{j \neq i} \hat{v}_{j,n}\right]\right] \cdot (1 - \delta_n)}{\Pr[z \leq p \leq z'] \cdot (\delta_n/n) + \mathbb{E}_{i,n}\left[\Pr_{i,n}\left[z \leq -v_{0,n} - \sum_{i \neq j} \hat{v}_{j,n} \leq z', \sum_{j \neq i} \hat{v}_{j,n}\right]\right] \cdot (1 - \delta_n)} \\
&&& \text{(LIE and defn. of } \tilde{v}_{i,n}\text{)} \\
&= \frac{\mathbb{E}_{i,n}\left[G_n\left(-z - \sum_{j \neq i} \hat{v}_{j,n}\right) - G_n\left(-z' - \sum_{j \neq i} \hat{v}_{j,n}\right)\right] \cdot (1 - \delta_n)}{\Pr[z \leq p \leq z'] \cdot (\delta_n/n) + \mathbb{E}_{i,n}\left[G_n\left(-z - \sum_{j \neq i} \hat{v}_{j,n}\right) - G_n\left(-z' - \sum_{j \neq i} \hat{v}_{j,n}\right)\right] \cdot (1 - \delta_n)} \\
&&& \text{(defn. of } G_n\text{)}
\end{aligned}$$

The upper bound is

$$\begin{aligned}
(56) &\leq \frac{\frac{1}{2}\mathbb{E}_{i,n}\left[e^{-\frac{|-z - \sum_{j \neq i} \hat{v}_{j,n}|}{\beta_n}}\right] \left(\frac{z' - z}{\beta_n} + e^{\left(\frac{z' - z}{\beta_n}\right)^2}\right) \cdot (1 - \delta_n)}{\Pr[z \leq p \leq z'] \cdot (\delta_n/n) + \frac{1}{2}\mathbb{E}_{i,n}\left[e^{-\frac{|-z - \sum_{j \neq i} \hat{v}_{j,n}|}{\beta_n}}\right] \left(\frac{z' - z}{\beta_n} - e^{\left(\frac{z' - z}{\beta_n}\right)^2}\right) \cdot (1 - \delta_n)} \\
&&& \text{(Claim 11)} \\
&\leq \eta_n^+(z' - z) && \text{(since } \Pr[\cdot] \geq 0 \text{ and defn. of } \eta_n^+\text{)} \\
&\leq \eta_n^+(v_H - v_L) && \text{(since } \eta_n^+ \text{ increasing)}
\end{aligned}$$

The lower bound is

$$\begin{aligned}
(56) &\geq \frac{\frac{1}{2}\mathbb{E}_{i,n}\left[e^{-\frac{|-z-\sum_{j\neq i}\tilde{v}_{j,n}|}{\beta_n}}\right]\left(\frac{z'-z}{\beta_n}-e\left(\frac{z'-z}{\beta_n}\right)^2\right)\cdot(1-\delta_n)}{\Pr[z\leq p\leq z']\cdot(\delta_n/n)+\frac{1}{2}\mathbb{E}_{i,n}\left[e^{-\frac{|-z-\sum_{j\neq i}\tilde{v}_{j,n}|}{\beta_n}}\right]\left(\frac{z'-z}{\beta_n}+e\left(\frac{z'-z}{\beta_n}\right)^2\right)\cdot(1-\delta_n)} \\
&\hspace{20em} \text{(Claim 11)} \\
&= \frac{\frac{1}{2}\mathbb{E}_{i,n}\left[e^{-\frac{|-z-\sum_{j\neq i}\tilde{v}_{j,n}|}{\beta_n}}\right]\left(\frac{z'-z}{\beta_n}-e\left(\frac{z'-z}{\beta_n}\right)^2\right)\cdot(1-\delta_n)}{\frac{z'-z}{v_H-v_L}\cdot\frac{\delta_n}{n}+\frac{1}{2}\mathbb{E}_{i,n}\left[e^{-\frac{|-z-\sum_{j\neq i}\tilde{v}_{j,n}|}{\beta_n}}\right]\left(\frac{z'-z}{\beta_n}+e\left(\frac{z'-z}{\beta_n}\right)^2\right)\cdot(1-\delta_n)} \\
&\hspace{20em} \text{(since } p\sim\text{UNIFORM}[v_L, v_H]) \\
&= \frac{\beta_n - e(z' - z)}{\frac{2}{1-\delta_n}\mathbb{E}_{i,n}\left[e^{-\frac{|-z-\sum_{j\neq i}\tilde{v}_{j,n}|}{\beta_n}}\right]\cdot\frac{\beta_n^2}{v_H-v_L}\cdot\frac{\delta_n}{n} + \beta_n + e(z' - z)} \\
&\geq \eta_n^-(z' - z) \hspace{10em} \text{(since } \frac{a}{w+b} \geq \frac{a}{b} - w \cdot \frac{a}{b^2}) \\
&\quad - \frac{2}{1-\delta_n}\mathbb{E}_{i,n}\left[e^{-\frac{|-z-\sum_{j\neq i}\tilde{v}_{j,n}|}{\beta_n}}\right]\cdot\frac{\beta_n^2}{v_H-v_L}\cdot\frac{\delta_n}{n}\cdot\frac{\beta_n - e(z' - z)}{(\beta_n + e(z' - z))^2} \\
&= \eta_n^-(z' - z) - \Theta(1) \cdot O(1) \cdot \Theta(\beta_n^2) \cdot \frac{\delta_n}{n} \cdot \Theta(\beta_n^{-1}) \\
&= \eta_n^-(z' - z) - O\left(\frac{\beta_n\delta_n}{n}\right) \\
&\geq \eta_n^-(v_H - v_L) - O\left(\frac{\beta_n\delta_n}{n}\right) \hspace{10em} \text{(since } \eta_n^+ \text{ decreasing)}
\end{aligned}$$

For convenience, let $\zeta_{i,n}(z, z') = \Pr_{i,n}[D = 0 \mid z \leq \tilde{p}_{i,n} \leq z']$. Observe that

$$\begin{aligned}
&\mathbb{E}_{i,n}[v_i \mid z \leq \tilde{p}_{i,n} \leq z'] \hspace{15em} (57) \\
&= \zeta_{i,n}(z, z')\mathbb{E}_{i,n}[v_i \mid z \leq \tilde{p}_{i,n} \leq z', D = 0] + (1 - \zeta_{i,n}(z, z'))\mathbb{E}_{i,n}[v_i \mid z \leq \tilde{p}_{i,n} \leq z', D = i] \\
&\hspace{20em} \text{(LIE)} \\
&= \zeta_{i,n}(z, z')\mathbb{E}_{i,n}[v_i \mid z \leq -n\tilde{v}_{i,n} \leq z', D = 0] + (1 - \zeta_{i,n}(z, z'))\mathbb{E}_{i,n}[v_i \mid z \leq p \leq z', D = i] \\
&\hspace{20em} \text{(defn. of } \tilde{p}_{i,n}) \\
&= \zeta_{i,n}(z, z')\mathbb{E}_{i,n}[v_i \mid z \leq -n\tilde{v}_{i,n} \leq z'] + (1 - \zeta_{i,n}(z, z'))\mathbb{E}_{i,n}[v_i] \hspace{5em} ((p, D_i) \perp (v_i, \tilde{v}_{i,n}))
\end{aligned}$$

We want to bound

$$\begin{aligned}
& \left| \mathbb{E}_{i,n}[v_i \mid z \leq \tilde{p}_{i,n} \leq z'] - \mathbb{E}_{i,n}[v_i \mid Q_{i,n}] \right| & (58) \\
& = \left| \mathbb{E}_{i,n}[v_i \mid z \leq \tilde{p}_{i,n} \leq z'] - \mathbb{E}_{i,n}[v_i \mid v_L \leq \tilde{p}_{i,n} \leq v_H] \right| & (\text{defn. of } Q_{i,n}) \\
& = \left| \zeta_{i,n}(z, z') \mathbb{E}_{i,n}[v_i \mid z \leq -n\tilde{v}_{i,n} \leq z'] + (1 - \zeta_{i,n}(z, z')) \mathbb{E}_{i,n}[v_i] \right. \\
& \quad \left. - \zeta_{i,n}(v_L, v_H) \mathbb{E}_{i,n}[v_i \mid v_L \leq -n\tilde{v}_{i,n} \leq v_H] - (1 - \zeta_{i,n}(v_L, v_H)) \mathbb{E}_{i,n}[v_i] \right| & (\text{eq. (57)}) \\
& = \left| \zeta_{i,n}(v_L, v_H) (\mathbb{E}_{i,n}[v_i \mid z \leq -n\tilde{v}_{i,n} \leq z'] - \mathbb{E}_{i,n}[v_i \mid v_L \leq -n\tilde{v}_{i,n} \leq v_H]) \right. \\
& \quad \left. + (\zeta_{i,n}(z, z') - \zeta_{i,n}(v_L, v_H)) \mathbb{E}_{i,n}[v_i \mid z \leq -n\tilde{v}_{i,n} \leq z'] \right. \\
& \quad \left. + \mathbb{E}_{i,n}[v_i] (\zeta_{i,n}(v_L, v_H) - \zeta_{i,n}(z, z')) \right| \\
& \leq \left| \zeta_{i,n}(v_L, v_H) \right| \cdot \left| \mathbb{E}_{i,n}[v_i \mid z \leq -n\tilde{v}_{i,n} \leq z'] - \mathbb{E}_{i,n}[v_i \mid v_L \leq -n\tilde{v}_{i,n} \leq v_H] \right| \\
& \quad + \left| \zeta_{i,n}(z, z') - \zeta_{i,n}(v_L, v_H) \right| \cdot \left| \mathbb{E}_{i,n}[v_i \mid z \leq -n\tilde{v}_{i,n} \leq z'] \right| \\
& \quad + \left| \zeta_{i,n}(v_L, v_H) - \zeta_{i,n}(z, z') \right| \cdot \left| \mathbb{E}_{i,n}[v_i] \right| & (\text{properties of } |\cdot|) \\
& \leq \left| \zeta_{i,n}(v_L, v_H) \right| \cdot \psi_n + 2 \left| \zeta_{i,n}(z, z') - \zeta_{i,n}(v_L, v_H) \right| \cdot (v_H - v_L) & (\text{Claim 12}) \\
& \leq \psi_n + 2 \left(\eta_n^+(v_H - v_L) - \eta_n^-(v_H - v_L) + O\left(\frac{\beta_n \delta_n}{n}\right) \right) \cdot (v_H - v_L) & (\text{bounds on (56)}) \\
& \leq \psi_n + O\left(\frac{1}{\beta_n} + \frac{\beta_n \delta_n}{n}\right) & (\text{since } \frac{w+1}{w-1} - \frac{w-1}{w+1} = O(1/w))
\end{aligned}$$

This completes the proof.

B.16 Proof of Claim 14

Agent i 's report $\hat{v}_{i,n}$ influences her allocation and the part of her transfers. attributable to the VCG mechanism (if $D = 0$) or the BDM mechanism (if $D = i$).

I claim that her payoffs from her report can be represented as purchasing alternative $x = 1$ at price $\tilde{p}_{i,n}$. This is clearly true if agent i is dictator ($D = i$) and participates in the BDM mechanism. Suppose instead that there is no dictator ($D = 0$) and agent i participates in the VCG mechanism. There are two cases.

1. If $\tilde{v}_{i,n} < 0$, then agent i chooses between (i) alternative $x = 1$ and paying $-n\tilde{v}_{i,n}$ and (ii) alternative $x = 0$ and paying nothing.
2. If $\tilde{v}_{i,n} \geq 0$, then she chooses between (i) alternative $x = 1$ and paying nothing and (ii) alternative $x = 0$ and paying $n\tilde{v}_{i,n}$. This is strategically equivalent to choosing between (i) $x = 1$ and paying $n\tilde{v}_{i,n}$ and (ii) $x = 0$ and paying nothing.

Next, we show that the agent's optimal report $\hat{v}_{i,n}$ is close to her expected value $E_{i,n}[v_i]$. Fix a constant $t > 0$. For the rest of this proof, we condition agent i having a signal realization such that

$$|E_{i,n}[v_i \mid v_L \leq -n\tilde{v}_{i,n} \leq v_H] - E_{i,n}[v_i \mid z \leq -n\tilde{v}_{i,n} \leq z']| \leq t + \psi_n + O\left(\frac{1}{\beta_n} + \frac{\beta_n \delta_n}{n}\right)$$

It follows from the triangle inequality, Claim 10, and Claim 13 that this holds with high probability. Let $z < E_{i,n}[v_i] - 2t'$ where

$$t' = t\psi_n + O\left(\frac{1}{\beta_n} + \frac{\beta_n \delta_n}{n}\right)$$

Agent i 's expected payoff from reporting $E_{i,n}[v_i] - 2t'$ is

$$\begin{aligned} & E_{i,n}[\mathbf{1}(v_L \leq \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t') \cdot (v_i - \tilde{p}_{i,n})] \\ &= E_{i,n}[\mathbf{1}(z < \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t') \cdot (v_i - \tilde{p}_{i,n})] \\ & \quad + E_{i,n}[\mathbf{1}(v_L \leq \tilde{p}_{i,n} \leq z) \cdot (v_i - \tilde{p}_{i,n})] \end{aligned}$$

This is greater than her expected payoff from reporting z if the following is positive:

$$\begin{aligned} & E_{i,n}[\mathbf{1}(z < \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t') \cdot (v_i - \tilde{p}_{i,n})] \\ &= \Pr_{i,n}[z < \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t'] \cdot E_{i,n}[v_i - \tilde{p}_{i,n} \mid z < \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t'] \quad (\text{LIE}) \\ &\geq \Pr_{i,n}[z < \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t'] \cdot (E_{i,n}[v_i \mid z < \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t'] - E_{i,n}[v_i] + 2t') \\ &\geq \Pr_{i,n}[z < \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t'] \cdot (E_{i,n}[v_i] - t' - E_{i,n}[v_i] + 2t') \quad (59) \\ &\geq \Pr_{i,n}[z < \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t'] \cdot t' \\ &\geq \frac{\delta_n}{n} \cdot \frac{E_{i,n}[v_i] - 2t' - z}{v_H - v_L} \cdot t' \quad (\text{defn. of } \tilde{p}_{i,n}, D \text{ and since } p \sim \text{UNIFORM}[v_L, v_H]) \\ &> 0 \quad (\text{defn. of } z) \end{aligned}$$

Line (59) follows from the definition of t' , provided that

$$E_{i,n}[v_i \mid z \leq \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t'] = E_{i,n}[v_i \mid z < \tilde{p}_{i,n} \leq E_{i,n}[v_i] - 2t']$$

This follows from the fact that, conditional on any value $D \in \{0, i\}$ and $\sum_{j \neq i} \hat{v}_{j,n}$, $\tilde{p}_{i,n}$ is continuously-distributed and therefore puts zero probability on any given point z .

Altogether, we have shown that the agent i 's optimal report $\hat{v}_{i,n} \geq E_{i,n}[v_i] - 2t'$. To complete the proof, we use a similar argument to show that $\hat{v}_{i,n} \leq E_{i,n}[v_i] + 2t'$.

B.17 Proof of Lemma 11

To reduce notation, we implicitly condition on the probability $1 - \delta_n$ event that no agent is chosen as dictator. This does not affect our conclusions since $\delta_n \rightarrow 0$.

The difference between the expected welfare given mechanism $(\mathbf{x}^n, \mathbf{t}^n)$ and expected welfare under the ex-ante optimal alternative is

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E}[v_i | \omega] \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) - \mathbb{E}[v_i | \omega] \cdot \mathbf{1}(\mathbb{E}[v_i] \geq 0) \right] \quad (60) \\ & \geq \mathbb{E} \left[\left(\frac{1}{n} \sum_{j=1}^n v_j \right) \cdot \mathbf{1} \left(v_{0x} + \sum_{j=1}^n \hat{v}_{jx} \geq 0 \right) - \mathbb{E}[v_i | \omega] \cdot \mathbf{1}(\mathbb{E}[v_i] \geq 0) \right] - o(1) \\ & \hspace{20em} \text{(Hoeffding's inequality)} \end{aligned}$$

We want to show that (60) $\geq -o(1)$. The key step is the following claim.

Claim 15. *The following inequality holds in every equilibrium $(\mathbf{s}^n, \hat{\mathbf{v}}^n, \hat{\mathbf{b}}^n)$:*

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{1}{n} \sum_{j=1}^n v_j \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) \right] \\ & \geq \mathbb{E} \left[\frac{1}{n} \cdot \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) \right] - o(1) \end{aligned}$$

Claim 15 relies critically on Lemmas 7 and 8. We leave the proof to Section B.18.

Given this claim, we observe that

$$(60) \geq \mathbb{E} \left[\frac{1}{n} \cdot \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) - \mathbb{E}[v_i | \omega] \cdot \mathbf{1}(\mathbb{E}[v_i] \geq 0) \right] - o(1)$$

At this point, there are two cases to consider.

1. Suppose that $E[v_i] < 0$. Then

$$\begin{aligned}
(60) &\geq \mathbb{E} \left[\frac{1}{n} \cdot \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) \right] - o(1) \\
&\geq -o(1) \qquad \qquad \qquad (\text{since } \mathbb{E}[X \cdot \mathbf{1}(X \geq 0)] \geq 0)
\end{aligned}$$

2. Suppose that $E[v_i] \geq 0$. Then

$$\begin{aligned}
(60) &\geq \mathbb{E} \left[\frac{1}{n} \cdot \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) - \mathbb{E}[v_i \mid \omega] \right] - o(1) \\
&\geq \mathbb{E} \left[\frac{1}{n} \cdot \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \right) - \mathbb{E}[v_i \mid \omega] \right] - o(1) \\
&\qquad \qquad \qquad (\text{since } \mathbb{E}[X \cdot \mathbf{1}(X \geq 0)] \geq \mathbb{E}[X]) \\
&= \mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n \hat{v}_{j,n} - \mathbb{E}[v_i \mid \omega] \right] - o(1) \qquad (\text{since } \mathbb{E}[v_{0,n}] = 0) \\
&\geq \mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n \mathbb{E}_{j,n}[v_j] - \mathbb{E}[v_i \mid \omega] \right] - O(\phi_{1,n} + (v_H - v_L) \cdot \phi_{2,n}) - o(1) \\
&\qquad \qquad \qquad (\text{Lemma 7}) \\
&= \mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n \mathbb{E}_{j,n}[v_j] - \mathbb{E}[v_i \mid \omega] \right] - o(1) \qquad (\text{since } \phi_{1,n} \rightarrow 0 \text{ and } \phi_{2,n} \rightarrow 0) \\
&= -o(1) \qquad \qquad \qquad (\text{LIE})
\end{aligned}$$

In both cases, the difference (60) is at least $-o(1)$. This completes the proof.

B.18 Proof of Claim 15

Recall that the scoring rule SR evaluates, among other things, agent i 's ability to predict the random variable

$$\mathbf{1}(\tilde{v}_{i,n} \geq 0) = \mathbf{1} \left(v_{0,n} + \sum_{j \neq i} \hat{v}_{j,n} \geq 0 \right)$$

Consider the following event:

$$\begin{aligned}
& \mathbf{1}\left(v_{0,n} + \sum_{j \neq i} \hat{v}_{j,n} \geq 0\right) \neq \mathbf{1}\left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0\right) \\
& \implies v_{0,n} + \sum_{j \neq i} \hat{v}_{j,n} \in [-v_H, -v_L] \\
& \implies v_{0,n} \in [-v_H - A_{i,n}, -v_L - A_{i,n}] \quad (\text{where } A_{i,n} = \sum_{j \neq i} \hat{v}_{j,n})
\end{aligned}$$

We want to bound the probability of this event. Observe that

$$\begin{aligned}
& \Pr[v_{0,n} \in [-v_H - A_{i,n}, -v_L - A_{i,n}]] \\
& = \mathbb{E}[\Pr[v_{0,n} \in [-v_H - A_{i,n}, -v_L - A_{i,n}] \mid A_{i,n}]] \quad (\text{LIE}) \\
& \leq \frac{v_H - v_L}{2\beta_n} \quad (\text{since } v_{0,n} \sim \text{LAPLACE}(0, \beta_n))
\end{aligned}$$

It follows from the preceding arguments that

$$\mathbf{1}\left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0\right) \geq \mathbf{1}\left(v_{0,n} + \sum_{j \neq i} \hat{v}_{j,n} \geq 0\right) - O_p\left(\frac{v_H - v_L}{2\beta_n}\right) \quad (61)$$

We will use this observation later on to bound the following expression.

$$\begin{aligned}
& \mathbb{E}\left[\left(\frac{1}{n} \sum_{j=1}^n v_j\right) \cdot \mathbf{1}\left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0\right)\right] \quad (62) \\
& = \mathbb{E}\left[\left(\frac{1}{n} \sum_{j=1}^n \mathbb{E}_{j,n}[v_j]\right) \cdot \mathbf{1}\left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0\right)\right] \\
& \quad + \mathbb{E}\left[\left(\frac{1}{n} \sum_{j=1}^n \epsilon_{j,n}\right) \cdot \mathbf{1}\left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0\right)\right] \quad (\text{defn. of } \epsilon_{j,n})
\end{aligned}$$

There are two terms on the right-hand side of equation (62). The first one is

$$\mathbb{E}\left[\left(\frac{1}{n} \sum_{j=1}^n \mathbb{E}_{j,n}[v_j]\right) \cdot \mathbf{1}\left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0\right)\right] \quad (63)$$

$$\begin{aligned}
&\geq \mathbb{E} \left[\left(\frac{1}{n} \sum_{j=1}^n \hat{v}_j \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) \right] - O(\phi_{1,n} + (v_H - v_L)\phi_{2,n}) \quad (\text{by Lemma 7}) \\
&\geq \mathbb{E} \left[\left(\frac{1}{n} \sum_{j=1}^n \hat{v}_j \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) \right] - o(1) \quad (\text{since } \phi_{1,n} \rightarrow 0, \phi_{2,n} \rightarrow 0) \\
&\geq \mathbb{E} \left[\left(\frac{v_{0,n}}{n} - \frac{v_{0,n}}{n} \cdot \mathbf{1}(v_{0,n} \geq 0) + \frac{1}{n} \sum_{j=1}^n \hat{v}_j \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) \right] - o(1) \\
&\hspace{20em} (\text{since } v_{0,n} \cdot \mathbf{1}(v_{0,n} \geq 0) \geq v_{0,n}) \\
&\geq \mathbb{E} \left[\left(\frac{v_{0,n}}{n} + \frac{1}{n} \sum_{j=1}^n \hat{v}_j \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) - \frac{v_{0,n}}{n} \cdot \mathbf{1}(v_{0,n} \geq 0) \right] - o(1) \\
&\hspace{20em} (\text{since } v_{0,n} \cdot \mathbf{1}(v_{0,n} \geq 0) \geq 0) \\
&= \mathbb{E} \left[\frac{1}{n} \cdot \left(v_{0,n} + \sum_{j=1}^n \hat{v}_j \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) \right] - \frac{\beta_n}{n} - o(1) \\
&\hspace{20em} (\text{since } v_{0,n} \sim \text{LAPLACE}(0, \beta_n)) \\
&= \mathbb{E} \left[\frac{1}{n} \cdot \left(v_{0,n} + \sum_{j=1}^n \hat{v}_j \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) \right] - o(1) \quad (\text{since } \beta_n/n \rightarrow 0)
\end{aligned}$$

The second term on the right-hand side of equation (62) is

$$\begin{aligned}
&\mathbb{E} \left[\left(\frac{1}{n} \sum_{j=1}^n \epsilon_{j,n} \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{j=1}^n \hat{v}_{j,n} \geq 0 \right) \right] \tag{64} \\
&\geq \mathbb{E} \left[\left(\frac{1}{n} \sum_{j=1}^n \epsilon_{j,n} \right) \cdot \mathbf{1} \left(v_{0,n} + \sum_{k \neq j} \hat{v}_{k,n} \geq 0 \right) \right] - O \left(\frac{(v_H - v_L)^2}{2\beta_n} \right) \quad (\text{inequality (61)}) \\
&\geq -O(\lambda_n^{-1/2}) - O \left(\frac{(v_H - v_L)^2}{2\beta_n} \right) \tag{Lemma 8} \\
&= o(1) \hspace{15em} (\text{since } \lambda_n \rightarrow \infty, \phi_{1,n} \rightarrow 0, \beta_n \rightarrow \infty)
\end{aligned}$$

Plugging inequalities (63) and (64) into equation (62) completes the proof.

References

- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359–378.
- Jensen, J. L. W. V. (1905). Om konvekse funktioner og uligheder imellem middelveerdier. *Nyt tidsskrift for matematik*, 16, 49–68.
- Kim, J. [Jeankyung], & Pollard, D. (1990). Cube Root Asymptotics. *The Annals of Statistics*, 18(1), 191–219.