Comparisons of Signals

Benjamin Brooks, Alexander Frankel, and Emir Kamenica*

March 2023

Abstract

A (Blackwell) experiment specifies the joint distribution of truth and the data generated by the experiment. A signal specifies the joint distribution of truth, the data generated by the signal, and the data generated by any other signal. Describing two experiments does not determine their joint informational content; describing two signals does. Blackwell (1953) studied (equivalent) comparisons of experiments; he characterized when one experiment is more valuable than another regardless of the preferences of the agent. We study (various, non-equivalent) comparisons of signals. Among other comparisons, we characterize when one signal is more valuable than another regardless of the preferences of the agent and regardless of what other information the agent may have. We show this comparison is equivalent to a new condition, termed reveal-or-refine, which says that for every piece of data that could be generated by the more valuable signal, either that data reveals the truth, or it refines the data generated by the less valuable signal.

JEL classification: C70; D82; D83; D85
Keywords: Experiments, signals, Blackwell order, sufficiency

*Brooks: Department of Economics, University of Chicago, babrooks@uchicago.edu. Frankel: Booth School of Business, University of Chicago, afrankel@uchicago.edu. Kamenica: Booth School of Business, University of Chicago, emir@uchicago.edu. We thank Piotr Dworczak, Drew Fudenberg, Elliot Lipnowski, Lones Smith, and many seminar audiences for comments. Karen Wu provided valuable research assistance.

1 Introduction

Economic theory has long been concerned with comparing the value of information sources. Blackwell (1953), for instance, gives conditions under which one source is more valuable than another source, regardless of preferences. Most prior work, however, considers the value of an information source in isolation, in absence of other – potentially correlated – information sources. Concretely, suppose we wish to judge whether a subscription to the *New York Times* (*NYT*) is more valuable than a subscription to the *Washington Post* (*WP)*,[1] regardless of the reader's interests. Blackwell's analysis tells us how to make this judgment, namely by comparing the distributions of beliefs induced by reading the *NYT* vs. the *WP*. However, a conclusion based on this procedure might be mistaken if the reader already has an existing subscription to a newspaper. Most obviously, if the reader already subscribes to the *NYT*, a subscription to the *WP* is likely more valuable than a duplicate subscription to the *NYT.* More subtly, a subscription to a third newspaper such as the *Wall Street Journal* (*WSJ*) might flip the comparison and make the *WP* more valuable than the *NYT*, either because the *WSJ* and the *NYT* report similarly, or because the content of the *WP* is somehow complementary to the *WSJ.* Could there be a way to establish that the one newspaper is more valuable than another, no matter what existing subscriptions the reader might have?

In this paper, we study the general version of this question. We derive comparisons of information sources that are robust to the presence of pre-existing information. (These comparisons also turn out to be robust to subsequent endogenous acquisition of additional information.)

Formally, Blackwell models an information source as an *experiment*: a collection of possible outcomes and a conditional distribution of outcomes given the state. Blackwell's foundational result is that experiment A is more valuable than experiment B, regardless of the decision problem, if and only if the distribution of beliefs induced by observing A is a mean-preserving spread of that induced by observing B. Importantly, however, an experiment does not specify how observations from one information source are correlated with observations from other sources. This suffices for Blackwell's purpose, since he implicitly assumes that the two experiments being compared are the only information sources potentially available to the agent.[2]

[1] Throughout this paper we focus on the instrumental value of information sources, ignoring the possibility that reading a newspaper might also provide entertainment (Ely *et al.*, 2015).

[2] Blackwell's result extends to situations where there might be other sources of information but only if those sources

In order to capture the joint informational content of multiple sources, we follow Green and Stokey (1978) in modeling an information source as a *signal*: a partition of an expanded state space $\Omega \times X$ that distinguishes payoff-relevant states (Ω) from those that govern the realization of observations conditional on the state (X). A signal induces an experiment, but it also pins down its correlation with other signals. In particular, the information generated by observing signals A and B is given by the join of the partitions, denoted $A \vee B$.

We say that signal A *Blackwell dominates* signal B if the experiment induced by A is more valuable than an experiment induced by B, regardless of agent's preferences.[3] We then introduce the strong Blackwell order, defined as follows: signal A *strongly Blackwell dominates* signal B if for every signal C, $A \vee C$ Blackwell dominates $B \vee C$. In other words, we extend Blackwell's agnosticism about the agent's preferences to agnosticism about what other information the agent might have.

Our first theorem characterizes the strong Blackwell order. Say that signal A *reveals-or-refines* signal B if every signal realization of A either (i) occurs in only one state (and thus "reveals" the state), or (ii) is a subset of some signal realization of B (and thus "refines" B, pinning down what information is observed by B).[4] We show that A strongly Blackwell dominates B if and only if A reveals-or-refines B.

Once we are within the formalism which allows for combining sources of information, other comparisons of signals become natural. Say that signal A is *sufficient* for signal B if the experiment induced by A is the same as the experiment induced by $A \vee B$. In other words, B does not contain additional information about the state beyond that contained in A. Thus, for any agent with access A, the marginal value of B is zero. As with strong Blackwell, we show how to determine whether one signal is sufficient for another. Sufficiency turns out to be a distinct relation; it is implied by strong Blackwell, and it implies Blackwell. We also illustrate that, perhaps surprisingly, sufficiency is not transitive.

We also consider another comparison of information sources. Say that signal A *martingale*

are conditionally independent of the two experiments being compared. This reflects the fact that the Blackwell comparison of experiments does not depend on prior beliefs.

[3] For most of our analysis, we treat the prior belief as fixed. As is well known, if one source of information Blackwell dominates another for some (interior) prior, then it does so for all priors. We discuss this at greater length in Section 4.2.

[4] So, if Alice observes A and Bob observes B, either Alice's first-order beliefs (about the state) or her second-order beliefs (about Bob's beliefs) are degenerate.

dominates signal B if an agent who forms some posterior belief μ after observing B, thinks that an agent who observes A will, in expectation, also hold belief μ. If B were much more informative than A, there would no reason to think this; for example, observing B might reveal the state (and thus result in a degenerate belief) while observing A might never reveal the state (and thus cannot lead to a degenerate belief in expectation). We discuss at greater length in Section 4.2 the sense in which the martingale relation captures a notion of being "more informative." We also show that sufficiency implies martingale which in turn implies Blackwell. As with sufficiency, the martingale relation is not transitive.

Like the Blackwell comparison, the sufficiency and martingale relations implicitly presume that the agent has no additional sources of information. Analogously to our strong Blackwell order, it is possible to strengthen sufficiency, or martingale, or in fact any relation on signals, to reflect robustness to other information. Given relation $\mathcal{P}$ on signals, let the strengthening of $\mathcal{P}$, denoted $\overline{\mathcal{P}}$, be defined as: $A\overline{\mathcal{P}}B$ if for any C, $(A \vee C)\,\mathcal{P}\,(B \vee C)$. Two properties of strengthening are worth noting, namely monotonicity (if $\mathcal{P}$ implies $\mathcal{P}'$ then $\overline{\mathcal{P}}$ implies $\overline{\mathcal{P}'}$) and idempotence ($\overline{\overline{\mathcal{P}}} = \overline{\mathcal{P}}$).

These two properties, coupled with our earlier observations, immediately yield the characterization of strong sufficiency and strong martingale. Since strong Blackwell implies sufficiency and martingale, we have that strong Blackwell also implies strong sufficiency and strong martingale. But since sufficiency and martingale imply Blackwell, strong sufficiency and strong martingale also imply strong Blackwell. Thus, strong sufficiency, strong martingale, and strong Blackwell are all equivalent (and characterized by reveal-or-refine). The overarching message is that reveal-or-refine is a natural ranking of signals that is robust to the presence of additional information.

Our paper is most closely related to the literature on ordinal comparisons of the ex-ante value of information sources, starting with Blackwell (1951).[5] Much of the developments in this line of research focus on ways to weaken the Blackwell order. Lehmann (1988), Persico (2000), and Athey and Levin (2018) consider comparisons that apply to a subset of decision problems and/or a subset of experiments. Moscarini and Smith (2002) and Mu *et al.* (2021b) compare the values of large numbers of independent draws of different experiments.

Another closely related literature focuses on the joint informational content of multiple infor-

[5]A smaller literature considers the ex-post value of information (Frankel and Kamenica (2019); Frankel and Kasy (2022)).

mation sources.[6] Börgers *et al.* (2013) consider the question of when signals are complements or substitutes. Gentzkow and Kamenica (2017a,b) consider the impact of competition when multiple senders provide potentially correlated signals in an attempt to influence a receiver. Liang and Mu (2020) and Liang *et al.* (2022) consider acquisition of potentially complementary information sources. Brooks *et al.* (2022) analyze the relationship between the comparison of information sources conceptualized as experiments, according to the Blackwell order, and the comparison of information sources conceptualized as signals, according to refinement, sufficiency, and martingale.[7] Specifically, they ask when a collection of Blackwell-ordered experiments can be induced by a collection of refinement-, sufficiency-, or martingale-ordered signals.

2 Signals and experiments

There is a finite state space Ω and an interior prior $\mu_0 \in \Delta\Omega$. We denote a typical state by ω.

An *experiment* τ is a distribution of beliefs – i.e., an element of $\Delta\Delta\Omega$ – that has finite support and satisfies $\mathbb{E}_\tau[\mu] = \mu_0$. (An alternative definition of an experiment is a map from Ω to distributions over signal realizations, but as is common, we simply identify each experiment with the distribution of beliefs it induces.) We write $\tau \succsim \tau'$ if τ is a mean-preserving spread of τ'.

A *signal* π is a finite partition of $\Omega \times [0,1]$ s.t. $\pi \subset S$, where S is the set of non-empty Lebesgue-measurable subsets of $\Omega \times [0,1]$ (Green and Stokey, 1978; Gentzkow and Kamenica, 2017a). An element $s \in S$ is a *signal realization*. The interpretation of this formalism is that a random variable x, drawn uniformly from $[0,1]$, determines the signal realization conditional on the state. Thus, the conditional probability of s given ω is $p^\omega(s) = \lambda(\{x | (\omega, x) \in s\})$ where $\lambda(\cdot)$ denotes the Lebesgue measure. Observing signal realization s induces the posterior μ_s.[8]

Given signal π, let $\tilde{s}_\pi$ be the associated S-valued random variable on $\Omega \times [0,1]$ induced by π.[9] Let $\tilde{\mu}_\pi \equiv \mu_{\tilde{s}_\pi}$ denote the associated belief-valued random variable that reflects the posterior

[6]Just as additional sources of information can alter the value of a signal, additional sources of income can alter the value of a monetary gamble. Mu *et al.* (2021a) explore how a decision maker's preferences over monetary gambles can depend on background risk, i.e., independent uncertainty over income.

[7]As we discuss in Section 4.2, the notion of martingale relation introduced in Brooks *et al.* (2022) is slightly different from the one we study here.

[8]The posterior probability of ω given s is $\mu_s^\omega \equiv \frac{p^\omega(s)\mu_0^\omega}{\sum_{\omega' \in \Omega} p^{\omega'}(s)\mu_0^{\omega'}}$ as long as the unconditional probability of s, $\sum_{\omega' \in \Omega} p^{\omega'}(s)\mu_0^{\omega'}$, is strictly positive. The specification of μ_s when s has zero probability is irrelevant for our results.

[9]Recall that an S-valued random variable on $\Omega \times [0,1]$ is simply a function from $\Omega \times [0,1]$ to S; $\tilde{s}_\pi$ maps each

induced by observing the realization from π. Finally, let $\langle\pi\rangle$ denote the distribution of $\tilde{\mu}_\pi$, i.e., the experiment induced by signal π. If $\langle\pi\rangle = \langle\pi'\rangle$, we say that π and π' are *Blackwell equivalent* and write $\pi \sim \pi'$.

We denote the set of all signals by Π. We say π *refines* π' and write $\pi\mathcal{R}\pi'$ if every element of π is a subset of some element of π'.[10] If $\pi\mathcal{R}\pi'$, an agent who observes π has access to all the information available to an agent who observes π'. The relation $\mathcal{R}$ is a partial order on Π and poset $(\Pi, \mathcal{R})$ is a lattice. We let $\vee$ denote the *join*, i.e., $\pi \vee \pi'$ is the coarsest refinement of both π and π'. Note that $\pi \vee \pi'$ is the signal that is equivalent to observing both π and π'.

Given two relations on signals, $\mathcal{P}$ and $\mathcal{P}'$, we denote that $\mathcal{P}$ implies $\mathcal{P}'$ (i.e., $\pi\mathcal{P}\pi' \Rightarrow \pi\mathcal{P}'\pi'$) by $\mathcal{P} \subseteq \mathcal{P}'$.[11] If $\mathcal{P}$ implies $\mathcal{P}'$ but not vice versa, we have $\mathcal{P} \subsetneq \mathcal{P}'$.

3 Strong Blackwell

3.1 Absence of other information

A *decision problem* $D = (A, u)$ consists of a compact action set A and a continuous utility function $u : A \times \Omega \to \mathbb{R}$. Let $\mathcal{D}$ denote the class of all decision problems. The *value* of an experiment τ in problem D is given by $\mathbb{E}_{\tilde{\mu}\sim\tau}\left[\max_{a\in A} \mathbb{E}_{\omega\sim\tilde{\mu}} u(a, \omega)\right]$; the value of signal π in problem D is the value of the induced experiment $\langle\pi\rangle$.[12] Blackwell's Theorem (1953) establishes that τ is more valuable than τ' for every $D \in \mathcal{D}$ if and only if $\tau \gtrsim \tau'$.

We are primarily interested in studying comparisons of signals, rather than experiments. We say that signal π *Blackwell dominates* signal π' and write $\pi\mathcal{B}\pi'$ if π has a weakly higher value than π' for every $D \in \mathcal{D}$. Hence, $\pi\mathcal{B}\pi'$ if and only if $\langle\pi\rangle \gtrsim \langle\pi'\rangle$.

Note that $\mathcal{B}$ is a relation on signals, but it is not a partial order. While it is reflexive and transitive, it is not antisymmetric: $\pi\mathcal{B}\pi'$ and $\pi'\mathcal{B}\pi$ implies that the two signals are Blackwell equivalent ($\pi \sim \pi'$) but does not imply that they are the same signal ($\pi = \pi'$).[13] As we will see down the line, some economically meaningful relations on signals will not even be transitive.

(ω, x) to the signal realization $s \in S$ that contains (ω, x) in partition π.

[10] We then also say π' *coarsens* π.

[11] Recall that a relation on Π is a subset of $\Pi \times \Pi$, with $\pi\mathcal{P}\pi'$ denoting that $(\pi, \pi') \in \mathcal{P} \subseteq \Pi \times \Pi$.

[12] One can of course subtract the payoff under the prior from the definition of this value, but since that is constant it would not change any comparisons of experiments.

[13] The Blackwell order on *experiments*—that is, the mean-preserving spread order—is of course a partial order.

Figure 1: Blackwell vs. Strong Blackwell

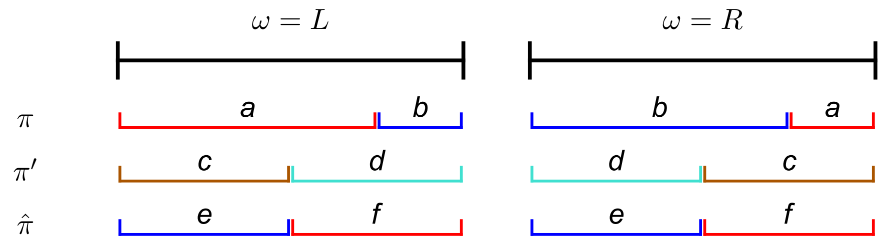

It holds that $\pi \mathcal{B} \pi'$, but it is not true that $\pi \overline{\mathcal{B}} \pi'$ because $\pi \vee \hat{\pi}$ is strictly Blackwell dominated by $\pi' \vee \hat{\pi}$.

3.2 Robustness to additional information

In the previous subsection, the analyst who compares the value of two signals is completely agnostic about the preferences of the agent but is implicitly dogmatic in her view that the signals whose value is being considered will be the only information available to the agent. We now extend agnosticism about preferences to agnosticism about what other information the agent has observed.

An *extended decision problem* $\hat{D} = (A, u, \hat{\pi})$ consists of a compact action set A, a continuous utility function $u : A \times \Omega \to \mathbb{R}$, and a signal $\hat{\pi}$. The interpretation is that we are considering an agent with action set A and utility function u who has observed signal $\hat{\pi}$. Let $\hat{\mathcal{D}}$ denote the class of all extended decision problems.

The value of a signal π in extended problem $\hat{D}$ is given by $\mathbb{E}_{\tilde{\mu} \sim \langle \pi \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega) \right]$. We say that signal π *strongly Blackwell dominates* signal π', denoted $\pi \overline{\mathcal{B}} \pi'$, if π has a higher value than π' for every $\hat{D} \in \hat{\mathcal{D}}$. We now describe some key properties of the strong Blackwell relation.

Remark 3.1. Strong Blackwell dominance implies but is not equivalent to Blackwell dominance, i.e., $\overline{\mathcal{B}} \subsetneq \mathcal{B}$. The fact that $\overline{\mathcal{B}} \subseteq \mathcal{B}$ follows from the observation that every extended decision problem is also a decision problem – we simply set $\hat{\pi}$ to be the trivial partition $\underline{\pi}$.[14] To see that $\overline{\mathcal{B}} \neq \mathcal{B}$, consider the three signals in Figure 1.

It is easy to see that $\pi \mathcal{B} \pi'$ since π is informative about the state and π' is not. But, it is not the case that $\pi \overline{\mathcal{B}} \pi'$ since $\pi \vee \hat{\pi}$ is only partially informative about the state while $\pi' \vee \hat{\pi}$ fully reveals the state.

Remark 3.2. There are two other natural ways we could ask whether the comparison of two signals

[14] The trivial partition is the one that contains a single signal realization.

is influenced by the presence of additional information. First, we could ask whether π is necessarily more valuable than π' if the agent had observed some specific signal realization. This would be an interim notion of more valuable, in contrast to the ex ante notion that is embodied in our definition. Formally, we could require that $\mathbb{E}_{\tilde{\mu}\sim\langle\pi|s\rangle}\left[\max_{a\in A}\mathbb{E}_{\omega\sim\tilde{\mu}}u\left(a,\omega\right)\right] \geq \mathbb{E}_{\tilde{\mu}\sim\langle\pi'|s\rangle}\left[\max_{a\in A}\mathbb{E}_{\omega\sim\tilde{\mu}}u\left(a,\omega\right)\right]$ for any triplet (A,u,s), where $\langle\pi|s\rangle$ denotes the distribution of posteriors induced by observing signal π after having previously observed signal realization s.

Second, we could consider the possibility that the agent, after obtaining a signal whose value we are interested in, could endogenously acquire additional costly information. Formally, we could require that $\max_{\hat{\pi}\in\Pi}\mathbb{E}_{\tilde{\mu}\sim\langle\pi\vee\hat{\pi}\rangle}\left[\max_{a\in A}\mathbb{E}_{\omega\sim\tilde{\mu}}u\left(a,\omega\right)\right]-c\left(\hat{\pi}\right) \geq \max_{\hat{\pi}\in\Pi}\mathbb{E}_{\tilde{\mu}\sim\langle\pi'\vee\hat{\pi}\rangle}\left[\max_{a\in A}\mathbb{E}_{\omega\sim\tilde{\mu}}u\left(a,\omega\right)\right]-c\left(\hat{\pi}\right)$ for any triplet (A,u,c), where $c:\Pi\to\mathbb{R}$ denotes the cost of acquiring additional information.

It turns out, however, that both of these alternative notions are equivalent to our definition of strong Blackwell dominance! We formalize and prove this claim in Appendix A.1.

Remark 3.3. Strong Blackwell dominance is transitive since Blackwell dominance is.

Our main result for this section is a characterization of strong Blackwell. This characterization can be motivated by considering two sufficient conditions for strong Blackwell.

First, it is immediate that refinement implies strong Blackwell: $\pi\mathcal{R}\pi'$ implies that $\langle\pi\vee\hat{\pi}\rangle \succsim \langle\pi'\vee\hat{\pi}\rangle$ for any $\hat{\pi}$. Second, consider any signal π that always reveals the state. It is immediate that $\pi\overline{\mathcal{B}}\pi'$ for every π': for any $\hat{\pi}$, the join $\pi\vee\hat{\pi}$ also always reveals the state, and therefore Blackwell dominates any other signal, including $\pi'\vee\hat{\pi}$.

Of course, neither of these sufficient conditions is necessary. Indeed, π' is strongly Blackwell dominated by any refinement of π', even if that refinement does not reveal the state, and π' is strongly Blackwell dominated by any signal that always reveals the state, even if that signal does not refine π'. Moreover, it might be that $\pi\overline{\mathcal{B}}\pi'$ even though π neither refines π' nor always reveals the state. The key insight is that if we consider elements of π one signal realization at a time, if it turns out that every signal realization of π either pins down the state or pins down the signal realization generated by π' (or both), then π must be more valuable than π' no matter what other information is available. Moreover, this condition is not merely sufficient for $\pi\overline{\mathcal{B}}\pi'$, it is also necessary.

Formally, say that π *reveals-or-refines* π', denoted $\pi\mathcal{O}\pi'$, if for every $s\in\pi$ either: (i) s reveals the state (i.e., $p\left(s|\omega\right)>0$ for at most one ω), or (ii) $s\subseteq s'$ for some $s'\in\pi'$. We then have the

Figure 2: Reveal-or-refine

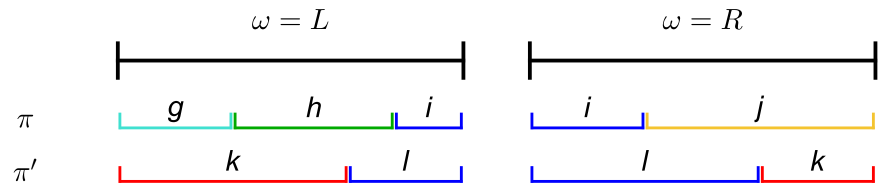

Note that π reveals-or-refines π', but it is not the case that π refines π' or that π always reveals the state.

following characterization.

Theorem 1. *Signal π strongly Blackwell dominates signal π' if and only if π reveals-or-refines π'.*

To see why reveal-or-refine implies strong Blackwell, first fix any extended decision problem. In order to show that π is more valuable than π', it suffices to show that π is more valuable than π' conditional on any signal realization s from π. If s reveals the state, nothing can be more valuable than π. If s refines $s' \in \pi'$, i.e., $s \subseteq s'$, then for any signal realization $\hat{s} \in \hat{\pi}$, $s \cap \hat{s} \subseteq s' \cap \hat{s}$, and thus π is more valuable than π'. To establish the converse, if π does not reveal-or-refine π', we explicitly construct $\hat{\pi}$ such that $\pi \vee \hat{\pi}$ does not Blackwell dominate $\pi' \vee \hat{\pi}$. Specifically, we build $\hat{\pi}$ such that $\pi \vee \hat{\pi} = \hat{\pi}$ while $\pi' \vee \hat{\pi}$ yields all the information in $\hat{\pi}$ but also sometimes reveals the state when $\hat{\pi}$ does not. Note that this argument establishes a slightly stronger result than Theorem 1: if π does not reveal-or-refine π', then there is a $\hat{\pi}$ such that $\pi' \vee \hat{\pi}$ strictly Blackwell dominates $\pi \vee \hat{\pi}$. Details on how to construct $\hat{\pi}$ are in the Appendix.

A key qualitative insight from Theorem 1 is that even though the definition of strong Blackwell involves a universal quantification over all decision problems and all signals, the universal quantifier can in fact be eliminated, and strong Blackwell is reduced to the much simpler reveal-or-refine comparison, which only requires checking a condition for each of the (finitely many) signal realizations. Indeed, using our graphical representation of signals, it is straightforward to check whether one signal reveals-or-refines another via "visual inspection." For example, to compare π and π' in Figure 2, we consider each signal realization of π in turn. Realization $g \in \pi$ both reveals the state and refines realization $k \in \pi'$ (i.e., $g \subseteq k$); realization h reveals the state; realization i refines l; finally, j reveals the state. Thus, π reveals-or-refines π'.

4 Other relations on signals

Given our shift in focus from comparisons of experiments to comparisons of signals, other natural comparisons besides the (strong or regular) Blackwell order arise.

4.1 Sufficiency

The Blackwell order is concerned with whether one source of information (e.g., the *New York Times*) is more valuable than another (e.g., the *Washington Post).* Another meaningful question is when one source of information might make another source of information moot. For example, how could we tell whether a subscription to some newspaper is worthless given an agent's existing subscriptions?

Formally, we say that π is *sufficient* for π', denoted $\pi\mathcal{S}\pi'$, if in any decision problem $D \in \mathcal{D}$, the value of signal $\pi \vee \pi'$ is the same as value of signal π.

This notion of sufficiency appears in various economic applications. For instance, Holmström (1979) shows that information about agent's effort in a moral hazard problem is valuable if and only if the observable output is not sufficient for that information.

Remark 4.1. Signal π is sufficient for signal π' if and only if $(\pi \vee \pi') \sim \pi$. If $(\pi \vee \pi') \sim \pi$, then the value of $\pi \vee \pi'$ is the same as value of π for any decision problem, so $\pi\mathcal{S}\pi'$. Conversely, if $\pi\mathcal{S}\pi'$, then the fact that π alone yields as much value as $\pi \vee \pi'$ implies that $\pi\mathcal{B}(\pi \vee \pi')$. Since we know $(\pi \vee \pi')\mathcal{B}\pi$, we have that $(\pi \vee \pi') \sim \pi$.

Yet another equivalent formulation of sufficiency is in terms of the induced random variables: $\pi\mathcal{S}\pi' \Leftrightarrow \tilde{\mu}_{\pi\vee\pi'} = \tilde{\mu}_\pi$. In general, $\tilde{\mu}_\pi = \tilde{\mu}_{\pi'} \Rightarrow \pi \sim \pi'$ but $\pi \sim \pi' \not\Rightarrow \tilde{\mu}_\pi = \tilde{\mu}_{\pi'}$. That said, we do have that $\pi\mathcal{S}\pi' \Leftrightarrow \pi \vee \pi' \sim \pi \Leftrightarrow \tilde{\mu}_{\pi\vee\pi'} = \tilde{\mu}_\pi$. This equivalence follows from the more general result that if π^* refines π and $\pi^* \sim \pi$, then $\tilde{\mu}_{\pi^*} = \tilde{\mu}_\pi$.

The formulation of sufficiency in terms of random variables provides a simple way to check whether one signal is sufficient for another. To compare π and π' in Figure 3, we consider each signal realization of $\pi \vee \pi'$ in turn. Realization $g = k \cap g \in \pi \vee \pi'$ clearly leads to the same belief as $g \in \pi$; realization $k \cap m \in \pi \vee \pi'$ leads to the same belief as $m \in \pi$ since $\frac{Pr(k\cap m|\omega=L)}{Pr(k\cap m|\omega=R)} = \frac{Pr(m|\omega=L)}{Pr(m|\omega=R)}$. The same applies to $l \cap m$ and m, and hence $\pi\mathcal{S}\pi'$.

Figure 3: Checking for sufficiency

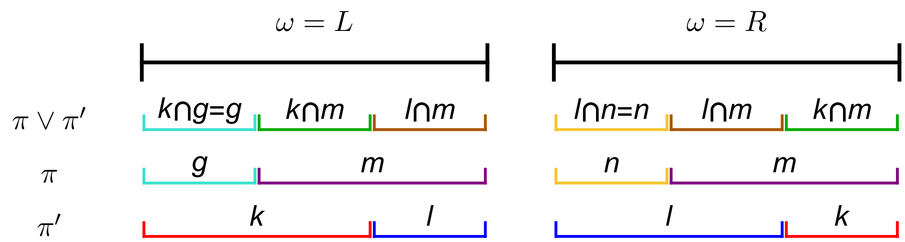

One can confirm that $\pi\mathcal{S}\pi'$ by comparing the likelihood ratios of each signal realization in π to the likelihood ratios of the overlapping signal realizations in $\pi \vee \pi'$.

Remark 4.2. Another equivalent definition of sufficiency is: for all $s \in \pi$ and all $s' \in \pi'$, $Pr\left(s'|s,\omega\right)$ is independent of ω. This formulation echoes Blackwell's (1953) notion of a garbling.[15] But unlike Blackwell, we have specified the underlying probability space, so we are not asking whether there exists a garbling that transforms experiment $\langle\pi\rangle$ into $\langle\pi'\rangle$. Rather, we ask whether – given their underlying correlation – the signal π' adds information about the state given signal π. Relatedly, it is worth noting that the following three conditions are equivalent: (i) $\pi\mathcal{B}\pi'$, (ii) $\exists\pi^*$ s.t. $\pi \sim \pi^*$ and $\pi^*\mathcal{R}\pi'$, and (iii) $\exists\pi^*$ s.t. $\pi \sim \pi^*$ and $\pi^*\mathcal{S}\pi'$. The equivalence of (i) and (ii) is Theorem 1 in Green and Stokey (1978).[16] The equivalence of (i) and (iii) is closely related to a standard formulation of Blackwell's theorem.

Remark 4.3. $\overline{\mathcal{B}} \subsetneq \mathcal{S} \subsetneq \mathcal{B}$.

First, it is easy to see that $\overline{\mathcal{B}} \subseteq \mathcal{S}$. If $\pi\overline{\mathcal{B}}\pi'$ we know $(\pi \vee \hat{\pi})\,\mathcal{B}\,(\pi' \vee \hat{\pi})$ for any $\hat{\pi}$, including $\hat{\pi} = \pi$; hence, $\pi\mathcal{B}\,(\pi \vee \pi')$. Since it's always the case that $(\pi \vee \pi')\,\mathcal{B}\pi$, we have $\pi \sim (\pi \vee \pi')$. Hence, $\pi\overline{\mathcal{B}}\pi'$ implies $\pi\mathcal{S}\pi'$.

Importantly, however, $\overline{\mathcal{B}} \neq \mathcal{S}$. This can be seen in Figure 5 where $\pi_{\mathcal{S}}\mathcal{S}\pi_0$ but $\neg\left(\pi_{\mathcal{S}}\overline{\mathcal{B}}\pi_0\right)$. The fact that $\overline{\mathcal{B}} \neq \mathcal{S}$ has a substantive economic interpretation. Suppose we know that $\pi\mathcal{B}\pi'$ but are not sure whether π remains more valuable than π' in presence of some additional information $\hat{\pi}$. One might think that the "worst case" scenario would be if $\hat{\pi} = \pi$ (e.g., in comparing the value of *NYT* to the value of *WP*, we worry that the reader already has a subscription to the *NYT*). This scenario, however, only tells us, however, whether $\pi\mathcal{S}\pi'$, which is a weaker condition than $\pi\overline{\mathcal{B}}\pi'$.

[15] This formulation also clarifies the relationship between our definition of sufficiency and the notion of a sufficient statistic in the field of statistics. Given some data $\vec{x}$, recall that a function $t\left(\vec{x}\right)$ is a sufficient statistic for ω if $Pr\left(\vec{x}|t\left(\vec{x}\right),\omega\right)$ is independent of ω.

[16] It is stated and proved using the formalism in our paper by Gentzkow and Kamenica (2017a).

Figure 4: Sufficiency is not transitive

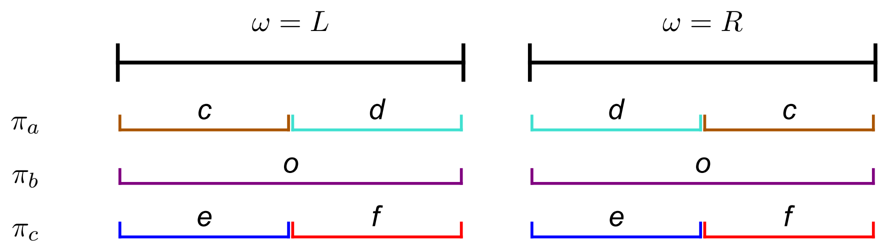

$\pi_a \mathcal{S} \pi_b$ and $\pi_b \mathcal{S} \pi_c$ but $\neg (\pi_a \mathcal{S} \pi_c)$.

Thus, $\hat{\pi} = \pi$ is not the most stringent test-case for strong Blackwell dominance. Instead, a greater concern is the possibility that $\hat{\pi}$ is complementary to π'.

It is also easy to see that $\mathcal{S} \subseteq \mathcal{B}$ since $\pi \mathcal{S} \pi'$ means the value of π in any decision problem is the same as the value of $\pi \vee \pi'$, which in turn must be weakly higher than the value of π'. Moreover, Figure 5 establishes that $\mathcal{S} \neq \mathcal{B}$ since $\pi_{\mathcal{B}} \mathcal{B} \pi_0$ but $\neg (\pi_{\mathcal{B}} \mathcal{S} \pi_0)$.

Remark 4.4. Sufficiency is not transitive. Consider Figure 4. Since $\pi_a \vee \pi_b = \pi_a$, π_a is *a fortiori* sufficient for π_b. Since both π_b and $\pi_b \vee \pi_c$ provide no information about the state, we have that π_b is sufficient for π_c. Yet, π_a is *not* sufficient for π_c; π_a on its own provides no information about the state while $\pi_a \vee \pi_c$ fully reveals the state.

4.2 Martingale

A widely-used and basic observation in information economics is that "beliefs are a martingale." If an agent with some current belief μ_0 observes additional data from some source of information, her expected posterior belief must be μ_0. This is a consequence of the Law of Iterated Expectations.

In the context of signals, one way to formulate this observation is to note that if π refines π', then it must be the case that $\mathbb{E}\left[\tilde{\mu}_\pi | \tilde{s}_{\pi'}\right] = \tilde{\mu}_{\pi'}$. In other words, additional information cannot change the beliefs on average.

In this paper, we take a novel perspective on the martingale property. Instead of treating it as an implication of Bayesian updating, we consider it as a relation between sources of information. When is it the case that, if I read the *Washington Post*, I think that in expectation, a reader of the *New York Times* would hold the same belief that I do? If the *Washington Post* were much more informative than the *NYT*, there would be no reason to think this: a reader of the *WP* might know the state of the world and yet expect the reader of the NYT to remain uninformed. By contrast, if the NYT contains all the information that the WP does, then the WP reader would in fact think that the expected belief of the NYT-reader is equal to her own. Thus, the martingale property ($\mathbb{E}\left[\tilde{\mu}_{\pi}|\tilde{s}_{\pi'}\right] = \tilde{\mu}_{\pi'}$) tells us that π is in some sense "more informative" than π'. In this section, we unpack what that means.

To formally define the martingale relation, we need to address a subtlety that was absent from the considerations of the Blackwell and sufficiency relations. We fixed an interior prior μ_0 at the outset, but as is well known, the Blackwell comparison (and the sufficiency comparison by extension) are prior independent. Thus, whether $\pi\mathcal{B}\pi'$ or whether $\pi\mathcal{S}\pi'$ does not depend on μ_0.

By contrast, for a given π and π', whether $\mathbb{E}\left[\tilde{\mu}_{\pi}|\tilde{s}_{\pi'}\right] = \tilde{\mu}_{\pi'}$ can depend on μ_0. (An example of this is given in Appendix A.4.)

Accordingly, we say π *martingale dominates* π', denoted $\pi\mathcal{M}\pi'$, if $\mathbb{E}\left[\tilde{\mu}_{\pi}|\tilde{s}_{\pi'}\right] = \tilde{\mu}_{\pi'}$ holds for any choice of μ_0.

Remark 4.5. In prior work (Brooks *et al.*, 2022), we introduced a similar relation, termed *belief-martingale,* defined by $\mathbb{E}\left[\tilde{\mu}_{\pi}|\tilde{\mu}_{\pi'}\right] = \tilde{\mu}_{\pi'}$. To understand the distinction between the two relations, it is helpful to introduce the idea of the belief-coarsening of a signal. Given any signal π, we let the *belief-coarsening* of π, denoted $C(\pi)$ be the signal that "pools together" any signal realizations in π that induce the same belief. Formally, $C(\pi)$ is the finest coarsening of π such that for any $s, s' \in C(\pi),\ s \neq s' \Rightarrow \mu_s \neq \mu_{s'}$. With this definition in hand, we have that π belief-martingale dominates π' if and only if π martingale dominates $C(\pi')$.[17] Moreover, $\pi\mathcal{M}\pi'$ implies that π belief-martingale dominates π'.

Remark 4.6. The notion of belief-coarsening also provides another way to characterize the martingale relation. In particular, it turns out that $\pi\mathcal{M}\pi'$ if and only if $C(\pi)\,\mathcal{S}\pi'$. It is easy to see that $C(\pi)\,\mathcal{S}\pi'$ implies $\pi\mathcal{M}\pi'$ because $C(\pi)\,\mathcal{S}\pi'$ implies $C(\pi)\,\mathcal{M}\pi'$ (because $\mathcal{S} \subseteq \mathcal{M}$) and $C(\pi)\,\mathcal{M}\pi'$ implies $\pi\mathcal{M}\pi'$ (because $\tilde{\mu}_{C(\pi)} = \tilde{\mu}_{\pi}$). The other direction is more subtle, and we provide a detailed argument in the Appendix. The equivalence of $\pi\mathcal{M}\pi'$ with $C(\pi)\mathcal{S}\pi'$ is illustrated in Figure 5 in

[17] For another illustration of belief-coarsening, if π refines $C(\pi')$ that means than an agent who observes π knows the first-order beliefs of an agent who observes π'. In Brooks *et al.* (2022), we discuss such "knowledge of first-order beliefs" as an example a proper relation on signals, i.e., a relation that is implied by refinement and implies belief-martingale.

Section 5, taking $\pi = \pi_{\mathcal{M}}$ and $\pi' = \pi_0$: $\pi_{\mathcal{M}}$ martingale dominates π_0 and $\pi_{\mathcal{M}}$ is not sufficient for π_0, but $C(\pi_{\mathcal{M}})$ is sufficient for π_0. More generally, the fact that $\pi \mathcal{M} \pi' \Leftrightarrow C(\pi) \mathcal{S} \pi'$ provides a simple way to check whether one signal martingale dominates another.

Remark 4.7. $\mathcal{S} \subsetneq \mathcal{M} \subsetneq \mathcal{B}$.

To see that $\mathcal{S} \subseteq \mathcal{M}$, suppose $\pi \mathcal{S} \pi'$. Since $\mathcal{R} \subseteq \mathcal{M}$, we know that $(\pi \vee \pi') \mathcal{M} \pi'$, i.e., $\mathbb{E}[\tilde{\mu}_{\pi \vee \pi'} | \tilde{s}_{\pi'}] = \tilde{\mu}_{\pi'}$, which in turn implies $\mathbb{E}[\tilde{\mu}_{\pi} | \tilde{s}_{\pi'}] = \tilde{\mu}_{\pi'}$ since $\tilde{\mu}_{\pi \vee \pi'} = \tilde{\mu}_{\pi}$. Thus $\pi \mathcal{M} \pi'$. To see that $\mathcal{M} \subseteq \mathcal{B}$, note that $\mathbb{E}[\tilde{\mu}_{\pi} | \tilde{s}_{\pi'}] = \tilde{\mu}_{\pi'}$ implies that the distribution of $\tilde{\mu}_{\pi}$ is a mean-preserving spread of the distribution of $\tilde{\mu}_{\pi'}$. Another way to see that $\mathcal{S} \subseteq \mathcal{M} \subseteq \mathcal{B}$ is to note the following analogous characterizations of these three relations (as shown in the Appendix).

- $\pi \mathcal{S} \pi'$ if and only if $\exists \pi^*$ s.t. $C(\pi^*) = C(\pi)$, $\pi^* \mathcal{R} \pi$, and $\pi^* \mathcal{R} \pi'$. (We can take $\pi^* = \pi \vee \pi'$.)
- $\pi \mathcal{M} \pi'$ if and only if $\exists \pi^*$ s.t. $C(\pi^*) = C(\pi)$ and $\pi^* \mathcal{R} \pi'$.
- $\pi \mathcal{B} \pi'$ if and only if $\exists \pi^*$ s.t. $\pi^* \sim \pi$ and $\pi^* \mathcal{R} \pi'$ (as noted in Remark 4.2).[18]

To see that $\mathcal{S} \neq \mathcal{M} \neq \mathcal{B}$, see Figure 5 in Section 5. In the figure, we see that $\pi_{\mathcal{M}} \mathcal{M} \pi_0$ but $\neg(\pi_{\mathcal{M}} \mathcal{S} \pi_0)$, and that $\pi_{\mathcal{B}} \mathcal{B} \pi_0$ but $\neg(\pi_{\mathcal{B}} \mathcal{M} \pi_0)$.

Remark 4.8. The martingale relation is not transitive. See example in Figure 7 in Appendix A.6.

5 Strengthening relations

So far, we discussed how to strengthen the Blackwell comparison to allow for the presence of additional information, and we introduced two new relations on signals, sufficiency and martingale. A natural question then, of course, is how to strengthen sufficiency or martingale, to make those comparisons robust to the presence of additional information. To do so, we introduce a general notion of strengthening: given any relation $\mathcal{P}$ on the set of signals Π, we define *strong* $\mathcal{P}$, denoted $\overline{\mathcal{P}}$, by $\pi \overline{\mathcal{P}} \pi'$ if $(\pi \vee \hat{\pi}) \mathcal{P} (\pi' \vee \hat{\pi})$ for all $\hat{\pi} \in \Pi$.

Of course, we have already seen one important case of the strong version of a relation, the strong Blackwell order $\overline{\mathcal{B}}$. It is easy to see that refinement is unaffected by strengthening, i.e., $\overline{\mathcal{R}} = \mathcal{R}$:

[18] Since $C(\pi^*) = C(\pi)$ implies that $\pi^* \sim \pi$, these characterizations make it clear that $\mathcal{M} \subseteq \mathcal{B}$.

if π refines π' then $\pi \vee \hat{\pi}$ refines $\pi' \vee \hat{\pi}$ for any $\hat{\pi}$. We now describe three important properties of strengthening.

Remark 5.1. Strengthening strengthens. For any $\mathcal{P}$, we have $\overline{\mathcal{P}} \subseteq \mathcal{P}$. If $(\pi \vee \hat{\pi})\,\mathcal{P}\,(\pi' \vee \hat{\pi})$ for all $\hat{\pi}$, then $\pi = (\pi \vee \underline{\pi})\,\mathcal{P}\,(\pi' \vee \underline{\pi}) = \pi'$.

Remark 5.2. Strengthening is idempotent. For any $\mathcal{P}$, $\overline{\overline{\mathcal{P}}} = \overline{\mathcal{P}}$. From the previous remark, $\overline{\overline{\mathcal{P}}} \subseteq \overline{\mathcal{P}}$. To show $\overline{\mathcal{P}} \subseteq \overline{\overline{\mathcal{P}}}$, suppose $\pi \overline{\mathcal{P}} \pi'$, i.e., $(\pi \vee \hat{\pi})\,\mathcal{P}\,(\pi' \vee \hat{\pi})$ for all $\hat{\pi} \in \Pi$. Then, for any $\hat{\pi}, \tilde{\pi} \in \Pi$, we have $(\pi \vee \hat{\pi} \vee \pi^*)\,\mathcal{P}\,(\pi' \vee \hat{\pi} \vee \tilde{\pi})$ since $\hat{\pi} \vee \tilde{\pi} \in \Pi$.

Remark 5.3. Strengthening is monotone. If $\mathcal{P} \subseteq \mathcal{P}'$, then $\overline{\mathcal{P}} \subseteq \overline{\mathcal{P}'}$. Suppose $\mathcal{P} \subseteq \mathcal{P}'$ and $\pi \overline{\mathcal{P}} \pi'$. For any $\hat{\pi}$, we have that $(\pi \vee \hat{\pi})\,\mathcal{P}\,(\pi' \vee \hat{\pi})$, which in turn implies $(\pi \vee \hat{\pi})\,\mathcal{P}'\,(\pi' \vee \hat{\pi})$. Since this holds for all $\hat{\pi}$, we have that $\pi \overline{\mathcal{P}'} \pi'$.

With these properties in hand, it turns out to be very easy to characterize the strong versions of both sufficiency and martingale relations.

Theorem 2. *Suppose $\mathcal{P}$ is a relation on Π and $\overline{\mathcal{B}} \subseteq \mathcal{P} \subseteq \mathcal{B}$. Then, $\pi \overline{\mathcal{P}} \pi'$ if and only if π reveals-or-refines π'. Hence, $\overline{\mathcal{B}} = \overline{\mathcal{S}} = \overline{\mathcal{M}}$.*

Proof. Suppose $\mathcal{P}$ is a relation on Π and $\overline{\mathcal{B}} \subseteq \mathcal{P} \subseteq \mathcal{B}$. That $\overline{\mathcal{P}} \subseteq \overline{\mathcal{B}}$ follows from $\mathcal{P} \subseteq \mathcal{B}$ by monotonicity of strengthening. To show $\overline{\mathcal{B}} \subseteq \overline{\mathcal{P}}$, we first observe that $\overline{\mathcal{B}} \subseteq \mathcal{P}$ implies $\overline{\overline{\mathcal{B}}} \subseteq \overline{\mathcal{P}}$ (by monotonicity) which in turn implies $\overline{\mathcal{B}} \subseteq \overline{\mathcal{P}}$ (by idempotence). Since $\overline{\mathcal{B}} \subseteq \mathcal{S} \subseteq \mathcal{M} \subseteq \mathcal{B}$, it follows that $\overline{\mathcal{B}} = \overline{\mathcal{S}} = \overline{\mathcal{M}}$. $\square$

There are various ways to compare the usefulness of a source of information – sufficiency, martingale, Blackwell. For any of these comparisons, we may wish to consider the strong version of the comparison that is robust to the potential presence of additional information. Theorem 2 delivers a remarkable message, namely that, even though sufficiency, martingale, and Blackwell are all distinct, their strong versions coincide! Moreover, the strong version of each of these comparisons is a simple relation, reveal-or-refine, that is very easy to check and involves no quantifiers over decision problems or signals.

Summarizing the relations that we have considered, we can order them as follows:

$$\overline{\mathcal{R}} = \mathcal{R} \subsetneq \overline{\mathcal{S}} = \overline{\mathcal{M}} = \overline{\mathcal{B}} \subsetneq \mathcal{S} \subsetneq \mathcal{M} \subsetneq \mathcal{B}.$$

Figure 5: Ranking the relations

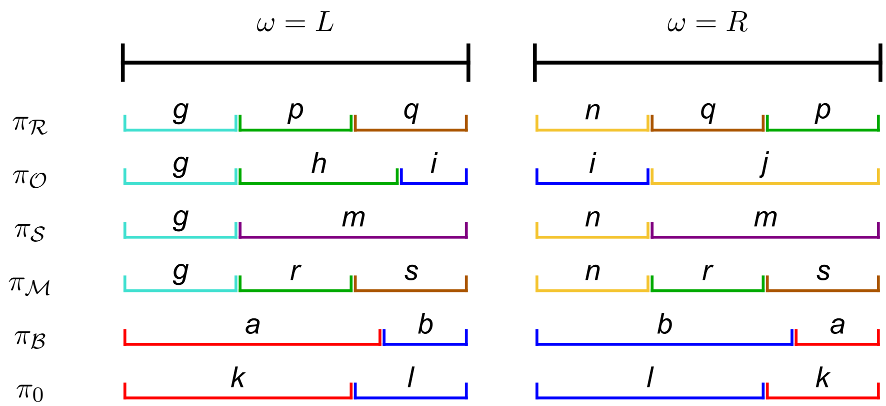

This figure illustrates the ranking of the relations $\mathcal{R} \subsetneq \mathcal{O} \subsetneq \mathcal{S} \subsetneq \mathcal{M} \subsetneq \mathcal{B}$, where for each relation $\mathcal{P}$, we have $\pi_{\mathcal{P}}\mathcal{P}\pi_0$. To confirm the strictness of this ranking, we see that $\pi_{\mathcal{O}}\mathcal{O}\pi_0$ but $\neg(\pi_{\mathcal{O}}\mathcal{R}\pi_0)$; $\pi_{\mathcal{S}}\mathcal{S}\pi_0$ but $\neg(\pi_{\mathcal{S}}\mathcal{O}\pi_0)$; $\pi_{\mathcal{M}}\mathcal{M}\pi_0$ but $\neg(\pi_{\mathcal{M}}\mathcal{S}\pi_0)$; and $\pi_{\mathcal{B}}\mathcal{B}\pi_0$ but $\neg(\pi_{\mathcal{B}}\mathcal{M}\pi_0)$. The fact that $\neg(\pi_{\mathcal{B}}\mathcal{M}\pi_0)$ follows from the fact that $C(\pi_{\mathcal{B}}) = \pi_{\mathcal{B}}$ is not sufficient for π_0.

Figure 5 illustrates the strict comparisons, providing examples where signals are ranked by reveal-or-refine but not refinement; sufficiency but not reveal-or-refine; martingale but not sufficiency; and Blackwell but not martingale.

6 Conclusion

Experiments have long been considered the natural formalism for modeling information sources. As we and others have argued, this formalism is incomplete, in that the definition of different experiments does not specify how they interact with one another. In contrast, we model information sources as *signals*, which provide a complete description of the joint distribution of data from one information source and all others.

With this shift in focus from experiments to signals, a number of natural questions emerge. In Brooks *et al.* (2022), we investigate the conditions under which a partial order on the information content of experiments can be made consistent with an analogous ordering on signals. In the present paper, we compare the value of signals with a focus on the robustness to potential presence of other information. We also argue for the study of relations on signals beyond the familiar Blackwell and refinement orders, including sufficiency and martingale. But many questions remain. What other relations on signals may be useful and/or meaningful in economic applications? What are the decision- or game-theoretic foundations for the different relations? (Sufficiency, for example, has a simple characterization that one signal not add value to another in any decision problem; refinement may be relevant in games, when a player cares not only about the underlying state but also about what other players know; we do not know of natural foundations for the martingale relation.) What are economically reasonable ways to model the cost of acquiring signals? We leave these issues for future work.

References

Athey, S. and Levin, J. (2018). The value of information in monotone decision problems. *Research in Economics*, **72** (1), 101–116.

Blackwell, D. (1951). Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, The Regents of the University of California.

— (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, pp. 265–272.

Börgers, T., Hernando-Veciana, A. and Krähmer, D. (2013). When are signals complements or substitutes? *Journal of Economic Theory*, **148** (1), 165–195.

Brooks, B., Frankel, A. and Kamenica, E. (2022). Information hierarchies. *Econometrica*, **90** (5), 2187–2214.

Ely, J., Frankel, A. and Kamenica, E. (2015). Suspense and surprise. *Journal of Political Economy*, **123** (1), 215–260.

Frankel, A. and Kamenica, E. (2019). Quantifying information and uncertainty. *American Economic Review*, **109** (10), 3650–80.

— and Kasy, M. (2022). Which findings should be published? *American Economic Journal: Microeconomics*, **14** (1), 1–38.

Gentzkow, M. and Kamenica, E. (2017a). Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, **104**, 411–429.

— and — (2017b). Competition in persuasion. *The Review of Economic Studies*, **84** (1), 300–322.

Green, J. R. and Stokey, N. L. (1978). Two representations of information structures and their comparisons.

Holmström, B. (1979). Moral hazard and observability. *The Bell journal of economics*, pp. 74–91.

Lehmann, E. (1988). Comparing location experiments. *The Annals of Statistics*, pp. 521–533.

Liang, A. and Mu, X. (2020). Complementary information and learning traps. *The Quarterly Journal of Economics*, **135** (1), 389–448.

—, — and Syrgkanis, V. (2022). Dynamically aggregating diverse information. *Econometrica*, **90** (1), 47–80.

Moscarini, G. and Smith, L. (2002). The law of large demand for information. *Econometrica*, **70** (6), 2351–2366.

Mu, X., Pomatto, L., Strack, P. and Tamuz, O. (2021a). Background risk and small-stakes risk aversion. *arXiv preprint arXiv:2010.08033*.

—, —, — and — (2021b). From blackwell dominance in large samples to rényi divergences and back again. *Econometrica*, **89** (1), 475–506.

Persico, N. (2000). Information acquisition in auctions. *Econometrica*, **68** (1), 135–148.

A Appendix

A.1 Alternative formulations of strong Blackwell dominance

An interim decision problem $\hat{D}^i = (A, u, s)$ consists of a compact action set A, a continuous utility function $u : A \times \Omega \to \mathbb{R}$, and a signal realization $s \in S$. The interpretation is that we are considering an agent with action set A and utility function u who has observed signal realization s. Let $\hat{\mathcal{D}}^i$ denote the class of interim decision problems. Let $\langle \pi | s \rangle$ denote the distribution of posteriors induced by observing signal π after having previously observed signal realization s: letting $Pr\left(\hat{s}\right) \equiv \sum_{\omega \in \Omega} p^{\omega}(\hat{s}) \mu_0^{\omega}$ denote the unconditional probability of realization $\hat{s}$ for any $\hat{s} \in S$, distribution $\langle \pi | s \rangle$ assigns probability $\sum_{\{s' \in \pi : \mu_{s \cap s'} = \mu\}} \frac{Pr(s \cap s')}{Pr(s)}$ to each belief μ. The value of a signal π in an interim decision problem $\hat{D}^i$ is given by $\mathbb{E}_{\tilde{\mu} \sim \langle \pi | s \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u\left(a, \omega\right)\right]$. We say π *strongly Blackwell dominates* π' *in the interim sense* and write $\pi \overline{\mathcal{B}}^i \pi'$ if π has a higher value than π' for every $\hat{D}^i \in \hat{\mathcal{D}}^i$.

A costly acquisition decision problem $\hat{D}^k = (A, u, c)$ consists of a compact action set A, a continuous utility function $u : A \times \Omega \to \mathbb{R}$, and a cost function $c : \Pi \to \mathbb{R}$. The interpretation is that we are considering an agent with action set A and utility function u who can, in addition to the signal whose value we are considering, acquire any additional signal $\hat{\pi}$ at cost $c\left(\hat{\pi}\right)$.[19] To simplify notation, we impose that the trivial partition is free, i.e., $c\left(\underline{\pi}\right) = 0$. Let $\hat{\mathcal{D}}^k$ denote the class of costly acquisition decision problems. The value of signal π in a costly acquisition decision problem is[20]

$$\max_{\hat{\pi} \in \Pi} \mathbb{E}_{\tilde{\mu} \sim \langle \pi \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u\left(a, \omega\right)\right] - c\left(\hat{\pi}\right).$$

We say π *strongly Blackwell dominates* π' *under costly information acquisition* and write $\pi \overline{\mathcal{B}}^k \pi'$ if π has a higher value than π' for every $\hat{D}^k \in \hat{\mathcal{D}}^k$.

These two alternative notions are equivalent to strong Blackwell:

Proposition 1. $\overline{\mathcal{B}} = \overline{\mathcal{B}}^i = \overline{\mathcal{B}}^k$.

Proof. Suppose $\pi \overline{\mathcal{B}} \pi'$. Consider some interim decision problem (A, u, s). Let $\hat{\pi}$ be a signal that

[19] We could also consider the possibility that the agent chooses what additional costly information to acquire only after she observes the realization of the signal whose value we are considering. Once again, this notion would be equivalent to the strong Blackwell order.

[20] We assume the maximum exists. Alternatively, we could replace *max* with *sup*, but then establishing the equivalence below would require a slightly more involved argument.

consists of s and, for each state, a signal realization (disjoint with s) that reveals that state, i.e., $\hat{\pi} = \{s\} \cup \{(\omega, x) \,|\, (\omega, x) \notin s\}_{\omega \in \Omega}$. Since $\pi \overline{\mathcal{B}} \pi'$, we know $\mathbb{E}_{\tilde{\mu} \sim \langle \pi \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] \geq \mathbb{E}_{\tilde{\mu} \sim \langle \pi' \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right]$, i.e.,

$$\sum_{\hat{s} \in \hat{\pi}} Pr(\hat{s}) \, \mathbb{E}_{\tilde{\mu} \sim \langle \pi | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] \geq \sum_{\hat{s} \in \hat{\pi}} Pr(\hat{s}) \, \mathbb{E}_{\tilde{\mu} \sim \langle \pi' | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right].$$

For any $\hat{s} \in \hat{\pi}$ with $\hat{s} \neq s$, we have $\mathbb{E}_{\tilde{\mu} \sim \langle \pi | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] = \mathbb{E}_{\tilde{\mu} \sim \langle \pi' | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right]$ since $\hat{s}$ fully reveals the state. Thus, we must have

$$\mathbb{E}_{\tilde{\mu} \sim \langle \pi | s \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] \geq \mathbb{E}_{\tilde{\mu} \sim \langle \pi' | s \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right].$$

Since the choice of (A, u, s) was arbitrary, we conclude $\pi \overline{\mathcal{B}}^i \pi'$. Thus, $\overline{\mathcal{B}} \subseteq \overline{\mathcal{B}}^i$.

Suppose $\pi \overline{\mathcal{B}}^i \pi'$. Consider some extended decision problem $(A, u, \hat{\pi})$. We have that

$$\begin{aligned} \mathbb{E}_{\tilde{\mu} \sim \langle \pi \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] - \mathbb{E}_{\tilde{\mu} \sim \langle \pi' \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] &= \\ \sum_{\hat{s} \in \hat{\pi}} Pr(\hat{s}) \, \mathbb{E}_{\tilde{\mu} \sim \langle \pi | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] - \sum_{\hat{s} \in \hat{\pi}} Pr(\hat{s}) \, \mathbb{E}_{\tilde{\mu} \sim \langle \pi' | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] &= \\ \sum_{\hat{s} \in \hat{\pi}} Pr(\hat{s}) \left(\mathbb{E}_{\tilde{\mu} \sim \langle \pi | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] - \mathbb{E}_{\tilde{\mu} \sim \langle \pi' | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right]\right). \end{aligned}$$

Since $\pi \overline{\mathcal{B}}^i \pi'$, we know $\mathbb{E}_{\tilde{\mu} \sim \langle \pi | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] - \mathbb{E}_{\tilde{\mu} \sim \langle \pi' | \hat{s} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] \geq 0$ for each $\hat{s}$ and thus $\mathbb{E}_{\tilde{\mu} \sim \langle \pi \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] \geq \mathbb{E}_{\tilde{\mu} \sim \langle \pi' \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right]$. Since the choice of $(A, u, \hat{\pi})$ was arbitrary, we conclude $\pi \overline{\mathcal{B}} \pi'$. Thus, $\overline{\mathcal{B}}^i \subseteq \overline{\mathcal{B}}$.

Suppose $\pi \overline{\mathcal{B}} \pi'$. Suppose the value of π' on some costly acquisition decision problem (A, u, c) is v. Let π^* be a signal that the agent acquires in addition to π' in the problem (A, u, c), i.e., $\pi^* \in \arg\max_{\hat{\pi} \in \Pi} \mathbb{E}_{\tilde{\mu} \sim \langle \pi' \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] - c(\hat{\pi})$. It must be the case that the value of π' in the extended decision problem (A, u, π^*) is at least $v + c(\pi^*)$. Therefore, since $\pi \overline{\mathcal{B}} \pi'$, the value of π in (A, u, π^*) is at least $v + c(\pi^*)$. Finally, this implies that the value of π in the costly acquisition decision problem (A, u, c) is at least v since $\max_{\hat{\pi} \in \Pi} \mathbb{E}_{\tilde{\mu} \sim \langle \pi \vee \hat{\pi} \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] - c(\hat{\pi}) \geq \mathbb{E}_{\tilde{\mu} \sim \langle \pi \vee \pi^* \rangle} \left[\max_{a \in A} \mathbb{E}_{\omega \sim \tilde{\mu}} u(a, \omega)\right] - c(\pi^*)$. Thus, $\overline{\mathcal{B}} \subseteq \overline{\mathcal{B}}^k$.

Suppose $\pi\overline{\mathcal{B}}^k\pi'$. Consider some extended decision problem $(A, u, \hat{\pi})$. Let

$$K = \max_{\omega}\left(\max_{a\in A} u(a,\omega) - \min_{a\in A} u(a,\omega)\right).$$

Clearly, the value of any signal (relative to $\underline{\pi}$) in the decision problem (A, u) is less than K. Let $c^* : \Pi \to \mathbb{R}$ be as follows: $c^*(\hat{\pi}) = c^*(\underline{\pi}) = 0$ and $c^*(\pi) = K$ for all $\pi \notin \{\hat{\pi},\ \underline{\pi}\ \}$. Then, the value of any signal in the extended decision problem $(A, u, \hat{\pi})$ must be the same as the value of that signal in the costly acquisition decision problem (A, u, c^*). Since $\pi\overline{\mathcal{B}}^k\pi'$, we know that π is more valuable than π' in (A, u, c^*); thus π is more valuable than π' in $(A, u, \hat{\pi})$. Since the choice of $(A, u, \hat{\pi})$ was arbitrary, we conclude $\pi\overline{\mathcal{B}}\pi'$. Thus, $\overline{\mathcal{B}}^k \subseteq \overline{\mathcal{B}}$. □

A.2 Proof of Theorem 1

We already provided a proof in the text that $\pi\mathcal{O}\pi$ implies $\pi\overline{\mathcal{B}}\pi'$. To establish the other direction, suppose $\neg(\pi\mathcal{O}\pi')$. We seek to show that $\neg\left(\pi\overline{\mathcal{B}}\pi'\right)$, i.e., there exists $\hat{\pi}$ such that $\neg(\pi \vee \hat{\pi}\mathcal{B}\pi' \vee \hat{\pi})$. Let $s \in \pi$ be a signal realization that does not reveal the state and there is no $s' \in \pi'$ such that $s \subseteq s'$. There must be distinct states ω_1 and ω_2 and distinct signal realizations $s'_1, s'_2 \in \pi'$ such that $p(s \cap s'_1|\omega_1) > 0$ and $p(s \cap s'_2|\omega_2) > 0$. Let E be the event $(s \cap s'_1 \cap (\{\omega_1\} \times [0,1])) \cup (s \cap s'_2 \cap (\{\omega_2\} \times [0,1]))$. Let E^c denote the complement of E in $\Omega \times [0,1]$. Let $\pi'' = \{E, E^c\}$. Let $\hat{\pi} = \pi \vee \pi''$.

Note that $\pi \vee \hat{\pi} = \pi \vee \pi''$ so $\pi \vee \hat{\pi}$ reveals: (i) everything that π reveals, and (ii) whether (ω, x) is in E or not. By contrast, $\pi' \vee \hat{\pi} = \pi' \vee \pi \vee \pi''$, so $\pi' \vee \hat{\pi}$ reveals: (i) everything that π reveals, (ii) whether (ω, x) is in E or not, and (iii) if $(\omega, x) \in E$, whether ω is ω_1 or ω_2. Hence, $\pi' \vee \hat{\pi}$ Blackwell dominates $\pi \vee \hat{\pi}$.

A.3 Characterization of martingale relation

In Remark 4.6, we referred to the following observation, which we now state formally:

Proposition 2. *$\pi\mathcal{M}\pi'$ if and only if $C(\pi)\mathcal{S}\pi'$.*

As mentioned earlier, it is straightforward to observe that $C(\pi)\mathcal{S}\pi'$ implies that $\pi\mathcal{M}\pi'$. Here, we provide the proof of the other direction, that $\pi\mathcal{M}\pi'$ implies $C(\pi)\mathcal{S}\pi'$.

We begin with some notation. Let $\mu_u \in \Delta(\Omega)$ indicate the uniform prior over states. Let $\Delta^o(\Omega)$ indicate the interior of the set of beliefs. Let $N = |\Omega|$. For a signal realization s, let $p(s)$ be the vector of probabilities of the signal realization s, i.e., $p(s) = (p^\omega(s))_{\omega\in\Omega}$. Denote its relative likelihood vector as $l(s) \equiv p(s)/\sum_\omega p^\omega(s)$. The relative likelihood vector is exactly the induced posterior from observing s under the uniform prior μ_u, and for future reference, we observe that $l(s)\cdot\mu_u = 1/N$. Notice that, fixing an interior prior $\mu_0 \in \Delta^o(\Omega)$, the posterior belief after observing s is generated by a one-to-one mapping from $l(s)$ into $\Delta(\Omega)$. Hence, $C(\pi)$ pools together all of the realizations $s \in \pi$ that have identical relative likelihood vectors $l(s)$.

Now suppose that $C(\pi)$ is not sufficient for π'. We seek to show that π does not martingale dominate π'.

Because $\neg(C(\pi)\mathcal{S}\pi')$, there exist $\underline{s} \in \pi'$ and $\overline{s} \in C(\pi)$ that have a non-trivial intersection (i.e., $p^\omega(\underline{s}\cap\overline{s}) > 0$ for some ω) and $l(\overline{s}) \neq l(\underline{s}\cap\overline{s})$, since $l(\overline{s}) \neq l(\underline{s}\cap\overline{s})$ implies that posterior beliefs are different after observing $\overline{s}$ versus $\overline{s}$ and $\underline{s}$. Fix this element $\underline{s} \in \pi'$. Denumerate the elements of $C(\pi)$ that non-trivially intersect $\underline{s}$ as $\{\overline{s}_i\}_{i\in Q}$, and for each $i \in Q$ define $\underline{s}_i = \overline{s}_i \cap \underline{s}$. Observe that for $i \neq j$ in Q, we have that $l(\overline{s}_i) \neq l(\overline{s}_j)$, because any two signal realizations in $C(\pi)$ have different relative likelihood vectors. Note that there is some $i \in Q$ for which $l(\overline{s}_i) \neq l(\underline{s}_i)$; let $Q' \subseteq Q$ be the (non-empty) set of indices i at which $l(\overline{s}_i) \neq l(\underline{s}_i)$.

Claim 1. If there exists $\mu_0 \in \Delta^o(\Omega)$ such that

$$\left(\sum_\omega \sum_{i\in Q} p^\omega(\overline{s}_i)\frac{\sum_{\omega'}\mu_0^{\omega'}p^{\omega'}(\underline{s}_i)}{\sum_{\omega'}\mu_0^{\omega'}p^{\omega'}(\overline{s}_i)}\right) \neq \sum_\omega p^\omega(\underline{s}) \tag{1}$$

then $\neg(\pi\mathcal{M}\pi')$.

To prove the claim, first observe that $C(\pi)\mathcal{M}\pi'$ if and only if for all $\overline{s} \in C(\pi)$, $\underline{s} \in \pi'$, we have $\mathbb{E}\left[\tilde{\mu}_{C(\pi)}|\underline{s}\right] = \mu_{\underline{s}}$, i.e., for all ω:

$$\sum_{i\in Q}\sum_{\omega'} \underbrace{\frac{\mu_0^{\omega'}p^{\omega'}(\underline{s})}{\sum_{\omega''}\mu_0^{\omega''}p^{\omega''}(\underline{s})}}_{=\mu_{\underline{s}}^{\omega'}=Pr(\omega'|\underline{s})} \underbrace{\frac{p^{\omega'}(\underline{s}_i)}{p^{\omega'}(\underline{s})}}_{=Pr(\overline{s}_i|\underline{s},\omega')} \underbrace{\frac{\mu_0^\omega p^\omega(\overline{s}_i)}{\sum_{\omega''}\mu_0^{\omega''}p^{\omega''}(\overline{s}_i)}}_{=\mu_{\overline{s}_i}^\omega} - \frac{\mu_0^\omega p^\omega(\underline{s})}{\sum_{\omega''}\mu_0^{\omega''}p^{\omega''}(\underline{s})} = 0$$

$$\iff \frac{\mu_0^\omega}{\sum_{\omega''}\mu_0^{\omega''}p^{\omega''}(\underline{s})}\left(\sum_{i\in Q} p^\omega(\overline{s}_i)\frac{\sum_{\omega'}\mu_0^{\omega'}p^{\omega'}(\underline{s}_i)}{\sum_{\omega''}\mu_0^{\omega''}p^{\omega''}(\overline{s}_i)} - p^\omega(\underline{s})\right) = 0.$$

This expression holds for all $\mu_0 \in \Delta^o(\Omega)$ if and only if the term in parentheses is zero for all ω. Summing across ω gives the result.

Importantly, the RHS of (1) does not depend on μ_0. So we can guarantee that there exists an interior prior μ_0 at which the two sides are not equal as long as the LHS is not constant in μ_0. Rewriting sums as dot products and simplifying further, we get the following implication.

Claim 2. Let $H_i(\mu) : [0,1]^N \to \mathbb{R}$ be defined as

$$H_i(\mu) \equiv \left(\sum_\omega p^\omega(\underline{s}_i) \right) \frac{\mu \cdot l(\underline{s}_i)}{\mu \cdot l(\overline{s}_i)}. \tag{2}$$

Let $H(\mu) \equiv \sum_{i \in Q'} H_i(\mu)$. If $H(\mu)$ is non-constant over the domain $\mu \in \Delta^o(\Omega)$, then $\neg(\pi \mathcal{M} \pi')$.[21]

We can consider two exhaustive and mutually exclusive cases.

Case 1: Q' is a singleton, which can be written as $\boldsymbol{Q' = \{i'\}}$; and $\boldsymbol{l(\overline{s}_{i'}) = \mu_u}$. In this case, $\mu \cdot l(\overline{s}_{i'}) = 1/N$ for all $\mu \in \Delta^o(\Omega)$, and hence $H(\mu) = N\left(\sum_\omega p^\omega(\underline{s}_{i'})\right) \mu \cdot l(\underline{s}_{i'})$. Moreover, because $l(\underline{s}_{i'}) \neq l(\overline{s}_{i'}) = \mu_u$, it also holds that $\mu \cdot l(\underline{s}_{i'})$ is linear and non-constant in μ. Hence, $H(\mu)$ is non-constant over the domain $\mu \in \Delta^o(\Omega)$.

Case 2: There exists some $\boldsymbol{i \in Q'}$ such that $\boldsymbol{l(\overline{s}_i) \neq \mu_u}$. We will find a direction $d_* \in \mathbb{R}^N$ with $\sum_\omega d_*^\omega = 0$ such that $H(\mu_u + \delta d_*)$ is nonconstant in δ in the neighborhood of $\delta = 0$, which will complete the proof.

Let

$$\hat{i} \in \arg \max_{\{i \in Q' | l(\overline{s}_i) \neq l(\underline{s}_i)\}} \| l(\overline{s}_i) - \mu_u \| \tag{3}$$

where $\| \cdot \|$ denotes the Euclidean norm, and set $d = \mu_u - l(\overline{s}_{\hat{i}})$. From the definition of $\hat{i}$, and the

[21] The function $H_i(\mu)$ may be undefined at points μ that lead to a 0 denominator, but H_i (and therefore H) is defined everywhere on $\Delta^o(\Omega)$.

fact that $\mu \cdot \mu_u = l(s) \cdot \mu_u = 1/N$ for all s and $\mu \in \Delta(\Omega)$, we have

$$\begin{aligned}
(\mu_u + \delta d) \cdot l(\overline{s}_i) &= \mu_u \cdot l(\overline{s}_i) + \delta \left(\mu_u - l\left(\overline{s}_{\hat{i}}\right)\right) \cdot l(\overline{s}_i) \\
&= 1/N - \delta (l(\overline{s}_{\hat{i}}) - \mu_u) \cdot (l(\overline{s}_i) - \mu_u) \\
&\geq 1/N - \delta \| l(\overline{s}_{\hat{i}}) - \mu_u \| \| l(\overline{s}_i) - \mu_u \| \\
&\geq 1/N - \delta \| l(\overline{s}_{\hat{i}}) - \mu_u \|^2 \\
&= (\mu_u + \delta d) \cdot l(\overline{s}_{\hat{i}}).
\end{aligned}$$

(Note that the inequality is strict if $i \neq \hat{i}$, because then $l(\overline{s}_i) \neq l(\overline{s}_{\hat{i}})$.) Now take δ^* to be the unique δ such that this last expression is equal to zero, i.e., $\delta^* \equiv 1/(N \| l(\overline{s}_i) - \mu_u \|^2)$, and for all $i \neq \hat{i}$, we have that $(\mu_u + \delta^* d) \cdot l(\overline{s}_i) > 0$.

If in addition we have

$$(\mu_u + \delta^* (\mu_u - l(\overline{s}_{\hat{i}}))) \cdot l(\underline{s}_{\hat{i}}) \neq 0,$$

then take $d^* = d$. Otherwise, let d' be the projection of $l(\underline{s}_{\hat{i}}) - \mu_u$ onto the null space of $l(\overline{s}_{\hat{i}}) - \mu_u$, and note that $d' \neq 0$ because $l(\underline{s}_{\hat{i}}) \neq l(\overline{s}_{\hat{i}})$ (per (3)). Hence, for ϵ sufficiently small,

$$\begin{aligned}
&(\mu_u + \delta^* (d + \epsilon d')) \cdot l(\overline{s}_i) > 0 \ \forall i \neq \hat{i}; \\
&(\mu_u + \delta^* (d + \epsilon d')) \cdot l(\underline{s}_{\hat{i}}) \neq 0; \\
&(\mu_u + \delta^* (d + \epsilon d')) \cdot l(\overline{s}_{\hat{i}}) = 0.
\end{aligned}$$

We then set $d^* = d + \epsilon d'$.[22]

Again using the fact that $\mu_u \cdot \mu = 1/N$ for any μ with $\sum_\omega \mu = 1$, we have that for all $\delta \in [0, \delta^*)$ and for all i, $(\mu_u + \delta d^*) \cdot l(\overline{s}_i) > 0$. Hence, $H(\mu_u + \delta d^*)$ is finite for all $\delta \in [0, \delta^*)$, since the denominators of 2 are nonzero for every $i \in Q'$; and because at δ^* the numerator at $\hat{i}$ is non-zero, the denominator at $\hat{i}$ is zero, and the denominators at $i \neq \hat{i}$ are all non-zero, we have that

$$\lim_{\delta \nearrow \delta^*} H(\mu_u + \delta d^*) = \pm\infty.$$

[22]Note that $\mu_u + \delta^* d^*$ need not be a probability vector in $\Delta(\Omega)$; the sum of components is 1, but it may have negative components. The rest of the proof shows that H is non-constant on the restricted domain of $\Delta^o(\Omega)$.

Figure 6: The martingale property can depend on priors

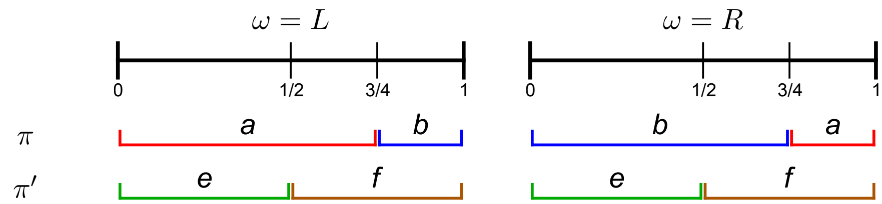

At prior μ_0 on $Pr(\omega = R)$, it holds that $\mu_a = \frac{\mu_0}{3-2\mu_0}$; $\mu_b = \frac{3\mu_0}{2\mu_0+1}$; $\mu_e = \mu_0$; and $\mu_f = \mu_0$. Moreover, $\mathbb{E}\left[\tilde{\mu}_\pi | e\right] = \mu_0\mu_a + (1-\mu_0)\mu_b = \frac{\mu_0(8\mu_0^2-14\mu_0+9)}{-4\mu_0^2+4\mu_0+3}$ and $\mathbb{E}\left[\tilde{\mu}_\pi | f\right] = \frac{1}{2}\mu_a + \frac{1}{2}\mu_b = \frac{\mu_0(5-2\mu_0)}{-4\mu_0^2+4\mu_0+3}$. It is easy to verify that $\mathbb{E}\left[\tilde{\mu}_\pi | e\right] = \mu_e$ and $\mathbb{E}\left[\tilde{\mu}_\pi | f\right] = \mu_f$ if $\mu_0 = 1/2$, and that these equalities do not hold at any other interior prior $\mu_0 \in (0, 1)$.

Finally, note that $H(\mu_u + \delta d^*)$ is a rational function of δ (and therefore analytic in δ), is defined for all $\delta \in [0, \delta^*]$, and only has a singularity at $\delta = \delta^*$. Thus, H must be non-constant in δ on every open set in the interval $[0, \delta^*]$, and in particular, it is non-constant in the neighborhood of $\delta = 0$.

A.4 Martingale property can depend on priors

We define the martingale relation $\mathcal{M}$ as follows: $\pi\mathcal{M}\pi'$ if $\mathbb{E}\left[\tilde{\mu}_\pi | \tilde{s}_{\pi'}\right] = \tilde{\mu}_{\pi'}$ holds for all priors μ_0. In this section, we note that there are signals for which $\mathbb{E}\left[\tilde{\mu}_\pi | \tilde{s}_{\pi'}\right] = \tilde{\mu}_{\pi'}$ holds at some interior priors but not others. See Figure 6 for an example. Hence, the for-all quantifier was important.

In Figure 6, π is informative about the state, while π' is informative about how informative π is. In particular, signal π realizes either a, indicating a higher chance that $\omega = L$ and yielding $\mu_a < \mu_0$ (with beliefs in $[0, 1]$ denoting the probability of $\omega = R$); or b, indicating a higher chance that $\omega = R$ and yielding $\mu_b > \mu_0$. When $e \in \pi'$ is realized, π is in fact perfectly informative: $a \in \pi$ implies that $\omega = L$ for sure, and $b \in \pi$ implies $\omega = R$. And when $f \in \pi'$ is realized, π is uninformative: a and b both have the same conditional likelihood across states. Since π' is uninformative about the state itself, though, the posterior after observing either realization from π' is always equal to the prior: $\mu_e = \mu_f = \mu_0$. But the expectation of $\tilde{\mu}_\pi$ (the posterior of π) given either realization of π' is equal to the prior only when beliefs are degenerate, or when the prior is uniform at $\mu_0 = 1/2$. This is easiest to see by considering $\mathbb{E}[\tilde{\mu}_\pi | f]$. Conditional on $f \in \pi'$, the signal π realizes a and b with equal probability, independently of the prior; but the posterior beliefs μ_a and μ_b are not equally distant from the prior. For instance, at priors $\mu_0 \in (0, 1/2)$, it holds that $\mu_b - \mu_0 > \mu_0 - \mu_a$.

Given that this martingale property can depend on the prior, we see that there is an alternative "martingale relation" on signals that we could have defined. Define the existence-martingale relation, denoted $\mathcal{M}^{\exists}$, as follows: $\pi\mathcal{M}^{\exists}\pi'$ if there exists an interior prior μ_0 at which $\mathbb{E}\left[\tilde{\mu}_\pi|\tilde{s}_{\pi'}\right] = \tilde{\mu}_{\pi'}$. It is easy to see that $\mathcal{M} \subsetneq \mathcal{M}^{\exists}$: the fact that $\mathcal{M} \subseteq \mathcal{M}^{\exists}$ follows immediately from definitions (for-all implies there-exists), and $\mathcal{M} \neq \mathcal{M}^{\exists}$ follows from the example in Figure 6. Moreover, it turns out that $\mathcal{M}^{\exists} \subsetneq \mathcal{B}$. The fact that $\mathcal{M}^{\exists} \subseteq \mathcal{B}$ can be seen by noting that if $\pi\mathcal{M}^{\exists}\pi'$, then for an interior prior μ_0 at which $\mathbb{E}\left[\tilde{\mu}_\pi|\tilde{s}_{\pi'}\right] = \tilde{\mu}_{\pi'}$, it holds that $\langle\pi\rangle$ is a mean-preserving spread of $\langle\pi'\rangle$; and if the posteriors of π are a mean-preserving spread of those of π' at any one interior prior, then they are a mean-preserving spread at all priors, i.e., $\pi\mathcal{B}\pi'$. The fact that $\mathcal{M}^{\exists} \neq \mathcal{B}$ can be established by observing that, in Figure 5, $\pi_{\mathcal{B}}\mathcal{B}\pi_0$, but, as can be directly calculated, $\neg(\pi_{\mathcal{B}}\mathcal{M}^{\exists}\pi_0)$.[23] Hence, we can expand our summary of the ranking of the relations to $\overline{\mathcal{R}} = \mathcal{R} \subsetneq \overline{\mathcal{S}} = \overline{\mathcal{M}} = \overline{\mathcal{B}} \subsetneq \mathcal{S} \subsetneq \mathcal{M} \subsetneq \mathcal{M}^{\exists} \subsetneq \mathcal{B}$.

A.5 An alternative characterization of $\mathcal{S}$, $\mathcal{M}$,$\mathcal{B}$

Remark 4.7 stated the following characterizations of $\mathcal{S}$, $\mathcal{M}$, and $\mathcal{B}$.

1. $\pi\mathcal{S}\pi'$ if and only if $\exists\pi^*$ s.t. $C(\pi^*) = C(\pi)$, $\pi^*\mathcal{R}\pi$, and $\pi^*\mathcal{R}\pi'$. (We can take $\pi^* = \pi \vee \pi'$.)

2. $\pi\mathcal{M}\pi'$ if and only if $\exists\pi^*$ s.t. $C(\pi^*) = C(\pi)$ and $\pi^*\mathcal{R}\pi'$.

3. $\pi\mathcal{B}\pi'$ if and only if $\exists\pi^*$ s.t. $\pi^* \sim \pi$ and $\pi^*\mathcal{R}\pi'$.

Part 3 had already been presented in Remark 4.2.

To see the only if direction of Part 1, first suppose that $\pi S\pi'$, and let $\pi^* = \pi \vee \pi'$. We see that $C(\pi^*) = C(\pi \vee \pi') = C(\pi)$, with the second equality following from the fact that $\pi\mathcal{S}\pi' \Leftrightarrow \tilde{\mu}_{\pi\vee\pi'} = \tilde{\mu}_\pi$. And by construction, $\pi^*\mathcal{R}\pi$ and $\pi^*\mathcal{R}\pi'$. Next, consider the if direction. Suppose that $\exists\pi^*$ s.t. $C(\pi^*) = C(\pi)$, $\pi^*\mathcal{R}\pi$, and $\pi^*\mathcal{R}\pi'$. The latter two properties imply that $\pi^*\mathcal{R}(\pi \vee \pi')$. Therefore, $C(\pi^*)\mathcal{B}C(\pi \vee \pi')\mathcal{B}C(\pi)$. This fact, coupled with $C(\pi) = C(\pi^*)$ implies that $C(\pi^*) \sim C(\pi \vee \pi') \sim C(\pi)$. Finally, $C(\pi \vee \pi') \sim C(\pi)$ implies that $\pi\mathcal{S}\pi'$.

Part 2 follows from Part 1 combined with the observation (Remark 4.6, Proposition 2) that $\pi\mathcal{M}\pi'$ if and only if $C(\pi)\mathcal{S}\pi'$. First, we suppose that $\pi\mathcal{M}\pi'$, and show that $\exists\pi^*$ s.t. $C(\pi^*) = C(\pi)$

[23] Writing all beliefs in terms of $Pr(\omega = R)$, we have $\mu_k = \frac{\mu_0 \cdot 1/3}{\mu_0 \cdot 1/3 + (1-\mu_0)\cdot 2/3}$ at $k \in \pi_0$, along with $\mu_a = \frac{\mu_0 \cdot 1/4}{\mu_0 \cdot 1/4 + (1-\mu_0)\cdot 3/4}$ and $\mu_b = \frac{\mu_0 \cdot 3/4}{\mu_0 \cdot 3/4 + (1-\mu_0)\cdot 1/4}$ at a and b in $\pi_{\mathcal{B}}$. Conditional on realization $k \in \pi_0$, the expected belief at $\pi_{\mathcal{B}}$ is given by $\mathbb{E}[\tilde{\mu}_{\pi_{\mathcal{B}}}|k] = (1-\mu_k)\mu_a + \mu_k(\mu_a \cdot 3/4 + \mu_b \cdot 1/4)$. The martingale property at prior μ_0 holds only if $\mathbb{E}[\tilde{\mu}_{\pi_{\mathcal{B}}}|k] - \mu_k = 0$, but the LHS simplifies to $-\frac{\mu_0(1-\mu_0)}{6+5\mu_0-12\mu_0^2+4\mu_0^3}$, which has no zeroes for $\mu_0 \in (0, 1)$.

Figure 7: Martingale is not transitive

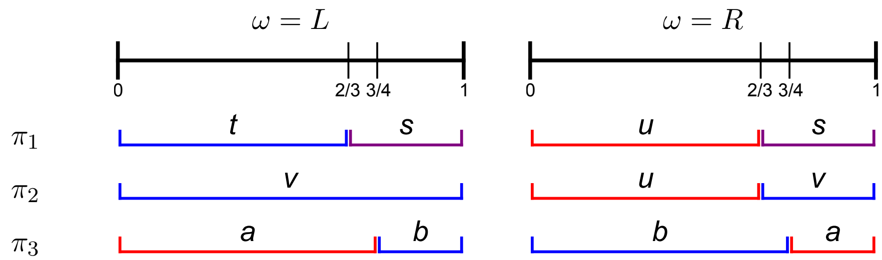

$\pi_1 \mathcal{M} \pi_2$ and $\pi_2 \mathcal{M} \pi_3$ but $\neg(\pi_1 \mathcal{M} \pi_3)$.

and $\pi^* \mathcal{R} \pi'$. This holds because $\pi \mathcal{M} \pi'$ implies $C(\pi) \mathcal{S} \pi'$, which implies by Part 1 that there exists π^* (including $\pi^* = \pi \vee \pi'$) satisfying these conditions. Next, we suppose that $\exists \pi^*$ s.t. $C(\pi^*) = C(\pi)$ and $\pi^* \mathcal{R} \pi'$, and show that $\pi \mathcal{M} \pi'$. Take such a π^*, and observe that it satisfies $\pi^* \mathcal{R} C(\pi^*)$ and $C(\pi^*) = C(\pi)$, and hence $\pi^* \mathcal{R} C(\pi)$. Because π^* by definition also satisfies $C(\pi^*) = C(C(\pi))$ (since $C(C(\pi)) = C(\pi)$) and $\pi^* \mathcal{R} \pi'$, Part 1 implies that $C(\pi) \mathcal{S} \pi'$, which then implies $\pi \mathcal{M} \pi'$.

A.6 Martingale dominance is not transitive

Consider Figure 7. We see that $\pi_1 \mathcal{M} \pi_2$ because $\pi_1 \mathcal{R} \pi_2$; and $\pi_2 \mathcal{M} \pi_3$ because $\pi_2 \mathcal{S} \pi_3$, which may not be immediately obvious.[24] However, it is not the case that $\pi_1 \mathcal{M} \pi_3$: with a prior of $\mu_0 = 1/2$ probability on $\omega = R$, we have that $\mu_a = 1/4$ while $\mathbb{E}[\tilde{\mu}_{\pi_1} | a] = \frac{2}{3} \cdot 0 + \frac{1}{3} \cdot \frac{1}{2} = \frac{1}{6}$.

[24] To see that $\pi_2 \mathcal{S} \pi_3$, or in other words that $\pi_2 \vee \pi_3$ induces the same beliefs as π_2, observe that $\pi_2 = \{u, v\}$ while $\pi_2 \vee \pi_3 = \{a \cap v = a, b \cap v, b \cap u = u\}$. So it suffices to show that the likelihood ratios of (and thus the beliefs at) a and of $b \cap v$ match that of v, which indeed they do: $Pr(v|L)/Pr(v|R) = Pr(b \cap v|L)/Pr(b \cap v|R) = Pr(a|L)/Pr(a|R) = 3$.