

Comparison of Decisions under Unknown Experiments

Andrew Caplin

New York University

Daniel Martin

Northwestern University

We take the perspective of an econometrician who wants to determine which of two experiments provides higher expected utility but only knows the decisions under each experiment. To compare these decisions, the econometrician must make inferences about what the experiment might have been for each set of decisions. We provide a necessary and sufficient condition that identifies when every experiment consistent with one set of decisions has a higher value of information than every experiment consistent with the other set of decisions.

I. Introduction

There are many facts about the world (or “states of the world”) that can be payoff relevant for decision makers (DMs). For example, their payoffs can depend on the fundamentals of a stock, the effectiveness of a vaccine, characteristics of a health plan, and so on. These facts can be presented to

We thank Edwin Muñoz Rodríguez for excellent research assistance and Jose Apesteguia, Victor Aguiar, Miguel Ballester, Henrique de Oliveira, Alex Frankel, Ben Handel, Emir Kamenica, Matt Kovach, Josh Schwartzstein, Charlie Sprenger, Bruno Strulovici, Dmitry Taubinsky, Gerelt Tserenjigmid, and seminar audiences at Alicante, Chicago, Northwestern, and University of California, Santa Barbara, for valuable feedback. Caplin thanks the NOMIS and Sloan Foundations for support. All mistakes are our own. This paper was edited by Emir Kamenica.

Electronically published August 31, 2021

Journal of Political Economy, volume 129, number 11, November 2021.

© 2021 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/716104>

DMs in a number of different ways by better-informed parties. For instance, advisors, firms, news networks, and governments can choose to selectively allocate information about the facts, as in Bayesian persuasion (e.g., Kamenica and Gentzkow 2011) or voluntary information disclosure (e.g., Milgrom 1981).¹ They can also choose the format that this information takes, making it easier or harder for DMs to understand.² In both of these cases, it has been well documented that the way the better-informed party chooses to present the facts can strongly influence how well informed DMs are when they make their choices.

We take the perspective of an econometrician who wants to compare different ways of presenting the facts based on how valuable that information was for DMs. For example, the econometrician might want to determine whether the advice from one financial advisor helped DMs make better portfolio allocations than the advice provided by a different advisor, whether watching one news program helped DMs choose better health behaviors than watching another news program (e.g., Bursztyn et al. 2020), or whether one description of fees led to better health plan choices made than a different description of fees (e.g., Bhargava, Loewenstein, and Sydnor 2017).

We model a presentation of the facts as an *experiment* (a joint distribution of signals and states). Traditionally, experiments have been used to model physical activities where observing signals is easy (e.g., drilling for oil or performing a medical test). However, in our application it is more challenging for the econometrician to observe the experiment itself. For instance, it can be hard to know what an advisor said to their clients if there are privacy concerns, advice is proprietary, or it is challenging to codify the advice provided. Further, even if we know the exact information they provided to their clients, it might be challenging to know what the clients understood about the facts based on that information. Yet it is often possible for the econometrician to observe the actions taken under each experiment. For instance, many data sets contain the stocks that were bought, the vaccines that were taken, the health plans that were selected, and so on.

Because of this, we assume that the econometrician knows only the actions taken under each experiment and nothing about the experiment itself (either the signal structure or the signal realizations). For example, all the econometrician might know about a particular experiment is that

¹ For reviews of the Bayesian persuasion literature, see Kamenica (2019); for reviews of the voluntary disclosure literature, see Dranove and Jin (2010); and for reviews of the disclosure experimental literature, see Jin, Luca, and Martin (2015).

² For example, see Hastings and Tejeda-Ashton (2008), Choi, Laibson, and Madrian (2009), Abeler and Jäger (2015), Carrera and Villas-Boas (2015), Ericson and Starc (2016), Jin, Luca, and Martin (2018), Esponda and Vespa (2019), Clippel and Rozen (2020), and Carpenter et al. (2021).

it results in the following joint distribution over actions (a_1, a_2, a_3) and states $(\omega_1, \omega_2, \omega_3)$:

$$P_g = \begin{pmatrix} \frac{22}{100} & 0 & 0 \\ 0 & \frac{22}{100} & \frac{18}{100} \\ 0 & \frac{18}{100} & \frac{22}{100} \end{pmatrix} \begin{matrix} a_1 \\ a_2 \\ a_3 \end{matrix}.$$

In practice, this joint distribution could be the frequency a stock is bought when it has certain fundamentals, a vaccine is taken when it has certain effectiveness, a preferred provider organization health plan is chosen when it has certain benefits, a television model is purchased when a competing model is on sale, a loan is made to someone who will default, a test is ordered for someone who actually has a disease, and so on.³

Thus, the econometrician wants to be able to determine which of two experiments provides higher expected utility for the DM based solely on the joint distributions of actions and states under each experiment. For example, does the experiment that produced P_g provide higher expected utility than the experiment that produced P_h ?

$$P_g = \begin{pmatrix} \frac{22}{100} & 0 & 0 \\ 0 & \frac{22}{100} & \frac{18}{100} \\ 0 & \frac{18}{100} & \frac{22}{100} \end{pmatrix} \begin{matrix} a_1 \\ a_2 \\ a_3 \end{matrix} \quad \text{and} \quad P_h = \begin{pmatrix} \frac{10}{100} & \frac{20}{100} & \frac{20}{100} \\ \frac{5}{100} & \frac{20}{100} & 0 \\ \frac{5}{100} & 0 & \frac{20}{100} \end{pmatrix} \begin{matrix} a_1 \\ a_2 \\ a_3 \end{matrix}.$$

When the DM's utility function is known, answering this question is easy because the econometrician can directly calculate expected utility using the probability of each action and state. However, we are interested in whether the econometrician can rank experiments without knowing the DM's utility function.

To accomplish this, the econometrician must make inferences about what the experiment might have been for each set of decisions. We use the same maintained assumptions as Blackwell (1953): for a given u , experiment π_g is *consistent* with P_g if P_g maximizes expected utility among the joint distributions of actions and states feasible under experiment π_g .

³ In all of these cases, the state impacts utility, the DM may not be fully informed about the state, and the econometrician knows the state.

Using the set of consistent experiments, we define a binary relation $\succsim_v$ that enables the econometrician to rank decisions. We say that $P_g \succsim_v P_h$ if for every utility function u , every experiment consistent with P_g has a higher value of information than every experiment consistent with P_h .

To characterize the relation $\succsim_v$, we leverage two features of the problem. First, for a given utility function u , every experiment consistent with a given P_g has the same value of information, which is the expected utility provided by P_g for that utility function. Second, there are utility functions for which no experiments are consistent with P_g or P_h . Because the condition for $P_g \succsim_v P_h$ is trivially satisfied for such utility functions, the econometrician does not need to consider them when comparing P_g and P_h .

Thus, the econometrician can safely conclude that $P_g \succsim_v P_h$ if they can rule out enough utility functions to ensure that P_g provides higher expected utility than P_h for all remaining utility functions. We establish the general logic by first showing that restricting utility functions has a clean geometric structure in the space of outcome lotteries. We then build on this structure to produce a necessary and sufficient condition for $P_g \succsim_v P_h$. The necessity and sufficiency of this condition follows as a direct consequence of the separating hyperplane theorem. In technical terms, this condition requires that a vector representing the difference in outcomes between P_g and P_h falls in the cone generated by the restrictions for a utility function to be consistent with P_g and P_h . This corresponds to solving a system of linear equations, so it is simple to check, and MATLAB programs that implement it are provided.⁴

The rest of the paper is organized as follows. In section II, we provide our framework, formally define our relation, and show how having a second set of decisions or knowing outcomes can allow the econometrician to rule out enough utility functions for decisions to be ordered according to the relation. In section III, we first introduce our geometric representation of ruling out utility functions and then leverage this representation to identify a testable condition that is both necessary and sufficient for decisions to be ranked according to $\succsim_v$. Section IV concludes by discussing related literature.

II. Framework

For each presentation of the facts, we assume that the DM starts with an interior prior over a finite set of states Ω given by $\mu \in \Delta(\Omega)$. The DM receives a signal realization, and as is now standard, a signal realization is represented by the posterior belief $\gamma \in \Delta(\Omega)$ that it generates (see Kamenica and Gentzkow 2011). This process is summarized by an experiment π that is a joint distribution over states Ω and posteriors $\Delta(\Omega)$, and

⁴ Programs are available at <https://github.com/danieljosephmartin>.

for notational simplicity we assume that the experiment has finite support over $\Delta(\Omega)$.

Given a posterior belief generated by an experiment, we assume that the DM implements a decision rule $\sigma: \Delta(\Omega) \rightarrow \Delta(A)$, where A represents a finite set of actions. The DM receives outcome $x(a, \omega) \in X$ when action $a \in A$ is chosen in state $\omega \in \Omega$, and the decision rule maximizes expected utility based on a utility function $u: X \rightarrow \mathbb{R}$.⁵

We consider an econometrician who wants to compare the expected utility provided by two experiments (two presentations of the facts). We assume that the unconditional probability of states is the same across the experiments, and our interpretation of this assumption is that a presentation of the facts cannot change the facts themselves. For instance, a financial advisor cannot change the actual financial conditions of the institutions their clients might invest in. In the case of choosing which way to describe fees, the fees themselves cannot be changed. This assumption is not required for our characterization, but it simplifies our analysis.

All the econometrician knows about these experiments is the joint distributions over actions and states they generate. We refer to an arbitrary joint distribution of actions and states as $P_f \in \{P_g, P_h\}$ and denote $P_f(a, \omega)$ as the probability of choosing action a and being in state ω for P_f .⁶

As with other stochastic choices, each P_g and P_h can be interpreted as watching the DM face a decision infinitely often. In practice, one might estimate it from repeated but finite choice data or by looking at a population rather than an individual, as in the literature on discrete choice following McFadden (1973). For notational simplicity, we assume that all outcomes can be obtained by taking some action in some state and that for each distribution of actions and states, each action is chosen in some state and an action is chosen in each state.⁷ This joins a growing literature that considers stochastic choice to be essential for studying information and utility (e.g., Manzini and Mariotti 2014; Apesteguia and Ballester 2018).

For what follows, it is not necessary for the econometrician to also know the outcome received from taking each action in each state, as it is without loss of generality for the econometrician to arbitrarily assign a distinct outcome to every action in every state. However, the presence of an outcomes space allows us to accommodate cases where the econometrician

⁵ The outcome space is also finite, and we denote its cardinality as M . It has generic element x or x_m .

⁶ The joint distributions P_g and P_h are state-dependent stochastic choice data, which were proposed for information-theoretic revealed preferences by Caplin and Martin (2015).

⁷ Our results would still go through without these assumptions, but doing so would require carefully specifying the support of each distribution of actions and states and adding technical regularity conditions, which would necessitate several pieces of additional notation while adding little additional economic insight.

knows that utility is equal across some states and actions. For example, the econometrician might know that an action is “safe” because it yields the same outcome in every state (for an example, see sec. III.A).

A. Comparison of Decisions

The econometrician would like to compare joint distributions of actions and states P_g and P_h based on the value of the information provided by the experiments that generated them. Without knowing anything about the structure of these experiments, the econometrician must determine the experiments that are consistent with P_g and P_h .

We define consistency using the same maintained assumptions as in Blackwell (1953). An experiment π_f is consistent with P_f if P_f maximizes expected utility among distributions of actions and states feasible under π_f .

A joint distribution of actions and states P_f is *feasible* under π_f if there exists a decision rule $\sigma_f: \Delta(\Omega) \rightarrow \Delta(A)$ such that

$$P_f(a, \omega) = \sum_{\gamma \in \text{supp}(\pi_f)} \pi_f(\gamma, \omega) \sigma_f(a|\gamma),$$

where the set of possible posterior beliefs is given by $\text{supp}(\pi_f)$. Given u , the highest expected utility for experiment π_f is

$$V(u, \pi_f) = \max_{P \in \Phi(\pi_f)} \sum_{a \in A} \sum_{\omega \in \Omega} P(a, \omega) u(x(a, \omega)),$$

where $\Phi(\pi_f)$ represents the set of all distributions of actions and states feasible under π_f . The function $V(u, \pi_f)$ is also known as the *value of information* for experiment π_f given utility function u .⁸ Thus, P_f is *consistent* with π_f if

$$P_f \in \operatorname{argmax}_{P \in \Phi(\pi_f)} \sum_{a \in A} \sum_{\omega \in \Omega} P(a, \omega) u(x(a, \omega)).$$

To allow the econometrician to rank distributions of actions and states based on the value of information for consistent experiments, we formally define the relation $\succeq_V$ as

$$P_g \succeq_V P_h$$

if for every u ,

$$V(u, \pi_g) \geq V(u, \pi_h)$$

for every π_g consistent with P_g and every π_h consistent with P_h .

⁸ An alternative way to define the value of information is as the improvement over the utility from taking actions at prior beliefs (see Frankel and Kamenica 2018; Lara and Gossner 2020). Since the prior is fixed across experiments in our framework, this definition would provide the same relative welfare assessments.

There are two features of the problem that help us in characterizing this relation. First, for a given u , every experiment π_f consistent with P_f has the same value of information, which is the expected utility provided by P_f :

$$V(u, \pi_f) = \sum_{\omega \in \Omega} \sum_{a \in A} P_f(a, \omega) u(x(a, \omega)).$$

Second, we need to consider only those utility functions for which there exist experiments consistent with P_g and P_h , as the condition for $P_g \succeq_v P_h$ is trivially satisfied for all other utility functions. Putting this together, $P_g \succeq_v P_h$ if and only if for all u for which there are experiments consistent with P_g and P_h ,

$$\sum_{\omega \in \Omega} \sum_{a \in A} P_g(a, \omega) u(x(a, \omega)) \geq \sum_{\omega \in \Omega} \sum_{a \in A} P_h(a, \omega) u(x(a, \omega)).$$

B. Ruling Out Utility Functions

To operationalize this restatement of $P_g \succeq_v P_h$, the econometrician needs to identify those u for which there are experiments consistent with P_g and P_h or, equivalently, to rule out those u for which there does not exist an experiment consistent with P_g or P_h .

For utility function u , there does not exist an experiment consistent with $P_f \in \{P_g, P_h\}$ if it is possible to improve utility by making a wholesale switch from any chosen action to another action. If the DM can improve utility by switching to an action $b \in A$ at all posteriors where they chose action $a \in A$, this means that whatever decision rule pairs with an experiment to make P_f feasible cannot maximize expected utility.

This is demonstrated in the following simple example, where utility equal to one when action a_3 is taken in state ω_1 , action a_2 is taken in state ω_2 , and action a_1 is taken in state ω_3 :

$$u = \begin{pmatrix} \omega_1 & \omega_2 & \omega_3 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \begin{matrix} a_1 \\ a_2 \\ a_3 \end{matrix} \quad \text{and} \quad P_f = \begin{pmatrix} \omega_1 & \omega_2 & \omega_3 \\ \frac{22}{100} & 0 & 0 \\ 0 & \frac{22}{100} & \frac{18}{100} \\ 0 & \frac{18}{100} & \frac{22}{100} \end{pmatrix} \begin{matrix} a_1 \\ a_2 \\ a_3 \end{matrix}.$$

For this utility function, P_f is not consistent with any experiment because the DM can improve utility by making a wholesale switch to choosing a_3 when a_1 was chosen. When the DM chooses a_1 , they get a utility of zero

for certain, but if they had chosen a_3 instead, the DM could have gotten a utility of one for certain.

Caplin and Martin (2015) formalize this logic by introducing the no improving action switches (NIAS) condition, which is a system of linear inequalities ensuring it is better not to make a wholesale switch from any chosen action a to any other action b . Utility function u satisfies NIAS for P_g and P_h if

$$\sum_{\omega \in \Omega} P_f(a, \omega) u(x(a, \omega)) \geq \sum_{\omega \in \Omega} P_f(a, \omega) u(x(b, \omega))$$

for all $P_f \in \{P_g, P_h\}$ and $a, b \in A$. The NIAS inequality for choosing a over b in P_f indicates that choosing a instead of b is optimal on average at the choice probabilities where a is chosen in P_f given the utility of the outcomes from choosing a instead of b .

As illustrated above, if a utility function u does not satisfy NIAS, then there does not exist an experiment consistent with P_f . Caplin and Martin (2015) show that the reverse is true as well. If u satisfies NIAS, then there always exists an experiment consistent with P_f for that u .⁹ Thus, the set of all u satisfying NIAS for P_g and P_h is precisely the set of u the econometrician should consider when comparing P_g and P_h .

We establish here a general feature of NIAS that enhances its analytical and computational tractability. The following lemma indicates which NIAS inequalities must hold with equality and which must hold strictly. In other words, it states that for $P_f \in \{P_g, P_h\}$ the NIAS inequality for choosing a over b holds with equality if and only if the outcomes associated with that NIAS inequality (the additional probability of each outcome gained by not switching from a to b) can be expressed as a nonpositive combination of the outcomes associated with other NIAS inequalities.

LEMMA 1. For every u that satisfies NIAS,

$$\sum_{\omega \in \Omega} P_f(a, \omega) u(x(a, \omega)) = \sum_{\omega \in \Omega} P_f(a, \omega) u(x(b, \omega)) \quad (1)$$

for $a, b \in A$ if and only if there exists a collection of N triples with generic element (P_n, a_n, b_n) having $P_n \in \{P_g, P_h\}$, $a_n \in A$, $b_n \in A$, and $(P_n, a_n, b_n) \neq (P_f, a, b)$ and nonpositive weights $w_1, \dots, w_N$ such that for every $x \in X$,

$$\begin{aligned} & \sum_{\omega \in \Omega} P_f(a, \omega) (\mathbf{1}_{\{x(a, \omega) = x\}} - \mathbf{1}_{\{x(b, \omega) = x\}}) \\ &= \sum_{n=1}^N w_n \left(\sum_{\omega \in \Omega} P_n(a_n, \omega) (\mathbf{1}_{\{x(a_n, \omega) = x\}} - \mathbf{1}_{\{x(b_n, \omega) = x\}}) \right), \end{aligned} \quad (2)$$

⁹ For instance, if u satisfies NIAS, then the revealed experiment for P_f is consistent with P_f for that u .

where $\mathbf{1}_{\{x(a,\omega)=x\}}$ is an indicator function that takes a value of one when the outcome from taking action a in state ω yields outcome x .

Proof. See the appendix. QED

C. Ruling Out Utility Functions to Rank Decisions

The following two examples demonstrate that NIAS can rule out enough utility functions to allow the value of information to be ranked between two distributions of actions and states.

1. Tracking Problems

We first consider “tracking” decision problems, in which the DM receives a state-specific outcome x_k if their action matches state ω_k and outcome x_B if they fail to match the action to the state. For this class of decision problems, the map $x(a, \omega)$ between actions, states, and outcomes is known by the econometrician and is given by

$$x(a_j, \omega_k) = \begin{cases} x_k & j = k, \\ x_B & j \neq k. \end{cases}$$

For the three-action and three-state version of this tracking problem, the map between actions, states, and outcomes can be represented as a matrix where actions a_1 – a_3 are given in the rows and states ω_1 – ω_3 are given in the columns:

$$\begin{matrix} & \omega_1 & \omega_2 & \omega_3 \\ \begin{pmatrix} x_1 & x_B & x_B \\ x_B & x_2 & x_B \\ x_B & x_B & x_3 \end{pmatrix} & a_1 \\ & a_2 \\ & a_3 \end{matrix}.$$

Imagine the following distribution of actions and states, which were given in the introduction:

$$P_g = \begin{pmatrix} \frac{22}{100} & 0 & 0 \\ 0 & \frac{22}{100} & \frac{18}{100} \\ 0 & \frac{18}{100} & \frac{22}{100} \end{pmatrix} \begin{matrix} a_1 \\ a_2 \\ a_3 \end{matrix} \quad \text{and} \quad P_h = \begin{pmatrix} \frac{10}{100} & \frac{20}{100} & \frac{20}{100} \\ \frac{5}{100} & \frac{22}{100} & 0 \\ \frac{5}{100} & 0 & \frac{20}{100} \end{pmatrix} \begin{matrix} a_1 \\ a_2 \\ a_3 \end{matrix}.$$

In the analysis that follows, we show that these distributions of actions and states reveal that the outcomes from matching actions to states are

“good” and the outcome from not matching actions is “bad.” Formally, this means that for all utility functions that rationalize P_g and P_h , $u(x_k) \geq u(x_B)$ for all $k \in \{1, 2, 3\}$. Given this, P_g will be revealed to have a higher value of information because the DM matches actions to states more often in every state. In terms of signal structures, it is as if P_g is generated by a DM with a signal structure that is perfectly informative about whether the state is ω_1 but is not as informative about the other states as the signal structure that generated P_h .

Without loss of generality we set $u(x_B) = 0$, so we can compute the value of information for P_g as

$$\frac{20}{100} u(x_0) + \frac{22}{100} u(x_1) + \frac{22}{100} u(x_2)$$

and for P_h as

$$\frac{10}{100} u(x_0) + \frac{20}{100} u(x_1) + \frac{20}{100} u(x_2).$$

Clearly, if $u(x_1)$, $u(x_2)$, or $u(x_3)$ are revealed to be nonnegative for all rationalizing utility functions, then P_g is revealed to have a higher value of information.

The fact that these utilities are nonnegative can be established through the NIAS inequalities for P_g . The NIAS inequality for P_g for a_1 chosen over action a_2 gives $u(x_1) \geq 0$ because

$$\begin{aligned} \sum_{\omega \in \Omega} P_g(a_1, \omega) u(x(a_1, \omega)) &\geq \sum_{\omega \in \Omega} P_g(a_2, \omega) u(x(a_2, \omega)), \\ \frac{20}{100} u(x_1) &\geq \frac{20}{100} u(x_B) = 0. \end{aligned}$$

Likewise, the NIAS inequality for P_g for a_2 chosen over action a_1 gives $u(x_2) \geq 0$, and the NIAS inequality for a_3 chosen over action a_1 gives $u(x_3) \geq 0$.

2. Problems with Distinct Outcomes

Next, we consider a common class of decision problems in which every action yields a distinct outcome in every state, so that $x(a, \omega) \neq x(b, \nu)$ if $a \neq b$ or $\omega \neq \nu$. As noted previously, this case covers the situation where the econometrician does not know the map between actions, states, and outcomes.

For the three-action and three-state version of this tracking problem, the map between actions, states, and outcomes can be represented as a matrix where actions a_1 – a_3 are given in the rows and states ω_1 – ω_3 are given in the columns:

$$\begin{array}{ccc} \omega_1 & \omega_2 & \omega_3 \\ \left(\begin{array}{ccc} x_{11} & x_{12} & x_{13} \\ x_{21} & x_{22} & x_{23} \\ x_{31} & x_{32} & x_{33} \end{array} \right) & \begin{array}{l} a_1 \\ a_2 \\ a_3 \end{array} \end{array}.$$

One example of this is given by

$$P_g = \begin{array}{ccc} \omega_1 & \omega_2 & \omega_3 \\ \left(\begin{array}{ccc} \frac{24}{72} & 0 & 0 \\ 0 & \frac{16}{72} & \frac{8}{72} \\ 0 & \frac{8}{72} & \frac{16}{72} \end{array} \right) & \begin{array}{l} a_1 \\ a_2 \\ a_3 \end{array} \end{array} \quad \text{and} \quad P_h = \begin{array}{ccc} \omega_1 & \omega_2 & \omega_3 \\ \left(\begin{array}{ccc} \frac{12}{72} & \frac{6}{72} & \frac{6}{72} \\ \frac{6}{72} & \frac{5}{72} & \frac{13}{72} \\ \frac{6}{72} & \frac{13}{72} & \frac{5}{72} \end{array} \right) & \begin{array}{l} a_1 \\ a_2 \\ a_3 \end{array} \end{array}.$$

Like the tracking example, these distributions of actions and states will reveal that the DM prefers the outcome obtained when choosing action a_1 when the state is ω_1 , prefers the outcomes obtained when choosing a_2 and a_3 in the other states, and is perfectly informed when taking action a_1 . Once again, it is as if for P_g the DM gets a signal realization that is perfectly informative of whether the state is ω_1 and so knows to take action a_1 if the state is ω_1 and not to choose action a_1 otherwise.

However, unlike the tracking example, this P_g and P_h reveal that the utility obtained from taking actions a_2 and a_3 is the same in every state.¹⁰ This follows from the fact that the NIAS inequalities for a_2 chosen over a_3 and a_3 chosen over a_2 hold with equality for both P_g and P_h , which is a consequence of lemma 1. Lemma 1 states that an NIAS inequality is equal to zero if and only if that NIAS inequality can be expressed as a nonpositive combination of other NIAS inequalities. For example, the NIAS inequality for a_2 chosen over a_3 for P_g is $(16/72)(u(a_2, \omega_2) - u(a_3, \omega_2)) + (8/72)(u(a_2, \omega_3) - u(a_3, \omega_3)) \geq 0$.¹¹ The negative of this can be obtained by simply adding together the outcome lotteries from the NIAS inequalities for a_3 chosen over a_2 for P_g , for a_2 chosen over a_3 for P_h , and for a_3 chosen over a_2 for P_h .

Given that the NIAS inequalities for a_2 chosen over a_3 and a_3 chosen over a_2 hold with equality for P_g , the utility differences between a_2 and a_3 in ω_2 and the utility differences between a_2 and a_3 in ω_3 are both equal to zero because those NIAS inequalities say

¹⁰ This example can also be generalized to any version of this problem with arbitrarily many actions and at least as many states as actions.

¹¹ Given that there are no common outcomes across states and actions in this decision problem, we will shorten $u(x(a, \omega))$ to $u(a, \omega)$.

$$\frac{16}{72}(u(a_2, \omega_2) - u(a_3, \omega_2)) + \frac{8}{72}(u(a_2, \omega_3) - u(a_3, \omega_3)) = 0$$

and

$$-\frac{8}{72}(u(a_2, \omega_2) - u(a_3, \omega_2)) - \frac{16}{72}(u(a_2, \omega_3) - u(a_3, \omega_3)) = 0,$$

which is possible only if $u(a_2, \omega_2) - u(a_3, \omega_2) = 0$ and $u(a_2, \omega_3) - u(a_3, \omega_3) = 0$. Likewise, given that the NIAS inequalities for a_2 chosen over a_3 and a_3 chosen over a_2 hold with equality for P_g , the utility difference between a_2 and a_3 in ω_1 is also equal to zero. Thus, the utility from taking a_2 is the same as the utility from taking a_3 in every state.

Given this, the value of information is higher for P_g if

$$\begin{aligned} & \frac{12}{72}(u(a_1, \omega_1) - u(a_2, \omega_1)) \\ & + \frac{6}{72}(u(a_2, \omega_2) - u(a_1, \omega_2) + u(a_2, \omega_3) - u(a_1, \omega_3)) \geq 0. \end{aligned}$$

To show that this holds, we first note that $u(a_1, \omega_1) \geq u(a_2, \omega_1)$ (the DM preferring to take action a_1 in state ω_1) follows directly from the NIAS inequality for a_1 chosen over a_2 for P_g . Second, because a_2 and a_3 give the same utility in every state, the NIAS inequalities for a_2 chosen over a_1 and a_3 chosen over a_1 for P_g yield

$$\frac{16}{72}(u(a_2, \omega_2) - u(a_1, \omega_2)) + \frac{8}{72}(u(a_2, \omega_3) - u(a_1, \omega_3)) \geq 0$$

and

$$\frac{8}{72}(u(a_2, \omega_2) - u(a_1, \omega_2)) + \frac{16}{72}(u(a_2, \omega_3) - u(a_1, \omega_3)) \geq 0.$$

Adding these together gives

$$u(a_2, \omega_2) - u(a_1, \omega_2) + u(a_2, \omega_3) - u(a_1, \omega_3) \geq 0.$$

With this, we have that P_g provides a higher value of information.

III. Characterizing the Relation

Is there a general approach to checking whether there are enough restrictions on u to ensure $P_g \succsim_V P_h$? We produce a necessary and sufficient condition for $P_g \succsim_V P_h$ by moving fully to the space of probabilities and probability differences over outcomes. There are three features that make this space important. First, it allows geometric representation of the NIAS inequalities. Second, it allows identification of all utility functions that

satisfy these inequalities. Third, it identifies all differences in outcome lotteries that are guaranteed to raise utility (which allows us to identify $\succeq_v$). Because of this geometric representation, we can reduce $P_g \succeq_v P_h$ to a single system of linear equations.

A. Ruling Out Utility Functions Geometrically

First, each NIAS inequality can be represented as an M -dimensional vector $\vec{d}_f(a, b)$ that gives the outcome lottery gained from not making a wholesale switch from action a to action b for P_f —in other words, the additional probability of receiving each outcome from not making this wholesale switch. Element m of this vector gives the additional probability of receiving outcome x_m in X from not making a wholesale switch from action a to action b for P_g , which is

$$\sum_{\omega \in \Omega} P_f(a, \omega) (\mathbf{1}_{\{x(a, \omega) = x_m\}} - \mathbf{1}_{\{x(b, \omega) = x_m\}}),$$

where $\mathbf{1}_{\{x(a, \omega) = x_m\}}$ is an indicator function that takes a value of one when the outcome from taking action a in state ω yields outcome x_m . The convex cone D formed by all NIAS inequalities is

$$D = \{\alpha_1 \vec{d}_{f_1}(a_1, b_1) + \cdots + \alpha_N \vec{d}_{f_N}(a_N, b_N) \mid \alpha_n \in \mathbb{R}_+, f_n \in \{g, h\}, a_n, b_n \in A\}.$$

A utility function can be represented as an M -dimensional vector $\vec{u}$, where element m gives the utility of outcome x_m . A utility vector $\vec{u}$ satisfies NIAS if $\vec{d} \bullet \vec{u} \geq 0$ for every vector $\vec{d} \in D$. We call the convex cone formed by all $\vec{u}$ that satisfy NIAS the *NIAS utility cone*.

We illustrate this with a simple decision problem that has a safe action, where the map between states, actions, and outcomes is given by

$$\begin{array}{cc} \omega_1 & \omega_2 \\ \left(\begin{array}{cc} x_1 & x_2 \\ x_3 & x_3 \end{array} \right) & \begin{array}{l} a \\ b \end{array} \end{array}$$

Imagine the following distribution of actions and states:

$$P_g = \begin{array}{cc} \omega_1 & \omega_2 \\ \left(\begin{array}{cc} 0.4 & 0.1 \\ 0.1 & 0.4 \end{array} \right) & \begin{array}{l} a \\ b \end{array} \end{array} \quad \text{and} \quad P_h = \begin{array}{cc} \omega_1 & \omega_2 \\ \left(\begin{array}{cc} 0.15 & 0.05 \\ 0.35 & 0.45 \end{array} \right) & \begin{array}{l} a \\ b \end{array} \end{array}.$$

Choosing a in P_g gets x_1 and x_2 with unconditional probabilities 0.4 and 0.1. Choosing b gets x_3 . Hence, sticking with a over b yields

$$\vec{d}_g(a, b) = (0.4, 0.1, -0.5).$$

Likewise, sticking with b over a in P_g gets x_3 rather than x_1 with unconditional probability 0.1 and x_2 with unconditional probability 0.4:

$$\vec{d}_g(b, a) = (-0.1, -0.4, 0.5).$$

Analogously,

$$\vec{d}_h(a, b) = (0.15, 0.05, -0.2),$$

$$\vec{d}_h(b, a) = (-0.35, -0.45, 0.8).$$

NIAS identifies rationalizing utility functions as all that have (weakly) positive dot products with all of these vectors. This can be visualized in two dimensions by normalizing $u(x_3) = 0$. With this normalization, the (x_1, x_2) space can illustrate both D and the NIAS utility cone, which is given in figure 1.

B. Ranking Decisions Geometrically

Let $\vec{d}(g, h)$ be an M -dimensional vector that gives the outcome lottery gained from encountering P_g instead of P_h —in other words, the additional probability of receiving each outcome from P_g . For outcome x_m in X , this is

$$\sum_{a \in A} \sum_{\omega \in \Omega} (P_g(a, \omega) - P_h(a, \omega)) \mathbf{1}_{\{x(a, \omega) = x\}}.$$




FIG. 1.—Illustration of the geometric structure of ruling out utility functions using NIAS.

The joint distribution P_g being revealed to have a higher value of information ($P_g \succsim_v P_h$) is equivalent to $\vec{d}(g, h) \cdot \vec{u} \geq 0$ for all $\vec{u}$ in the NIAS utility cone. Thus, $P_g \succsim_v P_h$ if and only if the vector $\vec{d}(g, h)$ is in D because a vector is in D if and only if it has a nonnegative dot product with all $\vec{u}$ in the NIAS utility cone.

Finally, because $\vec{d}(g, h)$ is in D if and only if it is a nonnegative weighted average of vectors in D , a necessary and sufficient condition for $P_g \succsim_v P_h$ corresponds to the outcome lottery gained from P_g being a nonnegative weighted average of the outcome lotteries gained from not making wholesale switches from any action for either P_f or P_g .

Returning to the example, P_g yields $(0.4, 0.1, 0.5)$, and P_h yields $(0.15, 0.05, 0.8)$. Hence, $\vec{d}(g, h) = (0.25, 0.05, -0.3)$. As illustrated in figure 2, $\vec{d}(g, h)$ is in D , so it has a positive dot product with the entire NIAS utility cone, so P_g and P_h are ranked by $\succsim_v$.

C. General Condition for Ranking Decisions

The condition that $\vec{d}(g, h)$ is in D , which we call *decision improvement without action switches* (DISI) for P_g over P_h , is defined for a weighting function $t_{gh} : A \times A \rightarrow \mathbb{R}_+$, which provides these nonnegative weights.

CONDITION 1 (DISI). Weighting function $t_{gh} : A \times A \rightarrow \mathbb{R}_+$ satisfies DIAS for P_g over P_h if for every $x \in X$,




FIG. 2.—Illustration of the geometric structure of ranking decisions.

$$\begin{aligned}
& \sum_{P_f \in \{P_g, P_h\}} \sum_{a \in A} \sum_{b \in A} \sum_{\omega \in \Omega} P_f(a, \omega) (\mathbf{1}_{\{x(a, \omega) = x\}} - \mathbf{1}_{\{x(b, \omega) = x\}}) t_{gh}(a, b) \\
&= \sum_{a \in A} \sum_{\omega \in \Omega} (P_g(a, \omega) - P_h(a, \omega)) \mathbf{1}_{\{x(a, \omega) = x\}}.
\end{aligned}$$

The following theorem formally shows that DISI provides a necessary and sufficient condition to reveal that every experiment consistent with one distribution of actions and states has a higher value of information than every experiment consistent with one distribution of actions and states. After restating NIAS and DISI in terms of matrix multiplication, the proof of this theorem follows as a direct consequence of Farkas's lemma (Farkas 1902).

THEOREM 1. $P_g \succsim_V P_h$ if and only if there exists a weighting function t_{gh} that satisfies DISI for P_g over P_h .

Proof. See the appendix. QED

This theorem has an economic interpretation in terms of preferences over outcome lotteries. DISI states that the difference in the outcome lotteries offered by the distributions of actions and states can be represented as the difference in two compound lotteries: one composed of the outcome lotteries from taking each action a for P_g and the other composed of outcome lotteries from taking each action b with the same probability as a . These compound lotteries have the same weights for all $P_f \in \{P_g, P_h\}$ and $a, b \in A$, which are given by a normalized version of t_{gh} . Because NIAS is satisfied, there exists a preference relation over lotteries such that every element of one compound lottery is weakly preferred to every element in the other compound lottery. Since all elements of the two compound lotteries are preference ordered, the compound lotteries are also preference ordered, which means that the outcome lotteries given by each distribution of actions and states are as well.

As noted previously, this theorem applies even when the econometrician does not know the outcomes to taking actions in each state. In this case, the econometrician can arbitrarily assign a distinct outcome to every action in every state. If DISI is satisfied given the unrestricted outcome mapping, then $P_g \succsim_V P_h$ also holds for all other outcome mappings that are consistent with at least one utility function that satisfies NIAS.

D. Testability

Determining whether there exists a t_f that satisfies DISI corresponds to determining whether there is a solution to a system of linear equations, so it is simple to check whether distributions of actions and states are welfare ranked. We provide MATLAB computer programs that determine whether a solution to this linear system exists for a given set of data.¹²

¹² Programs are available at <https://github.com/danieljosephmartin>.

Also, there are settings where some options are clearly dominant, and NIAS and DISI can be easily amended to account for these additional restrictions. Say, for example, that outcome x_1 clearly dominates outcome x_2 . This restriction on utility can be incorporated into NIAS by generating an additional linear inequality given by

$$\sum_{x \in X} (\mathbf{1}_{x=x_1} - \mathbf{1}_{x=x_2}) u(x) \geq 0.$$

Clearly, this restriction on the set of admissible utility functions can only reduce the set of u that satisfy NIAS.

Although DISI is not expressed in terms of utility, the dominance of outcome x_1 over outcome x_2 can be incorporated into DISI for P_g by requiring that, in addition to the weighting function t_{gh} , there exists a non-negative t that solves

$$\begin{aligned} & \sum_{P_f \in \{P_g, P_h\}} \sum_{a \in A} \sum_{b \in A} \sum_{\omega \in \Omega} P_f(a, \omega) (\mathbf{1}_{\{x(a, \omega)=x\}} - \mathbf{1}_{\{x(b, \omega)=x\}}) t_{gh}(a, b) \\ & + (\mathbf{1}_{x=x_1} - \mathbf{1}_{x=x_2}) t \\ & = \sum_{a \in A} \sum_{\omega \in \Omega} (P_g(a, \omega) - P_h(a, \omega)) \mathbf{1}_{\{x(a, \omega)=x\}} \end{aligned}$$

for every $x \in X$. If t is equal to zero, this reduces to the requirement for DISI, so this addition can only increase the proportion of P_g and P_h where there exists a weighting function t_{gh} that satisfies DISI. Sensibly, knowledge about dominance improves our ability to rank decisions according to their welfare.

IV. Related Literature

Our work is most closely related to three other papers. First, our relation draws natural parallels to the seminal informativeness relation provided by Blackwell (1953). In comparing experiments, Blackwell holds fixed π_g and π_h and for each u evaluates the expected utility provided by all P_g and P_h consistent with those experiments. On the other hand, in comparing decisions, we hold fixed P_g and P_h and for each u evaluate the expected utility provided by all π_g and π_h consistent with those decisions. This change in perspective produces an important technical difference that is illustrated in the preceding sections. For Blackwell's relation, every u must be considered because every experiment π_g has a consistent P_g for every u . However, every u does not need to be considered for our relation because P_g may not have a consistent π_g for some u .

Second, like our paper, Lu (2016) also orders unknown experiments, but his characterization requires much more than simply the joint distribution of actions and states P_g and P_h . He also requires that the econometrician

observes richer “test functions” F_g and F_h , which are not naturally occurring. A test function F indicates how often the actions are chosen when the set of actions is paired with every possible mixture between the best and worst action.

Third, our paper builds on the results of Caplin and Martin (2015) and the NIAS condition they introduce. As a result, our paper is also related to the work of Bergemann and Morris (2016), as their obedience condition is identical to the NIAS condition in a single-player setting with no initial signals. However, an important distinction is that Bergemann and Morris (2016) take the utility function as known and use the obedience condition to determine the set of joint distribution of actions and states that are consistent with Bayes correlated equilibrium, whereas Caplin and Martin (2015) use NIAS to determine the set of utility functions that are consistent with a joint distribution of actions and states.

We provide three innovations relative to Caplin and Martin (2015). First, we provide a novel geometric representation for using NIAS to identify consistent utility functions. Second, we provide a new result showing when NIAS holds strictly and weakly, which enhances the analytical and computational tractability of NIAS. Third, and most importantly, we provide an entirely new application of NIAS by showing exactly when it can rule out enough utility functions to allow the value of information to be ranked between two unknown experiments.

Appendix

A1. Proof of Lemma 1

First, for any u that satisfies NIAS, by definition

$$\sum_{x \in X} \left(\sum_{\omega \in \Omega} P_f(a, \omega) (\mathbf{1}_{\{x(a, \omega) = x\}} - \mathbf{1}_{\{x(b, \omega) = x\}}) \right) u(x) \geq 0 \quad (\text{A1})$$

for any triple (P_n, a_n, b_n) , where $P_n \in \{P_f, P_g\}$, $a_n \in A$, and $b_n \in A$. Thus, for any u that satisfies NIAS,

$$-1 \times \sum_{x \in X} \left(\sum_{n=1}^N w_n \left(\sum_{\omega \in \Omega} P_n(a_n, \omega) (\mathbf{1}_{\{x(a_n, \omega) = x\}} - \mathbf{1}_{\{x(b_n, \omega) = x\}}) \right) \right) u(x) \geq 0$$

for any collection of triples (P_n, a_n, b_n) , where $P_1, \dots, P_N \in \{P_f, P_g\}$, $a_1, \dots, a_N \in A$, and $b_1, \dots, b_N \in A$ with $(P_n, a_n, b_n) \neq (P_f, a, b)$ and nonpositive weights $w_1, \dots, w_N$. Assuming that equation (2) holds, this implies that for any u that satisfies NIAS,

$$-1 \times \sum_{x \in X} \left(\sum_{\omega \in \Omega} P_f(a, \omega) (\mathbf{1}_{\{x(a, \omega) = x\}} - \mathbf{1}_{\{x(b, \omega) = x\}}) \right) u(x) \geq 0.$$

Because of equation (A1), this must equal zero, so equation (1) must hold.

Second, if equation (1) holds, then for all u that satisfy NIAS, it cannot be that

$$-1 \times \sum_{x \in X} \left(\sum_{\omega \in \Omega} P_f(a, \omega) (\mathbf{1}_{\{x(a, \omega) = x\}} - \mathbf{1}_{\{x(b, \omega) = x\}}) \right) u(x) < 0. \quad (\text{A2})$$

By Farkas's lemma, equations (A1) and (A2) mean that there must exist non-positive weights on that collection of NIAS inequalities that give equation (2), completing the proof. QED

A2. Proof of Theorem 1

The NIAS inequality for $P_f \in \{P_g, P_h\}$ and actions $a, b \in A$ can be expressed as a $1 \times M$ -row vector $\vec{d}_f(a, b)$, where element m of this vector gives the difference in the probability of receiving outcome x_m from taking action a and from taking action b with the same probability. This is given by

$$\sum_{\omega \in \Omega} P_h(a, \omega) (\mathbf{1}_{\{x(a, \omega) = x_m\}} - \mathbf{1}_{\{x(b, \omega) = x_m\}}),$$

where $\mathbf{1}_{\{x(a, \omega) = x_m\}}$ is an indicator function that takes a value of one when the outcome from taking action a in state ω yields outcome x_m .

Stacking the row vectors for all NIAS inequalities for $P_f \in \{P_g, P_h\}$ produces a $J^2 \times M$ matrix D_h , where

$$D_h = \begin{bmatrix} \vec{d}_h(a_1, a_1) \\ \vec{d}_h(a_1, a_2) \\ \vdots \\ \vec{d}_h(a_j, a_{j-1}) \\ \vec{d}_h(a_j, a_j) \end{bmatrix},$$

and stacking the matrix of NIAS inequalities for both distributions of actions and states produces a $2 \times J^2 \times M$ matrix D , where

$$D = \begin{bmatrix} D_f \\ D_g \end{bmatrix}.$$

Based on this matrix D , NIAS can be restated as the $M \times 1$ -column vector $u \in \mathbb{R}^M$ satisfying $Du \geq 0$, with $Du(m) > 0$ for some $m \in \{1, \dots, M\}$.

In addition, the requirement for P_g to be revealed to have a higher value of information than P_h can be expressed as a $1 \times M$ -row vector $\vec{d}$, where element m gives the expected gain in outcome x_m from choosing with P_g instead of P_h , which is given by

$$\sum_{a \in A} \sum_{\omega \in \Omega} (P_f(a, \omega) - P_g(a, \omega)) \mathbf{1}_{\{x(a, \omega) = x_m\}}.$$

Here $P_g \succsim_V P_h$ can be restated as $\vec{d}u \geq 0$ for all $u \in \mathbb{R}^M$ that satisfy NIAS. With this notation, both directions of the theorem follow from Farkas's lemma.

1. There exists $t \in \mathbb{R}_+^{2 \times J^*}$ s.t. $D^T t = (\vec{d})^T \Rightarrow$. For all $u \in \mathbb{R}^M$ satisfying NIAS, $\vec{d}u \geq 0$. Assume not. Take $u \in \mathbb{R}^M$ such that NIAS is satisfied, so that $Du \geq 0$, but $\vec{d}u < 0$. By Farkas's lemma, there cannot exist a $t \in \mathbb{R}_+^{2 \times J^*}$ s.t. $D^T t = (\vec{d})^T$, which is a contradiction.
2. For all $u \in \mathbb{R}^M$ satisfying NIAS, $\vec{d}u \geq 0 \Rightarrow$. There exists $t \in \mathbb{R}_+^{2 \times J^*}$ s.t. $D^T t = (\vec{d})^T$. Assume there does not exist $t \in \mathbb{R}_+^{2 \times J^*}$ such that $D^T t = (\vec{d})^T$. By Farkas's lemma, there must exist a $u \in \mathbb{R}^M$ satisfying NIAS and with $\vec{d}u < 0$, which is a contradiction. QED

References

- Abeler, Johannes, and Simon Jäger. 2015. "Complex Tax Incentives." *American Econ. J.: Econ. Policy* 7 (3): 1–28.
- Apestegui, Jose, and Miguel A. Ballester. 2018. "Monotone Stochastic Choice Models: The Case of Risk and Time Preferences." *J.P.E.* 126 (1): 74–106.
- Bergemann, Dirk, and Stephen Morris. 2016. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Econ.* 11 (2): 487–522.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor. 2017. "Choose to Lose: Health Plan Choices from a Menu with Dominated Option." *Q.J.E.* 132 (3): 1319–72.
- Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *Ann. Math. Statist.* 24 (2): 265–72.
- Burszty, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2020. "Misinformation during a Pandemic." Working Paper no. 2020-44, Becker Friedman Inst. Econ., Univ. Chicago.
- Caplin, Andrew, and Daniel Martin. 2015. "A Testable Theory of Imperfect Perception." *Econ. J.* 125 (582): 184–202.
- Carpenter, Jeffrey, Emiliano Huet-Vaughn, Peter Hans Matthews, Andrea Robbett, Dustin Beckett, and Julian Jamison. 2021. "Choice Architecture to Improve Financial Decision Making." *Rev. Econ. and Statist.* 103 (1): 102–18.
- Carrera, Mariana, and Sofia Villas-Boas. 2015. "Generic Aversion and Observational Learning in the Over-the-Counter Drug Market." Working paper, Dept. Agricultural and Resource Econ., Univ. California, Berkeley.
- Choi, James J., David Laibson, and Brigitte C. Madrian. 2009. "Reducing the Complexity Costs of 401(k) Participation through Quick Enrollment." In *Developments in the Economics of Aging*, edited by David A. Wise, 57–82. Chicago: Univ. Chicago Press (for NBER).
- Clippel, Geoffroy de, and Kareen Rozen. 2020. "Communication, Perception, and Strategic Obfuscation." Working paper, Dept. Econ., Brown Univ., Providence, RI.
- Dranove, David, and Ginger Zhe Jin. 2010. "Quality Disclosure and Certification: Theory and Practice." *J. Econ. Literature* 48 (4): 935–63.
- Ericson, Keith M. Marzilli, and Amanda Starc. 2016. "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange." *J. Health Econ.* 50:71–85.
- Esponda, Ignacio, and Emanuel Vespa. 2019. "Contingent Preferences and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory." Working paper, Dept. Econ., Univ. California, Santa Barbara.

- Farkas, Julius. 1902. "Über die Theorie der Einfachen Ungeichungen." *J. Reine Angewandte Math.* 124:1–24.
- Frankel, Alexander, and Emir Kamenica. 2018. "Quantifying Information and Uncertainty." *A.E.R.* 109 (10): 3650–80.
- Hastings, Justine S., and Lydia Tejada-Ashton. 2008. "Financial Literacy, Information, and Demand Elasticity: Survey and Experimental Evidence from Mexico." Working Paper no. 14538, NBER, Cambridge, MA.
- Jin, Ginger Zhe, Michael Luca, and Daniel Martin. 2015. "Is No News (Perceived as) Bad News? An Experimental Investigation of Information Disclosure." Working Paper no. 21099, NBER, Cambridge, MA.
- . 2018. "Complex Disclosure." Working Paper no. 24675, NBER, Cambridge, MA.
- Kamenica, Emir. 2019. "Bayesian Persuasion and Information Design." *Ann. Rev. Econ.* 11:249–72.
- Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *A.E.R.* 101 (6): 2590–615.
- Lara, Michel De, and Olivier Gossner. 2020. "Payoffs-Beliefs Duality and the Value of Information." *SIAM J. Optimization* 30 (1): 464–89.
- Lu, Jay. 2016. "Random Choice and Private Information." *Econometrica* 84 (6): 1983–2027.
- Manzini, Paola, and Marco Mariotti. 2014. "Stochastic Choice and Consideration Sets." *Econometrica* 82 (3): 1153–76.
- McFadden, Daniel. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics*, edited by Paul Zarembka, 105–42. New York: Academic Press.
- Milgrom, Paul R. 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell J. Econ.* 12 (2): 380–91.