SSE Working Paper Series in Economics
No. 2019:2

FEASIBLE BEST-RESPONSE CORRESPONDENCES AND QUADRATIC SCORING RULES

Henk Norde [a]
Mark Voorneveld [b]

[a] Tilburg University
Center and Department of Econometrics and Operations Research

[b] Stockholm School of Economics
Department of Economics

Feasible best-response correspondences and quadratic scoring rules*

HENK NORDE

CentER and Department of Econometrics and Operations Research, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, the Netherlands, h.norde@uvt.nl

MARK VOORNEVELD

Department of Economics, Stockholm School of Economics, Box 6501, 113 83 Stockholm, Sweden, mark.voorneveld@hhs.se

May 7, 2019

SSE WORKING PAPER SERIES IN ECONOMICS 2019:2

Abstract

The rational choice paradigm in game theory and other fields of economics has agents best-responding to beliefs about factors that are outside their control. And making certain options a best response is a common problem in mechanism design and information elicitation. But not every correspondence can be made into a best-response correspondence. So what characterizes a feasible best-response correspondence? And once we know that, can we find some or even all utility functions that give rise to this best-response correspondence? We answer these three questions for an expected-utility maximizing agent with finitely many actions and probabilistic beliefs over finitely many states or opponents' strategies. We apply our results to information elicitation problems where contracts (scoring rules) are designed to financially reward an expected-payoff maximizing agent to truthfully reveal a property of her belief by sending a report from some finite set of messages. This leads to a number of new insights: firstly, we characterize exactly which properties can be elicited using scoring rules; secondly, we show that in this class of problems quadratic scoring rules are both necessary and sufficient methods of doing so.

Keywords: best-response correspondence, best-response equivalence, information elicitation, scoring rule

JEL codes: C72, D82, D83

*This paper has been a long time in the making; earlier versions were presented at SING2 (Spain, Italy, Netherlands conference on Game theory) in Foggia, the Conference on Mechanism and Institution Design in Durham, and seminars in Maastricht and Stockholm. We thank their audiences as well as Maria Kozlovskaya, Indrajit Ray, Lloyd Shapley, Elias Tsakas, Dries Vermeulen, and Günter Ziegler for helpful comments and discussions. Mark Voorneveld gratefully acknowledges financial support by the Wallander-Hedelius Foundation under grant P2016-0072:1.

1 Introduction

It is little wonder that economists are interested in best-response correspondences:

- *Homo economicus*, poster child of the rationality paradigm, is often defined as someone choosing/responding optimally given beliefs about factors outside its control.
- A variety of game-theoretic concepts (including Nash equilibria, rationalizable strategies, and adjustment processes like fictitious play) in mixed extensions of finite strategic games depend only on the game's best-response correspondences.
- Making certain options a best response is a common reverse-engineering problem in areas like mechanism design and information elicitation.

But not every correspondence can be made into a best-response correspondence of a well-defined game. So, (Q1) can we characterize whether a putative best-response correspondence is feasible? And once we know that, can we find (Q2) some utility function or (Q3) all utility functions with this best-response correspondence?

We answer these three questions in the finite setting, i.e., for an expected-utility maximizing agent with finitely many actions and probabilistic beliefs over finitely many states or opponents' strategies. Although this covers a variety of settings (see Remark 2.2 below), for the sake of concreteness we will mostly talk about the best-response correspondence of the row player in a bimatrix game. As another application, our theorems provide novel insights into information elicitation problems where contracts (scoring rules) are designed to financially reward an expected-payoff maximizing agent to truthfully reveal a property of her belief by sending a report from some finite set of messages. Let us briefly discuss our contributions.

Characterizing best-response correspondences. We provide two characterization results. An alleged best-response correspondence can be summarized by writing down, for each of the (say) m pure strategies of the row player, against which strategies of the column player they need to be a best response. This cuts the column player's mixed strategy space into m best-response sets $P_1, \ldots, P_m$.

Our first characterization provides a link to the relative-proximity literature on spatial competition in the spirit of Hotelling, where — all else being equal — businesses can sometimes compete by strategically locating themselves at a site that attracts customers simply because that site happens to be nearest: Theorem 3.1 says that feasible best-response sets coincide with nearest-neighbor regions for suitably chosen sites.

This is illustrated for a simple 3×3 game in Figure 1. The column player's mixed strategy space is the unit simplex with the standard basis vectors e_1, e_2, and e_3, where some column is chosen with probability one, as its vertices. Pick three sites v_1, v_2, and v_3 in $\mathbb{R}^3$. For simplicity, in the figure we take the three vertices: $v_i = e_i$ for $i = 1, 2, 3$. Using the usual Euclidean distance, let P_i (with $i = 1, 2, 3$) be the subset of the simplex for which v_i is the nearest of those sites. Theorem 3.1 assures that there is a payoff matrix (like matrix A in the figure) for the row player where these sets P_i are the beliefs against which row i is a best response; and conversely, every best-response correspondence of a 3×3 payoff matrix can be found in this way.

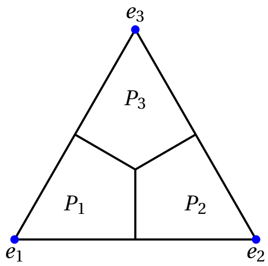

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Figure 1: Best responses and nearest neighbors in a simple 3×3 payoff matrix.

The second characterization, Theorem 4.1, identifies feasible best-response sets as solutions to a system of finitely many linear (in)equalities and provides a linear programming problem with optimal value zero if there is no solution to that system and optimal value one if there is.

These two characterizations are instrumental in answering our follow-up questions.

Finding some game with given best responses: quadratic scoring rules. The proof of the nearest-neighbor characterization (Theorem 3.1) not only tells us that feasible best-response sets are nearest-neighbor regions for some suitable sites; it also tells us how to transform those sites into a payoff matrix with the desired best responses. This transformation is of a special functional form (see (6)), known in the literature on information elicitation as a quadratic scoring rule. So, given feasible best-response sets, there is a corresponding payoff matrix that is a quadratic scoring rule (Corollary 3.2). Moreover, each payoff matrix with these best responses can be written as a quadratic scoring rule (Theorem 3.3).

Application to finite elicitation problems. In Section 2.2 and Remark 3.4 we discuss in detail an application of our results to finite elicitation problems where an expected-payoff maximizing agent needs to be incentivized to accurately reveal a property of her beliefs about the state of nature by sending one of finitely many messages. Scoring rules are contracts that pay the agent depending on the realized state and the message that she sent. Our three results in Section 3 lead to a number of new insights: firstly, by characterizing exactly which properties can be elicited using scoring rules; secondly, by showing that quadratic scoring rules are both necessary and sufficient methods of doing so. Especially the latter is unusual: quadratic scoring rules are known to perform well in a variety of settings, but often along with many other scoring rules — i.e., in those settings they are sufficient, but not necessary.

Finding all games with given best responses. The second characterization result, Theorem 4.1, says that finding all games with given best responses boils down to solving a system of homogeneous linear equalities and strict inequalities. These strict inequalities complicate the analysis: the set of solutions is a convex cone, but typically not closed and therefore not finitely generated. The appendix contains a technical Lemma A.1 that describes the set of solutions to such (in)equalities. In game-theoretic terms (Theorem 4.2), this assures that the set of all payoff matrices with the same best responses as some given matrix A can be written in terms of a dummy game (where the row player is indifferent between all her actions) and other games with a simpler (coarser) best-response structure than A.

A discussion of and precise references to related literature will be given in the main text; all proofs are in the appendix.

2 Problem statement and applications

2.1 Problem statement

In $\mathbb{R}^n$, vectors a and b have their usual inner product $\langle a, b\rangle = \sum_{i=1}^n a_i b_i$ and Euclidean distance $\|a-b\| = \sqrt{\langle a-b, a-b\rangle}$. Denote the unit simplex in $\mathbb{R}^n$ by

$$\Delta_n = \{y \in \mathbb{R}^n : y_1, \dots, y_n \geq 0, y_1 + \cdots + y_n = 1\}.$$

Consider an $m \times n$ matrix A, interpreted as the row player's von Neumann-Morgenstern utility function or, briefly, payoff function in a bimatrix game. Denote its rows by $a_1, \dots, a_m$ in $\mathbb{R}^n$. For each row $i = 1, \dots, m$, let ***best-response set***

$$P_i = \{y \in \Delta_n : \langle a_i, y\rangle \geq \langle a_k, y\rangle \text{ for all } k = 1, \dots, m\} \tag{1}$$

be the set of mixed strategies of player 2 against which pure strategy i is a best response. The next proposition summarizes some properties of best-response sets.

Proposition 2.1. *The best-response sets have the following properties:*

(a) Each best-response set is a polytope.

(b) Their union is Δ_n.

(c) The intersection of two or more best-response sets is either empty or a face of each of those sets.

We are concerned with the opposite direction: given putative best responses, do they really come from a well-defined game? Formally, the questions we set out to answer are: given a collection of sets $P_1, \dots, P_m$ whose union is Δ_n,

(Q1) are these ***feasible best-response sets***, i.e., does there exist a payoff matrix A that has these sets as best-response sets?

And if so, can we construct

(Q2) *some* payoff matrix or

(Q3) *all* payoff matrices with the desired best-response sets?

Remark 2.2. We stress that our setting of characterizing best-response correspondences in bimatrix games is more general than may at first appear. For instance, it includes mixed extensions of finite strategic games with more than two players if correlated beliefs are allowed. It also encompasses settings of decision-making under uncertainty, where the second player is replaced by nature. An important other application arises in information elicitation problems, which we discuss next. $\lhd$

2.2 Application: information elicitation using scoring rules

Consider a set $J = \{1,\dots,n\}$ of $n \in \mathbb{N}$, $n \geq 2$, exhaustive and mutually exclusive states. An agent's subjective beliefs about the probability of each state are modeled as usual by probability vectors in the unit simplex Δ_n. The agent can send a principal information about her belief by choosing a message from a set I. Some map, often called a ***property***, assigns to each belief $y \in \Delta_n$ a nonempty set $\text{PROP}(y) \subseteq I$ of messages considered the correct one(s) given those beliefs. Typical properties in the statistics literature on elicitation include the belief itself, moments of the probability distribution, or its mode. More prosaic, economic examples that come to mind are answers to questions like 'Given your assessment of tomorrow's weather, how many ice-cream cones should we buy?': pretty much any question related to forecasting or prediction.

A ***scoring rule*** is a contract mapping the agent's message and the realized state to a payoff, i.e., a function $S : I \times J \to \mathbb{R}$: after the agent reports some information $i \in I$, both the principal and the agent observe the true state $j \in J$ and the agent receives payoff $S(i, j) \in \mathbb{R}$.

The principal aims to find a scoring rule that makes it optimal for a payoff-maximizing agent to reveal the correct information. To facilitate comparison with our earlier notation, define for each message $i \in I$ the vector $a_i \in \mathbb{R}^n$ with coordinates $a_{i,j} = S(i, j)$ for each state $j \in J = \{1,\dots,n\}$. Given any belief $y \in \Delta_n$, the agent's expected payoff from sending information $i \in I$ to the principal is then

$$\sum_{j=1}^{n} y_j S(i, j) = \langle a_i, y\rangle.$$

A scoring rule S ***elicits*** the property PROP if the agent maximizes her expected payoff, given any belief y, if and only if she sends correct information, i.e., a message i in $\text{PROP}(y)$:

$$\text{for each } i \in I: \qquad \{y \in \Delta_n : i \in \text{PROP}(y)\} = \{y \in \Delta_n : \langle a_i, y\rangle \geq \langle a_k, y\rangle \text{ for all } k \in I\}. \tag{2}$$

Call the elicitation problem ***finite*** if both J (as presumed before) and I are finite; in that case we may enumerate I's elements by $I = \{1,\dots,m\}$ for some $m \in \mathbb{N}$ and (2) becomes

$$\text{for each } i \in I: \qquad \{y \in \Delta_n : i \in \text{PROP}(y)\} = \{y \in \Delta_n : \langle a_i, y\rangle \geq \langle a_k, y\rangle \text{ for all } k = 1,\dots,m\}.$$

On the right side we recognize the best-response sets (1) of a payoff matrix A with rows $a_1,\dots,a_m$. So for finite elicitation problems the question

> 'Is there a scoring rule that elicits our property?'

is equivalent with the question

> 'Are the sets $\{y \in \Delta_n : i \in \text{PROP}(y)\}$, with $i = 1,\dots,m$, feasible best-response sets?'

This makes finite elicitation problems a special case of one of the general questions we set out to answer.

The literature on scoring rules is too large to do serious justice to here; overview articles on different aspects include Savage (1971), Gneiting and Raftery (2007), Schotter and Trevino (2014), Schlag et al. (2015), and Carvalho (2016). A scoring rule that performs well in the elicitation of probabilities of specific events or of the whole probability distribution is the quadratic scoring rule. It dates back to a note on incentivizing meteorologists for accurate forecasts by Brier (1950), usually credited to be the starting point of this literature. This scoring rule plays an important role in our results too, so let us define it here.

Scoring rule $S : I \times J \to \mathbb{R}$ is a ***quadratic scoring rule*** if there are scalars α and β, with $\alpha > 0$, and a vector $v_i \in \mathbb{R}^n$ for each information report $i \in I$ such that the vector of scores associated with report i and the n distinct states is of the form

$$a_i = (S(i,1), \dots, S(i,n)) = \alpha \left(2v_i - \| v_i \|^2 \mathbf{1} \right) + \beta \mathbf{1}, \tag{3}$$

with $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^n$ the vector of ones.

After meteorology, the design and use of scoring rules caught on in statistics and many branches of the social sciences, including economics (accounting, management, finance), but also education, psychology, medicine, political and computer science; see the overviews cited earlier or Offerman et al. (2009, p. 1462). Applications in economics encompass, among others, incentive schemes in organizations (Thomson, 1979; Osband, 1985), information aggregation in markets (Ostrovsky, 2012), and strategic distinguishability (Bergemann et al., 2017).

3 Nearest neighbors and quadratic scoring rules

A crucial insight in the classical paper of Hotelling (1929) is that, all else being equal, businesses can sometimes compete by strategically locating themselves at a site that attracts customers simply because that site happens to be nearest. A similar point applies to politicians profiling themselves in an ideological space if voters choose candidates closest to their own ideals. Grofman (2004) reviews the literature on this proximity theory of voting. Our first characterization provides a link with this relative-proximity literature: feasible best-response sets are exactly the ones that cut the unit simplex into pieces corresponding to elements having suitably chosen sites as their nearest neighbor. Formally, consider m vectors or 'sites' $v_1, \dots, v_m$ in $\mathbb{R}^n$. Divide the unit simplex Δ_n into pieces, depending on which site happens to be nearest. The ***nearest-neighbor region*** of site v_i, denoted $\text{Near}(v_i)$, consists of those elements of the unit simplex that have v_i as nearest site:

$$\text{Near}(v_i) = \{ y \in \Delta_n : \| y - v_i \| \le \| y - v_k \| \text{ for all } k = 1, \dots, m \}.$$

In computational geometry, nearest-neighbor regions are sometimes called Voronoi regions. Aurenhammer et al. (2013, Ch. 4) and Borgwardt and Frongillo (2019) discuss algorithms to detect them.

Theorem 3.1 (Nearest-neighbor characterization)**.** *Polytopes $P_1, \dots, P_m$ with union Δ_n are feasible best-response sets if and only if they are nearest-neighbor regions:*

$$\textit{there are } v_1, \dots, v_m \in \mathbb{R}^n \textit{ with } P_i = \text{Near}(v_i) \textit{ for all } i = 1, \dots, m.$$

What drives the result is the following: given vectors y, v_k, and v_ℓ in $\mathbb{R}^n$, we have

$$\begin{aligned}\|y - v_k\| \leq \|y - v_\ell\| &\iff \langle y - v_k, y - v_k\rangle \leq \langle y - v_\ell, y - v_\ell\rangle \\ &\iff \|y\|^2 - 2\langle v_k, y\rangle + \|v_k\|^2 \leq \|y\|^2 - 2\langle v_\ell, y\rangle + \|v_\ell\|^2 \\ &\iff 2\langle v_k, y\rangle - \|v_k\|^2 \geq 2\langle v_\ell, y\rangle - \|v_\ell\|^2. \end{aligned} \tag{4}$$

Recall that $\mathbf{1}$ is the vector of ones. If r is a real number and y lies in the unit simplex, then $r = \langle r\mathbf{1}, y\rangle$, since the coordinates of y sum to one. With $\|v_k\|^2$ and $\|v_\ell\|^2$ instead of r, we can then rewrite (4):

$$\|y - v_k\| \leq \|y - v_\ell\| \iff \left\langle 2v_k - \|v_k\|^2\mathbf{1}, y\right\rangle \geq \left\langle 2v_\ell - \|v_\ell\|^2\mathbf{1}, y\right\rangle. \tag{5}$$

Hence, if we define $a_k = 2v_k - \|v_k\|^2\mathbf{1}$ and $a_\ell = 2v_\ell - \|v_\ell\|^2\mathbf{1}$, we see that y lies closer to v_k than to v_ℓ if and only if the expected payoff $\langle a_k, y\rangle$ exceeds the expected payoff $\langle a_\ell, y\rangle$. This tells us how to translate sites into payoffs; conversely, given a feasible payoff matrix with rows $a_1, \ldots, a_m$, one can transform them into sites of the form $\frac{1}{2}a_1 + t_1\mathbf{1}, \ldots, \frac{1}{2}a_m + t_m\mathbf{1}$ for suitable scalars $t_1, \ldots, t_m$.

So our proof establishes that if $P_1, \ldots, P_m$ are feasible best-response sets, then there is a corresponding payoff matrix — one that has the desired best-response sets — with rows of the form

$$a_i = 2v_i - \|v_i\|^2\mathbf{1} \tag{6}$$

for suitable v_i. By (3), with $\alpha = 1$ and $\beta = 0$, we see that these are quadratic scoring rules:

Corollary 3.2. *Given feasible best-response sets, there is a corresponding payoff matrix that is a quadratic scoring rule.*

A fortiori, *any* payoff matrix is just a quadratic scoring rule in disguise:

Theorem 3.3. *Each $m \times n$ payoff matrix A can be written as a quadratic scoring rule: there exist a constant $\alpha > 0$ and vectors $v_1, \ldots, v_m \in \mathbb{R}^n$ such that for each row $i = 1, \ldots, m$:*

$$a_i = \alpha\left(2v_i - \|v_i\|^2\mathbf{1}\right).$$

Remark 3.4. For our application (Sec. 2.2) to finite elicitation problems, the three results in this section paint a detailed picture of what scoring rules can achieve, including several new insights. By Theorem 3.1, a property PROP can be elicited if and only if the sets $\{y \in \Delta_n : i \in \text{PROP}(y)\}$ with $i = 1, \ldots, m$, specifying the beliefs for which the different messages are correct, are nearest-neighbor regions. Partial results along these lines have been around since at least de Finetti (1965), notably the link between small distances and large expected payoffs in (4) that establishes one implication in its proof. Also Friedman (1983, p. 450) points this link out, as do Lambert and Shoham (2009) in the computer science literature with additional restrictions on the properties being elicited. Our characterization telling for *all* finite elicitation problems *when* a property can be elicited — to our knowledge — is new. The other results tell much more, namely *how* to elicit them. Firstly (Cor. 3.2), the simple class of quadratic scoring rules is sufficient to elicit such properties. Secondly (Thm. 3.3), and perhaps more surprisingly, they are also necessary, in the sense that every other eliciting scoring rule can be rewritten as a quadratic one. This is in sharp contrast with elicitation of probabilities of events or distributions, where not only quadratic scoring rules do the right job, but so do scoring rules of different functional forms; recall, for instance, the discussion on quadratic, logarithmic, and spherical scoring rules in Winkler and Murphy (1968). ◁

4 Feasibility and best-response equivalence

The previous section addressed two of the three questions we set out to answer: whether there is a payoff matrix with given best responses and, if so, how to find one. Our final question concerned finding *all* of them, i.e., given a feasible best-response structure, characterizing all matrices with the desired best responses. We answer this question here, by reformulating it to finding the set of solutions to a system of finitely many linear (in)equalities (Thm. 4.1) and characterizing those solutions in game-theoretic terms (Thm. 4.2).

Consider polytopes $P_1, \ldots, P_m$ whose union is Δ_n: necessary conditions on best-response sets according to Proposition 2.1. Let $E = \cup_{i=1}^m \operatorname{ext}(P_i)$ denote the set of all extreme points of these polytopes. If we want the polytopes to be best-response sets of some $m \times n$ matrix A, then its rows $a_1, \ldots, a_m$ must satisfy

$$\langle a_i, y \rangle = \langle a_j, y \rangle \qquad \text{for all } y \in E \text{ and all } i, j \text{ with } y \in P_i \cap P_j, \tag{7}$$

$$\langle a_i, y \rangle > \langle a_j, y \rangle \qquad \text{for all } y \in E \text{ and all } i, j \text{ with } y \in P_i \setminus P_j. \tag{8}$$

The first condition simply says that if extreme point y lies in both P_i and P_j, then the expected payoffs from row i and row j should both be optimal and consequently the same. Likewise, the second condition requires that if extreme point y lies in P_i, but not in P_j, then row i should be a best response, but row j not, so the expected payoff from row i exceeds that of row j.

These finitely many linear (in)equalities assure that payoff matrix A gives the correct best responses in the extreme points of the polytopes; using a standard convexity argument, we will show that this gives the correct best responses against *any* strategy of the second player. In other words, conditions (7) and (8) are both necessary and sufficient for the existence of a payoff matrix with the desired best-response sets. Moreover, checking whether these conditions have a solution can be done via a linear-programming feasibility problem:

Theorem 4.1.

(a) Polytopes $P_1, \ldots, P_m$ with union Δ_n are best-response sets of an $m \times n$ matrix A if and only if its rows satisfy (7) and (8).

(b) The linear program (LP)

$$\begin{array}{lll} \textit{maximize} & t & \\ \textit{with} & (t, a_1, \ldots, a_m) \in \mathbb{R} \times (\mathbb{R}^n)^m & \\ & \langle a_i - a_j, y \rangle = 0 & \textit{for all } y \in E \textit{ and all } i, j \textit{ with } y \in P_i \cap P_j \\ & \langle a_i - a_j, y \rangle \geq t & \textit{for all } y \in E \textit{ and all } i, j \textit{ with } y \in P_i \setminus P_j \\ & t \leq 1 & \end{array}$$

has optimal value zero if there is no such matrix and optimal value one if there is. In the latter case, each optimum also gives a matrix with the desired properties.

If given best responses are feasible, finding all games with these best responses therefore boils down to solving a system (7) and (8) of homogeneous linear equalities and strict inequalities. These strict inequalities make the analysis tricky: the set of solutions is a convex cone, but typically not closed and consequently not finitely generated. In the appendix, we prove a technical Lemma A.1 that describes the set of solutions to such (in)equalities. We formulate its conclusions in game-theoretic terms here.

Let A be an $m \times n$ payoff matrix to the row player. We call A a ***dummy matrix*** if its rows are identical ($a_i = a_j$ for all $i, j = 1, \ldots, m$) or, equivalently, if its columns are constant. In a dummy matrix, each pure or mixed strategy gives the row player the same expected payoff, so each strategy is a best reply. This 'dummy' terminology comes from the literature on potential games; see Facchini et al. (1997, p. 195) and Voorneveld et al. (1999, Sec. 2). Dummy games play a crucial role in the axiomatization of the Nash equilibrium concept by Voorneveld (2019), who also discusses the role of such games for strategic invariance in the literature on equilibrium refinements and evolution.

If A and B are two $m \times n$ payoff matrices to the row player, we call them ***best-response equivalent*** if the row player has the same best-response correspondence in both cases; we say that payoff matrix B ***has a coarser best-response structure*** than A if best responses in A also are best responses in B: for each $y \in \Delta_n$, $\arg\max_i \langle a_i, y\rangle \subseteq \arg\max_i \langle b_i, y\rangle$.

Theorem 4.2 (Characterizing best-response equivalence)**.** *If there is a payoff matrix A with best-response sets $P_1, \ldots, P_m$, then there exist a positive integer k and matrices $C_1, \ldots, C_k$ with a coarser best-response structure than A such that the set of best-response equivalent matrices is of the form*

$$\mathscr{D} + \left\{ \sum_{\ell=1}^{k} \alpha_\ell C_\ell : \alpha_1, \ldots, \alpha_k > 0 \right\}.$$

Here, $\mathscr{D}$ is the set of $m \times n$ dummy matrices.

The closest result about best-response equivalence is in Morris and Ui (2004). Their Proposition 4 imposes an additional nondegeneracy assumption involving second-best replies and says the following about best-response equivalent payoff matrices A and B: if there are beliefs against which both the i-th and the j-th row are best responses, then the difference of these rows in A must be proportional to the difference of these rows in B.

Our result characterizes the best-response equivalent matrices without any additional restrictions and also shows that their proposition cannot be extended to cases where their nondegeneracy assumption is omitted. The following example illustrates this point, as well as the decomposition into matrices with a coarser best-response structure.

Example 4.3. Consider two payoff matrices for the row player:

$$A = \begin{bmatrix} 1 & -2 & 0 \\ 0 & 0 & 0 \\ -2 & 1 & 0 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 1 & -2 & 0 \\ 0 & 0 & 0 \\ -4 & 2 & 0 \end{bmatrix}$$

They are best-response equivalent: the best-response sets P_1, P_2, P_3 of the three rows are drawn in the unit simplex in the left panel of Figure 2. Its corners correspond with the standard basis vectors e_1, e_2, and e_3 where the column player picks some column with probability one; $y^1 = (2/3, 1/3, 0)$ and $y^2 = (1/3, 2/3, 0)$.

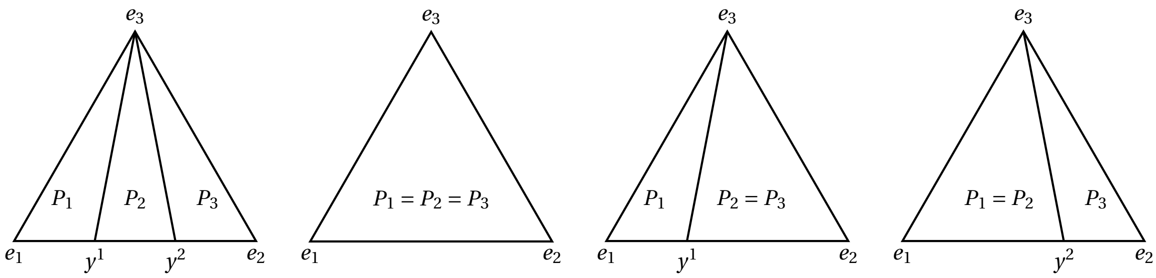

Figure 2: Best-response sets of A (and B) and the coarser ones from a dummy game, C_1, and C_2.

Both the first and the third row are best responses against belief $y = e_3$, the third column. But the difference of these rows in A is $(1,-2,0)-(-2,1,0) = (3,-3,0)$, whereas in B it is $(5,-4,0)$: these vectors are not proportional.

Writing down and simplifying conditions (7) and (8), the payoff matrices with these best-response sets can be written as

$$\underbrace{\begin{bmatrix} a & b & c \\ a & b & c \\ a & b & c \end{bmatrix}}_{\text{dummy}} + \alpha_1 \underbrace{\begin{bmatrix} 1 & -2 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{C_1} + \alpha_2 \underbrace{\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ -2 & 1 & 0 \end{bmatrix}}_{C_2} \text{ with } a,b,c \in \mathbb{R} \text{ and } \alpha_1, \alpha_2 > 0.$$

Matrix A above corresponds with $(a,b,c,\alpha_1,\alpha_2) = (0,0,0,1,1)$; for B, change α_2 to 2. Here we recognize the decomposition from Theorem 4.2: we start with a dummy game and matrices C_1 and C_2 have a coarser best-response structure consisting of suitable unions of the 'old' P_1, P_2, P_3 that we started with. Also these are illustrated in Figure 2. $\triangleleft$

In some games there are no nontrivial coarser best-response structures:

Example 4.4. Consider again the payoff matrix A from Figure 1. It is not possible to obtain a coarser best-response structure by, for instance, taking the union of P_1 and P_2. That would make the best-response structure look like

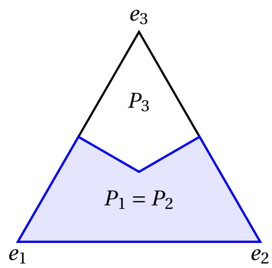

But that is infeasible: the shaded region where rows 1 and 2 would be best responses is not a polytope, in contradiction with Proposition 2.1. The payoff matrices where the row player has the same best-response correspondence as in A are of the form

$$\underbrace{\begin{bmatrix} a & b & c \\ a & b & c \\ a & b & c \end{bmatrix}}_{\text{dummy}} + \alpha_1 \underbrace{\begin{bmatrix} 0 & 0 & 0 \\ -1 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix}}_{C_1} \text{ with } a,b,c \in \mathbb{R} \text{ and } \alpha_1 > 0.$$

The matrix A is obtained by taking $(a, b, c, \alpha_1) = (1, 0, 0, 1)$. $\triangleleft$

A brief concluding remark: by changing the domain and codomain of potential best-response correspondences, our central questions of characterizing what constitutes a feasible best-response correspondence and of finding some or all utility functions giving rise to such a correspondence can be varied endlessly. A similar observation can be made for our application to information elicitation problems. We opted for a setting where we could provide full characterizations, leaving other settings as possible directions for future research.

A Proofs

A.1 Proof of Proposition 2.1

(a) The best-response set of pure strategy $i \in \{1, \dots, m\}$ can be rewritten as

$$P_i = \left\{ y \in \mathbb{R}^n : y_1, \dots, y_n \geq 0, \sum_{k=1}^n y_k = 1, \text{ and } \langle a_i, y \rangle \geq \langle a_j, y \rangle \text{ for all } j = 1, \dots, m \right\}.$$

So P_i is a bounded set of solutions to linear (in)equalities: it is a polytope.
(b) Trivial: there is a best response against each mixed strategy of player 2.
(c) Suppose two best-response sets P_i and P_j have a nonempty intersection. Each point $y \in P_i \cap P_j$ satisfies $\langle a_i, y \rangle = \langle a_j, y \rangle$, so $\langle a_j - a_i, y \rangle = 0$. And each point $y \in P_i \setminus P_j$ satisfies $\langle a_i, y \rangle > \langle a_j, y \rangle$, so $\langle a_j - a_i, y \rangle < 0$. Hence, $P_i \cap P_j$ is the set of maximizers of the linear function $y \mapsto \langle a_j - a_i, y \rangle$ over P_i, making it a face of P_i. Likewise, it is a face of P_j.

Next, suppose the intersection $\bigcap_{j \in J} P_j$ over an index set $J \subseteq \{1, \dots, m\}$ of more than two elements is nonempty. Let $k \in J$. Then $\bigcap_{j \in J} P_j = \bigcap_{j \in J, j \neq k} (P_k \cap P_j)$ is, by our previous step, the intersection of faces of P_k and consequently (Schrijver, 1986, Sec. 8.6) itself a face of P_k.

A.2 Proof of Theorem 3.1

First, assume that polytopes $P_1, \dots, P_m$ are nearest-neighbor regions of m sites $v_1, \dots, v_m$ in $\mathbb{R}^n$: $P_i = \text{Near}(v_i)$ for all $i = 1, \dots, m$. Define $m \times n$ matrix A with rows

$$a_i = 2v_i - \|v_i\|^2 \mathbf{1}, \qquad (i = 1, \dots, m) \tag{9}$$

Using (5) we have for each site v_i:

$$\begin{aligned} P_i &= \{ y \in \Delta_n : \|y - v_i\| \leq \|y - v_k\| \text{ for all } k = 1, \dots, m \} \\ &= \{ y \in \Delta_n : \langle a_i, y \rangle \geq \langle a_k, y \rangle \text{ for all } k = 1, \dots, m \}. \end{aligned}$$

By definition (1), these polytopes are the best-response sets of matrix A.

Conversely, assume that polytopes $P_1, \dots, P_m$ are best-response sets of $m \times n$ matrix A. Positive affine transformations of its entries do not affect best responses: without loss of generality we may take the entries to be nonnegative and so small that for each row $i = 1, \dots, m$:

$$0 \leq \langle a_i, \mathbf{1} \rangle \leq 1 \qquad \text{and} \qquad \|a_i\| \leq 1/\sqrt{n}. \tag{10}$$

Our previous step shows that if there are vectors $v_1, \dots, v_m \in \mathbb{R}^n$ satisfying (9), then for each i, the nearest-neighbor region $\text{Near}(v_i)$ equals the corresponding best-response set P_i and we're done. We construct such vectors of the form

$$v_i = \frac{1}{2} a_i + t_i \mathbf{1}, \qquad (i = 1, \dots, m)$$

for suitable scalars $t_1, \dots, t_m$. Substituting those v_i into (9), we must solve, for each row $i = 1, \dots, m$:

$$a_i = a_i + 2t_i \mathbf{1} - \left\| \frac{1}{2} a_i + t_i \mathbf{1} \right\|^2 \mathbf{1} = a_i + \left(2t_i - \left\| \frac{1}{2} a_i + t_i \mathbf{1} \right\|^2 \right) \mathbf{1}.$$

So for each i, the term in parentheses must be zero. Expand this:

$$0 = 2t_i - \left\| \frac{1}{2} a_i + t_i \mathbf{1} \right\|^2 = 2t_i - \left\langle \frac{1}{2} a_i + t_i \mathbf{1}, \frac{1}{2} a_i + t_i \mathbf{1} \right\rangle = 2t_i - \frac{1}{4} \|a_i\|^2 - \langle a_i, \mathbf{1} \rangle t_i - \|\mathbf{1}\|^2 t_i^2. \tag{11}$$

The last expression is a quadratic function of t_i. Using (10), its discriminant

$$\underbrace{(2 - \langle a_i, \mathbf{1} \rangle)^2}_{\geq 1} - \underbrace{\|\mathbf{1}\|^2}_{=n} \underbrace{\|a_i\|^2}_{\leq 1/n}$$

is nonnegative, so for each i there is indeed a root t_i solving (11).

A.3 Proof of Theorem 3.3

Consider an $m \times n$ matrix A with rows $a_1, \ldots, a_m$. We will find $\alpha > 0$ and $v_1, \ldots, v_m \in \mathbb{R}^n$ solving

$$a_i = \alpha \left(2v_i - \|v_i\|^2 \mathbf{1} \right), \qquad (i = 1, \ldots, m) \tag{12}$$

with the vectors v_i of the form

$$v_i = \frac{1}{2\alpha} a_i + t_i \mathbf{1}$$

for suitable scalars $t_1, \ldots, t_m$. Substituting those v_i into the right side of (12), we must solve for each row $i = 1, \ldots, m$:

$$a_i = \alpha \left(2v_i - \|v_i\|^2 \mathbf{1} \right) = a_i + \alpha \left(2t_i - \left\| \frac{1}{2\alpha} a_i + t_i \mathbf{1} \right\|^2 \right) \mathbf{1}.$$

Since we will choose $\alpha > 0$, the term in the big parentheses must be zero for each i. Expand this:

$$\begin{aligned} 0 &= 2t_i - \left\| \frac{1}{2\alpha} a_i + t_i \mathbf{1} \right\|^2 \\ &= 2t_i - \left\langle \frac{1}{2\alpha} a_i + t_i \mathbf{1}, \frac{1}{2\alpha} a_i + t_i \mathbf{1} \right\rangle \\ &= 2t_i - \frac{1}{4\alpha^2} \|a_i\|^2 - \frac{1}{\alpha} t_i \langle a_i, \mathbf{1} \rangle - t_i^2 \|\mathbf{1}\|^2 \\ &= -n t_i^2 + \left(2 - \frac{1}{\alpha} \langle a_i, \mathbf{1} \rangle \right) t_i - \frac{1}{4\alpha^2} \|a_i\|^2. \end{aligned}$$

The final expression is a quadratic function of t_i with discriminant

$$\left(2 - \frac{1}{\alpha} \langle a_i, \mathbf{1} \rangle \right)^2 - \frac{n}{\alpha^2} \|a_i\|^2.$$

This discriminant tends to 4 as $\alpha \to \infty$, so we can pick $\alpha > 0$ sufficiently large and make the discriminant nonnegative for each row i. Nonnegativity of the discriminants then implies that the quadratic functions in t_i have a root: we found our desired solution!

A.4 Proof of Theorem 4.1

(a) We discussed necessity of (7) and (8) already before Theorem 4.1. As for sufficiency, suppose $m \times n$ matrix A with rows $a_1, \ldots, a_m$ satisfies these conditions. Let $i \in \{1, \ldots, m\}$. We show that P_i is the best-response set of row a_i.

First, let $y \in P_i$. Then it is a convex combination of the extreme points of P_i. By (7) and (8), for each such extreme point $\tilde{y}$ and each $j = 1, \ldots, m$: $\langle a_i, \tilde{y} \rangle \geq \langle a_j, \tilde{y} \rangle$. Hence also the convex combination y satisfies $\langle a_i, y \rangle \geq \langle a_j, y \rangle$. So if $y \in P_i$, then a_i is a best response to y.

Next, let $y \in \Delta_n \setminus P_i$. Since $\cup_{j=1}^m P_j = \Delta_n$, there is a $k \neq i$ with $y \in P_k$. Write y as a convex combination of the extreme points of P_k: $y = \sum_{z \in \text{ext}(P_k)} \alpha_z z$ for some nonnegative scalars $(\alpha_z)_{z \in \text{ext}(P_k)}$ summing to one. Since y does not belong to P_i, there must be a $z^* \in \text{ext}(P_k)$ with $\alpha_{z^*} > 0$ such that z^* does not belong to P_i. By (7) and (8), we have $\langle a_k, z^* \rangle > \langle a_i, z^* \rangle$ and for all other $z \in \text{ext}(P_k)$: $\langle a_k, z \rangle \geq \langle a_i, z \rangle$. So

$$\langle a_k - a_i, y \rangle = \underbrace{\sum_{z \in \text{ext}(P_k), z \neq z^*} \alpha_z \langle a_k - a_i, z \rangle}_{\geq 0} + \underbrace{\alpha_{z^*} \langle a_k - a_i, z^* \rangle}_{>0} > 0:$$

row k gives a strictly higher payoff. So if $y \in \Delta_n \setminus P_i$, then a_i is not a best response to y.

From these two observations, we see that $y \in P_i$ if and only if row a_i is a best response to y.

(b) If (7) and (8) have a solution A, define

$$t = \min \left\{ \langle a_i - a_j, y \rangle : y \in E, i, j \text{ with } y \in P_i \setminus P_j \right\} > 0.$$

Then $(t, a_1, \ldots, a_m)$ satisfies all conditions of the LP, except possibly the final constraint $t \leq 1$. But rescaling by $1/t > 0$ gives a feasible point $(1, (1/t)a_1, \ldots, (1/t)a_m)$ with value 1 of the goal function. Since all feasible points of the LP have $t \leq 1$, this is clearly optimal and the matrix $\frac{1}{t}A$ has the desired best-response sets.

If (7) and (8) have no solution, there is no feasible $(t, a_1, \ldots, a_m)$ in the LP with $t > 0$. The zero vector in $\mathbb{R} \times (\mathbb{R}^n)^m$ is feasible in the LP. So in this case the maximal value of the goal function is zero.

A.5 Proof of Theorem 4.2

For real vectors a and b of the same dimension, write $a \geq b$ if $a_i \geq b_i$ for all coordinates i; likewise, $a > b$ if $a_i > b_i$ for all coordinates i. The proof relies on the following lemma:

Lemma A.1. *Let A and B be real matrices with the same number n of columns. Assume there is a solution $x \in \mathbb{R}^n$ to the system of linear (in)equalities*

$$Ax = \mathbf{0}, Bx > \mathbf{0}. \tag{13}$$

Then there are a positive integer k and vectors $y_1, \ldots, y_k$ in $\mathbb{R}^n$ such that its set of solutions is

$$\{x \in \mathbb{R}^n : Ax = \mathbf{0}, Bx > \mathbf{0}\} = \{x \in \mathbb{R}^n : Ax = \mathbf{0}, Bx = \mathbf{0}\} + \left\{ \sum_{i=1}^k \alpha_i y_i : \alpha_1, \ldots, \alpha_k > 0 \right\}.$$

Proof. For notational convenience, write

$$V = \{x \in \mathbb{R}^n : Ax = \mathbf{0}\} \qquad \text{and} \qquad W = \{x \in V : Bx = \mathbf{0}\}.$$

Both are null spaces of linear transformations, so V is a linear subspace of $\mathbb{R}^n$ and W is a linear subspace of V. Since V is finite-dimensional, it is the direct sum of W and its orthogonal complement $W^\perp = \{v \in V : \langle v, w \rangle = 0 \text{ for all } w \in W\}$. We prove:

$$\{x \in \mathbb{R}^n : Ax = \mathbf{0}, Bx > \mathbf{0}\} = W + \{z \in W^\perp : Bz > \mathbf{0}\}. \tag{14}$$

If $x \in \mathbb{R}^n$ has $Ax = \mathbf{0}$ and $Bx > \mathbf{0}$, then $x \in V$, so we can write $x = w + w^\perp$ with $w \in W$ and $w^\perp \in W^\perp$. Then

$$\mathbf{0} < Bx = Bw + Bw^\perp = \mathbf{0} + Bw^\perp.$$

So x is indeed the sum of an element $w \in W$ and an element $w^\perp \in \{z \in W^\perp : Bz > \mathbf{0}\}$.

Conversely, if $w \in W$ and $w^\perp \in \{z \in W^\perp : Bz > \mathbf{0}\}$, then $w + w^\perp \in W + W^\perp = V$, so

$$A(w + w^\perp) = \mathbf{0} \qquad \text{and} \qquad B(w + w^\perp) = \mathbf{0} + Bw^\perp > \mathbf{0}.$$

So $w + w^\perp$ solves the linear (in)equalities in (13). This proves (14).

We assumed that (13) has a solution. So $\{z \in W^\perp : Bz > \mathbf{0}\}$ and hence $\{z \in W^\perp : Bz \geq \mathbf{0}\}$ are nonempty. The latter is a polyhedral cone in a finite-dimensional space and consequently a finitely generated cone: there are finitely many $y_1, \ldots, y_k \in W^\perp$ such that

$$\{z \in W^\perp : Bz \geq \mathbf{0}\} = \left\{ \sum_{i=1}^k \alpha_i y_i : \alpha_1, \ldots, \alpha_k \geq 0 \right\}. \tag{15}$$

For ease of notation, abbreviate '$\sum_{i=1}^k \cdots$' by '$\sum_i \cdots$'. We next prove that

$$\{z \in W^\perp : Bz > \mathbf{0}\} = \left\{ \sum_i \alpha_i y_i : \alpha_1, \ldots, \alpha_k > 0 \right\}. \tag{16}$$

Let $z' \in \{z \in W^\perp : Bz > \mathbf{0}\}$. Defined by strict linear inequalities, $\{z \in W^\perp : Bz > \mathbf{0}\}$ is an open subset of $W^\perp$. So for $\varepsilon > 0$ sufficiently small, it also contains $z' - \varepsilon \sum_i y_i$. In particular, $B(z' - \varepsilon \sum_i y_i) \geq \mathbf{0}$. By (15), there are nonnegative $\alpha'_1, \ldots, \alpha'_k$ with $z' - \varepsilon \sum_i y_i = \sum_i \alpha'_i y_i$. So $z' = \sum_i (\alpha'_i + \varepsilon) y_i$, showing that z' a strictly positive combination of vectors $y_1, \ldots, y_k$.

Conversely, let $z = \sum_i \alpha_i y_i$ for strictly positive $\alpha_1, \ldots, \alpha_k$. Since $y_1, \ldots, y_k \in W^\perp$, also $z \in W^\perp$. By assumption, there is a $z^* \in W^\perp$ with $Bz^* > \mathbf{0}$. In particular, $Bz^* \geq \mathbf{0}$, so by (15), $z^* = \sum_i \alpha_i^* y_i$ for nonnegative $\alpha_1^*, \ldots, \alpha_k^*$. For each row b of B we have

$$0 < \langle b, z^* \rangle = \sum_i \alpha_i^* \langle b, y_i \rangle.$$

By (15), $\langle b, y_i \rangle \geq 0$ for all i. Therefore, at least one inequality is strict. Since all scalars $\alpha_1, \ldots, \alpha_k$ are strictly positive, this in turn implies that

$$\langle b, z \rangle = \sum_i \alpha_i \langle b, y_i \rangle > 0.$$

This holds for each row of B, so $z \in W^\perp$ solves $Bz > \mathbf{0}$, finishing the proof of (16).

Finally, substituting (16) into (14) proves the lemma. □

With this lemma in place, we are ready for the proof of Theorem 4.2: By Theorem 4.1, a payoff matrix $A \in \mathbb{R}^{m \times n}$ has the desired best-response sets if and only if it satisfies

$$\begin{aligned} &\langle a_i - a_j, y \rangle = 0 && \text{for all } y \in E \text{ and all } i, j \text{ with } y \in P_i \cap P_j, \\ &\langle a_i - a_j, y \rangle > 0 && \text{for all } y \in E \text{ and all } i, j \text{ with } y \in P_i \setminus P_j. \end{aligned}$$

This is a system of linear equations and strict inequalities over the $m \cdot n$ real entries of A. By Lemma A.1, its set of solutions is of the form

$$\mathscr{D} + \left\{ \sum_{i=1}^{k} \alpha_i C_i : \alpha_1, \ldots, \alpha_k > 0 \right\}, \tag{17}$$

where

$$\begin{aligned} \mathscr{D} = \{A \in \mathbb{R}^{m\times n} : &\langle a_i - a_j, y\rangle = 0 \text{ for all } y \in E \text{ and all } i, j \text{ with } y \in P_i \cap P_j, \\ &\langle a_i - a_j, y\rangle = 0 \text{ for all } y \in E \text{ and all } i, j \text{ with } y \in P_i \setminus P_j\}. \end{aligned}$$

$\mathscr{D}$ is the set of dummy matrices: Clearly, each dummy matrix lies in $\mathscr{D}$. Conversely, if A lies in the set on the right, pick a standard basis vector e_ℓ with $\ell \in \{1, \ldots, n\}$. This is an extreme point of Δ_n and hence of one of the polytopes, say P_i. Then we must have $\langle a_i - a_j, e_\ell\rangle = a_{i\ell} - a_{j\ell} = 0$ for all $j = 1, \ldots, m$: the ℓ-th column of A is constant. This holds for each column, so A is a dummy matrix.

Finally, each C_i lies in the closure of the set (17) of games that are best-response equivalent with A. So if a strategy is a best response to some belief in A, it is also a best response in C_i: payoff matrix C_i has a coarser best-response structure.

References

F. Aurenhammer, R. Klein, and D.-T. Lee. *Voronoi diagrams and Delaunay triangulations*. World Scientific Publishing, Singapore, 2013.

D. Bergemann, S. Morris, and S. Takahashi. Interdependent preferences and strategic distinguishability. *Journal of Economic Theory*, 168:329–371, 2017.

S. Borgwardt and R. M. Frongillo. Power diagram detection with applications to information elicitation. *Journal of Optimization Theory and Applications*, 181:184–196, 2019.

G. W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78: 1–3, 1950.

A. Carvalho. An overview of the applications of proper scoring rules. *Decision Analysis*, 13:223–242, 2016.

B. de Finetti. Methods for discriminating levels of partial knowledge concerning a test item. *British Journal of Mathematical and Statistical Psychology*, 18:87–123, 1965.

G. Facchini, F. van Megen, P. Borm, and S. Tijs. Congestion models and weighted potential games. *Theory and Decision*, 42:193–206, 1997.

D. Friedman. Effective scoring rules for probabilistic forecasts. *Management Science*, 29:447–454, 1983.

T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359–378, 2007.

B. Grofman. Downs and two-party convergence. *Annual Review of Political Science*, 7:25–46, 2004.

H. Hotelling. Stability in competition. *The Economic Journal*, 39:41–57, 1929.

N. Lambert and Y. Shoham. Eliciting truthful answers to multiple-choice questions. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 109–118, 2009.

S. Morris and T. Ui. Best response equivalence. *Games and Economic Behavior*, 49:260–287, 2004.

T. Offerman, J. Sonnemans, G. van de Kuilen, and P. P. Wakker. A truth serum for non-Bayesians: Correcting proper scoring rules for risk attitudes. *Review of Economic Studies*, 76:1461–1489, 2009.

K. H. Osband. *Providing incentives for better cost forecasting*. PhD thesis, University of California, Berkeley, 1985.

M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80:2595–2647, 2012.

L. J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66:783–801, 1971.

K. H. Schlag, J. Tremewan, and J. J. van der Weele. A penny for your thoughts: a survey of methods for eliciting beliefs. *Experimental Economics*, 18:457–490, 2015.

A. Schotter and I. Trevino. Belief elicitation in the laboratory. *Annual Review of Economics*, 6: 103–128, 2014.

A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, Inc., New York, NY, 1986.

W. Thomson. Eliciting production possibilities from a well-informed manager. *Journal of Economic Theory*, 20:360–380, 1979.

M. Voorneveld. An elementary axiomatization of the Nash equilibrium concept. SSE Working Paper Series in Economics, No. 2019:1. Stockholm School of Economics, 2019.

M. Voorneveld, P. Borm, F. van Megen, S. Tijs, and G. Facchini. Congestion games and potentials reconsidered. *International Game Theory Review*, 1:283–299, 1999.

R. L. Winkler and A. H. Murphy. “Good” probability assessors. *Journal of Applied Meteorology*, 7: 751–758, 1968.