

Bounded Rationality and the Choice of Jury Selection Procedures

Martin Van der Linden *Utah State University*

Abstract

A peremptory-challenge procedure allows the parties to a jury trial to dismiss some prospective jurors without justification. Complex challenge procedures offer an unfair advantage to parties who are better able to strategize. I introduce a new measure of strategic complexity based on level- k thinking and use this measure to compare challenge procedures often used in practice. In applying this measure, I overturn some commonly held beliefs about which jury selection procedures are strategically simple.

1. Introduction

It is customary to let the parties involved in a jury trial dismiss some of the potential jurors without justification. Procedures for dismissal are known as peremptory-challenge procedures. Such procedures are used in many countries, including the United States.¹ A variety of procedures are used in practice. These procedures differ notably in their strategic complexity. More complex procedures give an unfair advantage to parties who are strategically skilled or can devote ample resources to hiring jury consultants. It is therefore important to iden-

I am grateful to the editor and a referee for useful remarks and suggestions. I am also grateful to Geir Asheim, Benoit Decerf, Paul Edelman, Hanna Frank, Eun Jeong Heo, Greg Leo, Clayton Masterman, Claudia Rei, Roberto Serrano, Yves Sprumont, Rodrigo Velez, John Weymark, and Myrna Wooders for helpful discussions and comments. I also thank the participants at presentations at Aix-Marseille University, Catholic University of Louvain, KU Leuven, Vanderbilt University, the 13th meeting of the Society for Social Choice and Welfare, and the Fifth World Congress of the Game Theory Society for their questions and suggestions. I am thankful to the Center for Operations Research and Econometrics and the Center for Research in Economics for their hospitality during visits. Support from the Kirk Dornbush summer research grant and National Science Foundation grant IIS-1526860 is gratefully acknowledged.

¹ In *Swain v. Alabama* (380 U.S. 202 [1965]), the Supreme Court affirmed that “[t]he right to challenge a given number of jurors without showing cause is one of the most important of the rights secured to the accused” (*Pointer v. United States*, 151 U.S. 396, 408 [1894]). Following *Batson v. Kentucky* (476 U.S. 79 [1986]), a party can disqualify a peremptory challenge by her opponent if she can prove that the challenge was based on race. (Challenges based on gender were later also prohibited following *J. E. B. v. Alabama* [511 U.S. 127 (1994)].) However, *Batson v. Kentucky* is notoriously hard to implement, and judges rarely rule in favor of *Batson* challenges (Marder 2012; Daly 2016).

[*Journal of Law and Economics*, vol. 61 (November 2018)]

© 2018 by The University of Chicago. All rights reserved. 0022-2186/2018/6104-0023\$10.00

tify the procedures that are strategically simple in order to level the playing field among parties. In this paper, I introduce the concept of a dominance threshold, a new measure of strategic complexity based on level- k thinking, and I use this measure to compare the complexity of some challenge procedures commonly used in practice (see Crawford, Costa-Gomes, and Iriberry [2013] for a survey of the level- k literature).

Fairness is an important issue in jury selection. One feature of a procedure that impacts fairness is its strategic complexity. If a procedure is complex, parties with better strategic skills are likely to secure more favorable juries. This is particularly relevant in jury selection in which the parties invest significant resources for developing an effective strategy. For example, jury selection consultancy has become a well-established industry.² Using strategically simple procedures limits the impact on the selected jury of differences in the parties' ability to strategize or in their financial means to hire jury consultants.

1.1. Comparing Strategic Complexity

Comparing the strategic complexity of jury selection procedures presents two challenges. First, jury selection procedures are indirect mechanisms because the parties' actions consist of dismissing jurors rather than revealing their preferences. Second, in some procedures commonly used in practice, the parties submit their challenges simultaneously, which induces games of imperfect information. These two difficulties make it impossible to apply measures of strategic complexity previously developed in the literature (see Section 1.2).

I overcome these difficulties by introducing the concept of a dominance threshold. Given some assumption about the strategies her opponent could play—henceforth, a model of her opponent—a party has a dominant strategy if one of her strategies is a best response to any strategy of her opponent that is consistent with her model. The objective is to identify models for which the parties have dominant strategies. This is accomplished by iteratively eliminating strategies that are never-best responses. The dominance threshold is the number of rounds of elimination needed to reach models in which both parties have a dominant strategy. The dominance threshold measures the complexity of the model of opponents that the parties need in order to have a dominant strategy. For example, a dominance threshold of 1 corresponds to the parties having a dominant strategy given any model of their opponent. When the dominance threshold is 2, the parties need to know only that their opponent is a best responder in order to have a dominant strategy.

² The widespread use of jury consultants is evidenced by the existence of the American Society of Trial Consultants and its journal *The Jury Expert: The Art and Science of Litigation Advocacy*. Jury consultants explicitly describe how part of their job is concerned with the strategic use of challenges. Jury consultant Roy Futterman, for example, writes, “Caditz argues that . . . jury selectors pay . . . little to no attention to the strategic use of strikes [that is, peremptory challenges]. . . . [I]t is a bit of a reach to say that strategy is barely utilized. In my experience, . . . [jury selection] comes closer to a long battle of stealth, counter-punches, misdirection, and hand-to-hand combat than a lofty academic experience” (Futterman 2014).

Many judges appear to share the concern about selecting strategically simple procedures and have developed procedures that attempt to limit the parties' ability to strategize. In a report on judges' practices regarding peremptory challenges, Shapard and Johnson (1994, p. 6) write, "Some judges require that peremptories be exercised [following procedure X]. . . . This approach . . . makes it more difficult to pursue a strategy prohibited by *Batson* (or any other strategy)." "A more extreme approach to the same end . . . is [procedure Y]. . . . This approach imposes maximum limits on counsel's ability to employ peremptories in a strategic manner" (Shapard and Johnson [1994], p. 6 n. 6).

Using the dominance threshold as a measure of strategic complexity provides new insights and challenges some commonly held beliefs about jury selection procedures. Shapard and Johnson (1994, p. 6) write, "Other judges, for the same purposes [that is, limiting the parties' ability to strategize], allow all peremptories to be exercised after all challenges for cause, but with the parties making their choices 'blind' to the choices made by opposing parties (in contrast to alternating 'strikes' from a list of names of panel members)." I show that, contrary to these judges' beliefs, procedures in which challenges are sequential tend to be strategically simpler than procedures in which challenges are simultaneous: by generating imperfect-information games, simultaneous procedures increase the amount of guesswork needed to determine optimal strategies.

I also study the design of maximally simple jury selection procedures. I show that it is impossible to construct a reasonable procedure that allows the parties to challenge jurors and always have a dominant strategy. Hence, the smallest achievable dominance threshold is 2. Such a minimal dominance threshold is attained by a procedure that I call the sequential one-shot procedure, in which the parties sequentially submit a single list of jurors whom they want to challenge.

1.2. *Related Literature*

This paper differs from the previous game-theoretic literature on jury selection procedures in at least two ways (see Flanagan [2015] for a recent review). First, the literature focuses on subgame-perfect equilibrium as a solution concept (two exceptions are Bermant [1982] and Caditz [2015]). Subgame perfection requires a high level of strategic sophistication, especially in complex procedures. By relying on the concept of a dominance threshold, this paper accounts for the possibility of boundedly rational parties. I show how the dominance threshold, which measures the amount of common knowledge and rationality needed to have a dominant strategy, can be used to measure the strategic complexity of a procedure.

Second, most of the literature focuses on the characterization and properties of equilibria of different procedures (see Roth, Kadane, and DeGroot 1977; Brams and Davis 1978; DeGroot and Kadane 1980; Kadane, Stone, and Wallstrom 1999; Alpern and Gal 2009; Alpern, Gal, and Solan 2010). When the performance of procedures is compared, it is typically in terms of their effects on the composition

of the jury. These comparisons have yielded few policy recommendations (see, however, Bermant 1982; Flanagan 2015, sec. 4.2). In contrast, this paper compares procedures with respect to the standard objective of limiting the parties' ability to strategize. This latter approach enables a clear comparison of some of the procedures used in practice.

Focusing on strategic complexity also implies that this paper falls short of providing a full-fledged implementation analysis. Bounded rationality has been considered in implementation theory. Abreu and Matsushima (1992) notably show that any social choice function can be (virtually) implemented in iteratively undominated strategies. Because the authors do not restrict the number of iterations, the mechanism they propose can, depending on the application, have a very high dominance threshold.³ Instead of fixing a solution concept and investigating the social choice functions it allows to be implemented, this paper focuses on solution concepts themselves. I argue that among a particular class of solution concepts, the weakest solution concept that enables solving a procedure is a useful measure of a game's strategic complexity. A natural subsequent question—which is left open—is to determine the best procedure (according to some outcome-oriented objective) that can be solved by a reasonably weak solution concept (as in de Clippel, Eliaz, and Knight 2014).⁴

This paper also focuses heavily on procedures that are relevant in practice for jury selection. As a consequence, my analysis incorporates a number of institutional constraints that are specific to jury selection. In particular, I consider only procedures in which the parties' actions are limited to challenging some prospective jurors. Selection procedures based on other action spaces (for example, direct mechanisms) are known to have interesting properties (Barberà and Coelho 2008).⁵ For restricted domains, some of those procedures even have dominant strategies (Barberà, Sonnenschein, and Zhou 1991; Barberà, Massó, and Neme 2005) or can be solved by finitely rational players (Abreu and Matsushima 1992; de Clippel, Eliaz, and Knight 2014). However, those procedures involve—sometimes complex—action spaces that go beyond the simple challenge of jurors. My focus on challenge procedures is motivated by the fact that, in the context of jury selection, the law and legal customs often specifically limit the parties' actions to challenges. One contribution of this paper is to show that the few selection procedures with dominant strategies identified by Barberà, Massó, and Neme (2005) do not survive conferring the parties with minimal challenge abilities (see also Van der Linden 2017).

The dominance threshold relates to a recent strand of the literature that com-

³ Another difference is that Abreu and Matsushima (1992) rely on iteratively undominated strategies, whereas most of this paper deals with iteratively never-best responses. See Section 8 for a discussion of how my results generalize to iterated undominated strategies.

⁴ In this respect, it would be useful to develop a parametrized outcome-oriented objective for jury selection procedures. One could then weigh trade-offs between the degree to which a procedure satisfies the objective and the procedure's level of strategic complexity.

⁵ For a general characterization of mechanisms with a dominance threshold of 2, see Börgers and Li (2017).

pare the incentive properties of mechanisms that fail to have dominant strategies.⁶ For example, Pathak and Sönmez (2013) and Arribillaga and Massó (2016) define comparison criteria for direct games. In indirect games, de Clippel, Eliaz, and Knight (2014) recommend focusing on procedures that can be solved in two rounds of backward induction. More generally, this last recommendation suggests using the number of rounds of backward induction needed to solve a procedure as a measure of its complexity.

Measures developed for direct mechanisms are not well suited for the comparison of jury selection procedures because those procedures induce indirect games. In addition, an important question in the choice of a jury selection procedure is whether challenges should be simultaneous or sequential. Because backward induction is defined only for perfect-information games, a different measure is therefore required. (When games are modeled in extensive form, any simultaneous move implies that the game is of imperfect information.)⁷ Unlike previous measures in the literature, the dominance threshold can be used to compare the strategic complexity of any pair of games, including indirect games and games of imperfect information. Although the focus of this paper is the study of jury selection procedures, the dominance threshold applies more generally as a general measure of strategic complexity.

The paper is organized as follows. Section 2 introduces the model, several examples of jury selection procedures, and a general class of procedures. In Section 3, I show that most reasonable jury selection procedures do not have dominant strategies. Section 4 formally introduces the concept of a dominance threshold. The dominance threshold is then applied to comparing the strategic complexity of jury selection procedures in Sections 5 and 6. Whereas the previous sections considered a complete-information setting, Section 7 considers an extension to situations of incomplete information. In Section 8, I consider further extensions and discuss some open questions. Proofs are in the Online Appendix.

2. Model and Procedures

I focus on struck procedures. In addition to peremptory challenges, which require no justification, the parties can raise challenges for cause, which must be

⁶ In contrast, Li (2015) proposes a criterion to compare the incentive properties of different mechanisms that all have dominant strategies.

⁷ Even under complete information, one must be careful to prune game trees before using the number of rounds of backward induction as a measure of strategic complexity. For example, consider the divide-and-choose procedure for the fair division of a divisible endowment. The procedure has a dominance threshold of 2 and can be solved in two rounds of backward induction. In contrast, consider the divide-and-choose-and-raise-your-hand procedure. In this alternate procedure, after the endowment has been shared, players sequentially raise or lower their hands, with the show of hands having no impact on the final allocation. Unlike the divide-and-choose procedure, the divide-and-choose-and-raise-your-hand procedure requires four rounds of backward induction to be solved because of the addition of the inconsequential raise-your-hand action. The latter actions should therefore be pruned before applying the measure of backward-induction complexity. In contrast, the dominance threshold is insensitive to the addition of the inconsequential action and does not require pruning the game tree.

based on some bias recognized by law, such as being a direct relative of one of the parties. As explained by Bermant and Shapard (1981, p. 92), the defining feature of a struck procedure “is that the judge rules on all challenges for cause before the parties claim any peremptories. Enough potential jurors are examined to allow for the size of the jury plus the number of peremptory challenges allotted to both sides. In a federal felony trial, for example, the jury size is twelve; the prosecution has six peremptories, and the defense has ten. Under the struck jury method, therefore, 28 potential jurors are cleared through challenges for cause before the exercise of peremptories.”⁸

Struck procedures are commonly used in federal courts. In a 1977 survey of judges’ practices regarding the exercise of peremptory challenges, 55 percent of federal district judges reported using a struck procedure (Bermant and Shapard 1981). Today, the use of a struck procedure is, for example, recommended by law as the preferred method for criminal cases other than first-degree murder in Minnesota (Minn.R. Crim. P. 26.02, subdiv. 4.[3]b).

2.1. *The Model*

The set of prospective jurors left after all challenges for cause have been raised is $N = \{1, \dots, n\}$. The defendant D and the plaintiff P are allowed c_D and c_P peremptory challenges, respectively. From N , a jury J of b jurors must be selected. The jurors in J are the impaneled jurors. As explained above, when struck procedures are considered, $n = b + c_D + c_P$ to allow the parties to challenge up to c_D and c_P jurors.

Let \mathcal{J} be the set of juries containing b jurors, and let $\Delta\mathcal{J}$ be the set of lotteries on \mathcal{J} . In some cases, it is possible that after all challenges have been raised, more than b jurors remain unchallenged. In this case, I assume that the b impaneled jurors are chosen at random among the unchallenged jurors. As a consequence, the parties have expected utility preferences R_D and R_P on $\Delta\mathcal{J}$ with corresponding Bernoulli utility functions u_D and u_P on \mathcal{J} . A pair of preferences (R_P, R_D) is called a preference profile (hereafter, profile), and a quintuple (R_P, R_D, c_P, c_D, b) is called a jury selection problem (hereafter, problem).

Throughout, I assume for simplicity that preferences on juries are separable; that is, if replacing juror h by juror j in jury J is an improvement according to u_P , then the same is true when h is replaced by j in any other jury J' . Formally, for any $i \in \{D, P\}$, any $J, J' \in \mathcal{J}$, any $h \in J \cap J'$, and any $j \in N \setminus (J \cup J')$, we have $u_i(J \cup \{j\} \setminus \{h\}) \geq u_i(J)$ if and only if $u_i(J' \cup \{j\} \setminus \{h\}) \geq u_i(J')$.

When describing their approach to jury selection, jury consultants often suggest the use of separable (and even additive) preferences relying on some rating or scoring of the individual jurors (Caditz 2014; Leibold 2015). However, other

⁸ This contrasts with strike-and-replace procedures, in which challenges for cause and peremptory challenges are intertwined. In a strike-and-replace procedure, prospective jurors who are challenged (either for cause or peremptorily) are replaced by new jurors from the pool and “to one degree or another, counsel exercise their challenges without knowing the characteristics of the next potential juror to be interviewed” (Bermant and Shapard 1981, p. 93).

consultants point to the importance of interactions between jurors that cannot be captured by separable preferences (see, for example, Futterman 2014).

Importantly, most of the results in this paper still hold when separability is not assumed. Separability is mainly used to facilitate the formulation of a condition on profiles, which, if satisfied, make simultaneous procedures infinitely complex. I discuss the role of separability in more detail after the statements of propositions 4, 5, and 6.

Separability eases the exposition because it implies that the preferences R_D and R_P induce well-defined preferences for individual jurors. It is also assumed that the preferences for jurors induced by R_D and R_P are strict. With a slight abuse of notation, R_i serves to denote i 's preferences for individual jurors and for juries.

An extreme kind of profile is a juror-inverse profile. A profile is juror inverse if R_D and R_P induce inverse preferences in jurors (that is, for all $j, h \in N$, $j R_P h$ if and only if $h R_D j$). Unlike separability, which is assumed throughout the paper, juror-inverse profiles are considered only as a special case.

2.2. Procedures

As attested to by Bermant and Shapard (1981), a wide variety of struck procedures are used by judges. One common type of struck procedures are those I call one-shot procedures. In a one-shot procedure, each party $i \in \{D, P\}$ submits a single list of up to c_i jurors in N that i wants to challenge. Depending on the procedure, the parties submit their lists simultaneously (the one-shot_M procedure) or sequentially (the one-shot_Q procedure). The impaneled jurors are the jurors in N who have not been challenged. If more than b jurors are left unchallenged, the b impaneled jurors are drawn at random from among the unchallenged jurors.⁹

Another common type of struck procedure are the procedures that I call alternating. Alternating procedures proceed through a succession of rounds in which the parties can challenge as many jurors in N as they have challenges left. Again, an alternating procedure can be either simultaneous (alternating_M) or sequential (alternating_Q) depending on whether challenges are submitted simultaneously or sequentially in each round. In the alternating_M procedure, if both parties challenge the same juror in a given round, both parties are charged with the challenge and can challenge one less juror.

Alternating procedures end when neither party has challenges left or when both parties abstain from challenging jurors in a single round. The impaneled jurors are the jurors left unchallenged in N or a random draw of b of these jurors if more than b jurors are left unchallenged.¹⁰

⁹ The use of the one-shot_M procedure is documented by Bermant (1982, pp. 17–18, comments of Judges John Feikens and Donald S. Voorhees). Bermant (1982, p. 17, comments of Judge William B. Enright) shows that a procedure in which the parties alternate challenges twice has been used in practice, with each party allowed to challenge up to $c_i/2$ jurors in each round. Other than the fact that multiple jurors are selected (instead of a single arbitrator), the one-shot_Q procedure is strategically equivalent to the shortlisting procedure proposed by de Clippel, Eliaz, and Knight (2014).

¹⁰ Simultaneous challenges are used in alternating procedures for civil cases in Tennessee (Tenn. R. Civ. P. 47.03), although the mandated procedure in these cases is of the strike-and-replace type (see note 8).

One-shot and alternating procedures are members of the class of N -struck procedures in which parties take turns challenging jurors from N for a number of rounds.¹¹ Formally, every N -struck procedure consists of a maximum of $f \geq 1$ rounds, where f differs between procedures. Each round $r \in \{1, \dots, f\}$ is characterized by a maximum number of challenges $x_i^r \geq 1$ for each party, with $\sum_{r=1}^f x_i^r \geq c_i$. The number of challenges party i has left in round r is ℓ_i^r , with $\ell_i^1 = c_i$. In each round r , (A) the parties can challenge up to $\min\{x_i^r, \ell_i^r\}$ jurors among the jurors in N who have not yet been challenged. Challenges are sequential if the procedure is sequential and simultaneous if the procedure is simultaneous. (B) For each party $i \in \{D, P\}$, the number of challenges left is decreased by the number of jurors that the party challenged in A (that is, ℓ_i^{r+1} equals ℓ_i^r minus the number of jurors that the party challenged in A).

The procedure terminates when no party has challenges left, when round f is reached, or when both parties abstain from challenging jurors in a single round. The jurors left unchallenged when the procedure terminates are the impeaneled jurors. If more than b jurors are left unchallenged when the procedure terminates, the b impeaneled jurors are drawn at random from the unchallenged jurors.¹²

One-shot procedures are N -struck procedures with $f = 1$ and $x_i^1 = c_i$ for both parties $i \in \{D, P\}$. Alternating procedures are N -struck procedures with $x_i^r = c_i$ for both $i \in \{D, P\}$ and all $r \in \{1, \dots, f\}$, and $f = 2 \max_{i \in \{D, P\}} c_i$. Besides one-shot and alternating procedures, the class of N -struck procedures includes, for example, the two-round procedure documented by Bermant (1982, p. 17, comments of Judge Enright) and described in note 2.

From a game-theoretic point of view, a (jury selection) procedure is an extensive game form $\Gamma : \mathcal{S}_D \times \mathcal{S}_P \rightarrow \Delta \mathcal{J}$ that associates any pair of strategies (s_D, s_P) in some strategy space $\mathcal{S}_D \times \mathcal{S}_P$ with a lottery on juries in \mathcal{J} . In this paper, I restrict attention to pure strategies in any extensive game form Γ , although the results also hold when mixed strategies are allowed.

3. Impossibility Results

Given preference R_i , a best response for party i to some strategy s_{-i} of her opponent is a strategy $t_i(s_{-i})$ such that

$$\Gamma[t_i(s_{-i}), s_{-i}] R_i \Gamma(s'_i, s_{-i}) \quad \text{for all } s'_i \in \mathcal{S}_i. \tag{1}$$

When $-i$ plays s_{-i} and i plays $t_i(s_{-i})$, party i plays a best response to $-i$. A strategy $s_i \in \mathcal{S}_i$ is dominant for i given some model $\mathcal{S}_{-i} \subseteq \mathcal{S}_{-i}$ of her opponent if s_i is a best response to every strategy $s_{-i} \in \mathcal{S}_{-i}$. A dominant strategy is a strategy $s_i^* \in \mathcal{S}_i$ that is a best response for i to any strategy $s_{-i} \in \mathcal{S}_{-i}$. In other words, a

¹¹ The term “ N -struck procedure” emphasizes the fact that, in each round, the parties can challenge any juror in N who has not been challenged yet. This is not the case in every struck procedure (Bermant 1982, pp. 16–17, comments of Judge C. Clyde Atkins).

¹² The distribution is arbitrary as long as any remaining juror has a strictly positive probability of being selected.

dominant strategy is a strategy that is dominant for i given any model of her opponent.

Given some domain of preferences, a dominant-strategy procedure is a procedure in which both parties have a dominant strategy for every profile in the domain. Dominant-strategy procedures are strategically simple because each party can determine an optimal strategy regardless of her guess about the strategy of her opponent. Dominant-strategy procedures guarantee a form of equality among equals: two parties having the same preferences but different abilities to form expectations about their opponent’s strategy should be able to secure similar outcomes.

It is useful to relate dominant strategies with level- k thinking (see the survey in Crawford, Costa-Gomes, and Iriberry [2013]). In level- k terminology, an L_i^0 party is a nonstrategic party who could potentially play any strategy. An L_i^1 party assumes that her opponent is L_{-i}^0 , makes a guess about the L_{-i}^0 strategy s_{-i}^0 that her opponent will employ, and plays a best response to s_{-i}^0 .¹³ Similarly, an L_i^k party assumes that her opponent is L_{-i}^{k-1} , makes a guess about the L_{-i}^{k-1} strategy s_{-i}^{k-1} that her opponent will employ, and plays a best response to s_{-i}^{k-1} .

Observe that, because an L_{-i}^0 strategy can be any of $-i$ ’s strategies, i has a dominant strategy if and only if i has an L_i^1 strategy that is a best response to every L_{-i}^0 strategy of her opponent. In the language of level- k thinking, a dominant-strategy procedure limits the impact of differences in strategic skills because i can determine an optimal strategy regardless of her belief about her opponent’s level of rationality k_{-i} or her guess about which L_{-i}^{k-i} strategy her opponent will employ.

Unfortunately, most reasonable procedures that permit challenges do not have a dominant strategy. Consider the one-shot_M procedure. In the one-shot_M procedure, i ’s only best response to any s_{-i} is to challenge her c_i worst jurors among the jurors whom $-i$ does not challenge in s_{-i} . As illustrated in example 1, such a best response is highly dependent on the challenges chosen by $-i$. Hence, the one-shot_M procedure is not a dominant-strategy procedure.

Example 1. Suppose that each juror has four challenges ($c_D = c_P = 4$) and one juror must be selected ($b = 1$). A set of nine prospective jurors $N = \{1, \dots, 9\}$ will therefore remain after all challenges for cause have been raised. Let D ’s preference for these nine jurors be $1 R_D 2 R_D \dots R_D 9$. If P challenges the circled jurors in preference (P1), then D ’s best response is to challenge the boxed jurors:

R_D : 1 (2) (3) [4] (5) (6) [7] [8] [9] (P1)

If P instead challenges the circled jurors in preference (P2), D ’s best response is again to challenge the boxed jurors:

R_D : 1 (2) (3) (4) (5) [6] [7] [8] [9] (P2)

¹³ Recall that I consider only pure strategies. Hence, the set of i ’s level-0 strategies is the set of i ’s pure strategies. Again, all the results in this paper hold when mixed strategies are allowed. In particular, the results hold when i ’s level-1 strategies include i ’s best responses to probabilistic beliefs about the (pure) level-0 strategy that $-i$ will employ, as in Ho, Camerer, and Weigelt (1998).

Clearly, challenging the boxed jurors in preference (P2) is not a best response for D to P's challenging the circled jurors in preference (P1), which shows that the one-shot_M procedure is not a dominant-strategy procedure in any domain that has a profile containing R_D .

As shown in proposition 1, the preceding example generalizes to the whole class of N -struck procedures and to any problem. Intuitively, in any N -struck procedure, if $-i$ does not challenge any jurors, then i 's best response is to challenge her c_i worst jurors. On the other hand, if $-i$ challenges one of the c_i worst jurors of i , say w , then i is better off not challenging w and challenging one of her other c_i worst jurors. Recall that a (jury selection) problem is a quintuple (R_D, R_P, c_D, c_P, b) .

Proposition 1. For any problem, (i) the first party does not have a dominant strategy in the one-shot_Q procedure, and (ii) neither party has a dominant strategy in any N -struck procedure different from the one-shot_Q procedure.

Note that the one-shot_M procedure is an N -struck procedure different from the one-shot_Q procedure. Hence, proposition 1 shows that, for every problem, neither party has a dominant strategy in the one-shot_M procedure. The one-shot_Q procedure is the exception among N -struck procedures: it is the only N -struck procedure in which one of the parties—the second party to challenge—has a dominant strategy, although the other party does not for the reason explained before proposition 1. (See the proof of proposition 1 in the Online Appendix for more detail.)

Of course, N -struck procedures are only a small subset of all possible jury selection procedures. Other procedures used in practice include strike-and-replace procedures (see note 8) and other struck procedures in which the parties can challenge only from subsets of N in each round (Bermant 1982, pp. 16–17, comments of Judge Atkins). To my knowledge, the law recognizes only the challenge of jurors as an appropriate way for the parties to influence the composition of the jury, and procedures based on such challenges are the only ones used in practice. However, one could think of using other types of procedures, such as the classical selection procedures mentioned in Section 1 and involving the direct revelation of preferences (see, for example, Barberà, Massó, and Neme 2005). Other indirect procedures could also be considered, such as procedures in which the parties nominate jurors rather than challenge them, maybe with the option for a party to refuse some of the other party's nominations. Flanagan (2015) argues that procedures in which the parties are allowed to challenge entire juries (rather than individual jurors) may have desirable properties.¹⁴

In view of this variety of non- N -struck procedures, it is natural to ask whether

¹⁴For example, the parties can be repeatedly presented with potential juries of b jurors drawn from N . In each round, a new potential jury is drawn, and the parties are given the ability to reject the proposed jury a number of times. The jury proposed in a given round then becomes the effective jury if neither party challenges it in that round (either because the parties have no challenges left or because they abstain from using them).

there exist dominant-strategy procedures for jury selection outside the N -struck class. Proposition 2 shows that if such procedures exist, then they must either deprive a party of her right to challenge at least one juror in N or be so intricate that they are unlikely to be used in practice.

A procedure satisfies finiteness if the set of its decision nodes is finite for both parties and for nature. A procedure satisfies minimal challenge if for every prospective juror $j \in N$, both parties $i \in \{D, P\}$ have a strategy $s_i^j \in \mathcal{S}_i$ such that j is never part of the chosen jury when i plays s_i^j .¹⁵ Every N -struck procedure satisfies both finiteness and minimal challenge (strategy s_i^j can, for example, involve challenging juror j —and only juror j —in the first round).¹⁶

Proposition 2. In the domain of separable preferences, no dominant-strategy procedure satisfies both finiteness and minimal challenge.

In the Online Appendix, I show that proposition 2 is, in fact, true for smaller domains of profiles, including the domain of additive profiles.

4. A Measure of Strategic Complexity

Propositions 1 and 2 show that most procedures are not strategically simple in the sense that both parties cannot always follow the simple recommendation of playing a dominant strategy. This does not mean, however, that judges should give up on the idea of using procedures that are as simple as possible. This section and Section 5 show that, although procedures generally fail to feature dominant strategies, not all procedures are equal in terms of strategic complexity.

4.1. Motivating Example

Brams and Davis (1978, p. 969) argue that, when the parties have juror-inverse preferences, one-shot procedures raise “no strategic questions of timing: given that each side can determine those veniremen [that is, potential jurors] it believes least favorably disposed to its cause, it should challenge these up to the limit of its peremptory challenges.” This quote may be interpreted in different ways, but it suggests that challenging one’s least-preferred jurors is a clear optimal strategy in this case. This may come as a surprise given example 1 and proposition 1. Certainly the one-shot_M procedure is not a dominant-strategy procedure. In particular, challenging one’s least-preferred jurors is not a dominant strategy. How can one then make sense of Brams and Davis’s claim? Example 2 suggests one possible answer.

Example 2. Consider the one-shot_M procedure with $c_D = c_P = 2$ and $b = 5$. Let D have preference $1 R_D \dots R_D 9$. Also suppose that the parties have juror-

¹⁵ That is, the probability that j is chosen given that i plays s_i^j is 0 for all s_i^j .

¹⁶ The last two procedures outlined above also satisfy both properties, at least if the parties are allowed to oppose nominations in the first one. Procedures proposed in Barberà, Massó, and Neme (2005), however, do not satisfy minimal challenge; see Van der Linden (2017).

inverse preferences. If D believes that P is playing a best response to one of her strategies, D knows that P will challenge two of the circled jurors in profile (P3):

R_D :	①	②	③	④	5	6	7	8	9	(P3)
R_P :	9	8	7	6	5	④	③	②	①	

Indeed, a best response by P always involves challenging her two worst jurors among the seven jurors that she believes D will not challenge. Therefore, regardless of the jurors whom P believes D will challenge, a best response by P can never include P challenging a juror in $\{5, \dots, 9\}$. Thus, a best response by D to the minimal belief that P is a best responder always consists in challenging her two worst jurors (boxed in profile [P3]). By symmetry, the same is true for P.

In example 2, the one-shot_M procedure raises no strategic question because a party needs to know only that her opponent is a best responder in order to have a dominant strategy. For each party i , challenging her c_i worst jurors is a best response to any strategy of party $-i$ that is itself a best response to one of i 's strategies. In this sense, each party i has a dominant strategy given a minimal model of the strategic behavior of her opponent: the model $S_{-i} = L_{-i}^1$.

In the rest of this section, I generalize this logic to obtain a measure of strategic complexity. I then apply this measure in Sections 5 and 6 to compare struck procedures for different assumptions on the problem (R_D, R_P, c_D, c_P, b) .

4.2. The Dominance Threshold

As argued above, first-best procedures are procedures in which each party has a dominant strategy no matter what model she has of her opponent. It is then natural to call a procedure second best if each party has a dominant strategy given a minimal model of her opponent. As suggested in example 2, a meaningful concept of a minimal model is for a party to assume that her opponent will play a best response to some of her strategies.

In the language of level- k thinking, a procedure is second best if each party i has an L_i^2 strategy that is a best response to every L_{-i}^1 strategy of her opponent. Such second-best procedures limit the impact of differences in strategic skills because i 's optimal strategy depends minimally on her model of $-i$: i needs to assume only that $-i$ is L_{-i}^1 to have a dominant strategy.

The difference between first-best and second-best procedures is demonstrated in Figure 1, where an arrow from strategy s_i to strategy s_{-i} means that s_i is a best response to s_{-i} . In the first-best procedure represented in Figure 1A, party i has a strategy— s_i^6 —that is a best response to every strategy of her opponent (that is, to every L_{-i}^0 strategy). In the second-best procedure in Figure 1B, party i has a strategy— s_i^4 —that is a best response to every strategy of her opponent that is itself a best response (that is, to every L_{-i}^1 strategy). However, s_i^4 does not need to be a best response to every L_{-i}^0 strategy. For example, in Figure 1B, s_i^4 is not a best response to s_{-i}^1 .




Figure 1. First- (A), second- (B), and third-best (C) procedures

A second-best procedure guarantees a form of second-best equality among equal parties. Consider two defendants with the same preferences who both believe that P is L_p^1 and play a best response to one of her strategies. The two defendants might differ in other strategic aspects, such as their ability to guess which of their strategies P plays a best response to. In a second-best procedure, these differences have no impact: the two defendants play equivalent strategies and secure the same outcome.

Similarly, third-best procedures feature dominant strategies given a model that is minimally stronger than in second-best procedures. A natural candidate for such a minimally stronger model is for i to assume that $-i$ is L_{-i}^2 (see Figure 1C). This logic extends to higher-level reasoning.

In procedures with multiple rounds, it is important to ensure that best responses be enforced throughout the game tree. Therefore, the measure of strategic complexity defined below relies on the iterated elimination of strategies that are never-best responses in any subgame of an extensive game. That is, in each round of elimination, any strategy that fails to be a best response when restricted to any subgame of the game is discarded.

Definition 1: Iterated Elimination of Never-Best Responses. For any procedure Γ and any profile (R_D, R_P) , the process of iterated elimination of never-best responses is defined as follows:

Step 0. For each $i \in \{D, P\}$, the set of L_i^0 (level-0) strategies is S_i .

Step 1. For each $i \in \{D, P\}$, eliminate from L_i^0 the strategies s_i for which there exist a subgame γ of Γ such that the restriction $s_{-i|\gamma}$ of s_i to γ is not a best response to any $s_{-i|\gamma}$ in γ .

The remaining set of strategies is denoted L_i^1 . Any $s_i \in L_i^1$ is called an L_i^1 (level-1) strategy for i .

⋮

Step k . For each $i \in \{D, P\}$, eliminate from L_i^{k-1} the strategies s_i for which

there exist a subgame γ of Γ such that the restriction $s_{i|\gamma}$ of s_i to γ is not a best response to $s_{-i|\gamma}$ for any $s_{-i} \in L_{-i}^{k-1}$.

The remaining set of strategies is denoted L_i^k . Any $s_i \in L_i^k$ is called an L_i^k (level- k) strategy for i .

Observe that the sets of level- k strategies are nested ($L_i^0 \supseteq L_i^1 \supseteq \dots$). Observe also that, for every procedure Γ that satisfies finiteness, the set of level- k strategies is nonempty for every k .¹⁷ The argument at the beginning of this section suggests using the following concept of a dominance threshold as a measure of strategic complexity.

Definition 2: Dominance Threshold. For any procedure Γ and any profile (R_D, R_P) , the dominance threshold is the smallest integer r^* such that, for each $i \in \{D, P\}$, there exists an $L_i^{r^*}$ strategy s_i^* that is a best response to every $L_{-i}^{r^*-1}$ strategy.

If there exists no such integer, then the dominance threshold of Γ is ∞ ; that is, the procedure cannot be solved by iterated elimination of never-best responses. Note that if the dominance threshold r^* is finite, then there exists a strategy profile $(s_i, s_{-i}) \in L_i^{r^*} \times L_{-i}^{r^*}$ that is a subgame-perfect equilibrium.

Throughout this paper, the parties' knowledge of each other's preferences and levels of rationality is left unspecified. The idea behind the dominance threshold is precisely to measure the amount of common knowledge needed for the parties to have dominant strategies. For example, when the dominance threshold of a game is 1, each party has a dominant strategy regardless of her knowledge of her opponent's preferences and level of rationality (the parties need to know only the structure of the game). When the dominance threshold is 2, each party needs to know only her opponent's preferences and the fact that her opponent is a best responder in order to have a dominant strategy.

The related concept of a rationality threshold was introduced by Ho, Camerer, and Weigelt (1998). For a given assumption about the strategies of unsophisticated players, the rationality threshold measures the number of rounds of iterated best responses needed to reach an equilibrium. In contrast, the dominance threshold relies on iterated elimination of never-best responses and does not require a specific assumption about the nature of unsophisticated plays.

The dominance threshold also relates to rationalizability. For a dominance threshold r^* , any strategy in $L_i^{r^*}$ is rationalizable. The set of i 's rationalizable strategies can be larger than $L_i^{r^*}$ if some rationalizable strategies are not best responses to all the strategies of $-i$ that survive the iterated elimination of never-best responses. In addition, the rationality threshold is finite if for every player i , the set of i 's rationalizable strategies contains a strategy that is a best response to all the strategies of $-i$ that survive the iterated elimination of never-best responses. The rationality threshold is infinite otherwise.

¹⁷When Γ satisfies finiteness, for any $i \in \{D, P\}$ and any strategy s_{-i} the set $\{\Gamma(s_i, s_{-i}) \mid s_i \in S_i\}$ is finite because S_i is finite. Hence, there must exist a strategy $t_i(s_{-i})$ such that $\Gamma[t_i(s_{-i}), s_{-i}]R_i\Gamma(s'_i, s_{-i})$ for all $s'_i \in S_i$ and L_i^1 is nonempty. The nonemptiness of L_i^k then follows by induction.

5. One-Shot Procedures

In this section, I show that the one-shot_Q procedure is strategically simpler than the one-shot_M procedure in the following sense:

Proposition 3. (i) For every problem, the dominance threshold of the one-shot_M procedure is no smaller than the dominance threshold of the one-shot_Q procedure. (ii) For some problems, the dominance threshold of the one-shot_M procedure is larger than the dominance threshold of the one-shot_Q procedure.

In the rest of this section, I prove and illustrate proposition 3.

5.1. The One-Shot_Q Procedure Is Always Maximally Simple

Example 3 illustrates how to compute the dominance threshold of the one-shot_Q procedure for a particular problem:

Example 3. This example is represented in Figure 2, where L_T for $T \subseteq N$ represents a lottery in which one juror is drawn at random from T . The labels on the branches of the tree indicate the juror who is challenged in the corresponding action. Suppose that $c_D = c_P = b = 1$ and D is the first party to challenge. Suppose also that the parties have aligned preferences $1 R_D 2 R_D 3$ and $1 R_P 2 R_P 3$. Because preferences for jurors are strict, P has a unique dominant strategy s_P^* , which consists of challenging juror 3 if D did not challenge juror 3 and challenging juror 2 if D did challenge juror 3 (dotted branches). Strategy s_P^* is, therefore, the unique L_P^1 strategy. It directly follows from uniqueness that s_P^* is a best response to all L_D^1 strategies. Because there is a unique L_P^1 strategy s_P^* , any L_D^2 strategy that best responds to s_P^* (either dashed branch in Figure 2) is a best response to all L_P^1 strategies. Hence, the dominance threshold of the one-shot_Q procedure is at most 2 for this problem. But by proposition 1, because the one-shot_M procedure is an N -struck procedure, the dominance threshold of the one-shot_Q procedure is at least 2 for every problem. Thus, the dominance threshold of the one-shot_Q procedure is 2 for this problem.

It is not hard to see how the argument in example 3 generalizes to any problem. In general, the party $-i$ who challenges second in the one-shot_Q procedure has a unique dominant strategy s_{-i}^* . Then any best response by i to s_{-i}^* is a best response to every L_{-i}^1 strategy.

Proposition 4. For any problem, the dominance threshold of the one-shot_Q procedure is 2.

Proposition 4 does not depend on the separability assumption. Instead, the proof relies on the fact that preferences for the outcomes of the procedure are strict. Proposition 4 also extends to situations in which complete information (which is implicit in the definition of a dominance threshold of 2) is relaxed. Consider example 3. To have a dominant strategy, D needs to know only that P will challenge juror 3 if she challenges juror 2. Hence, D needs to know only

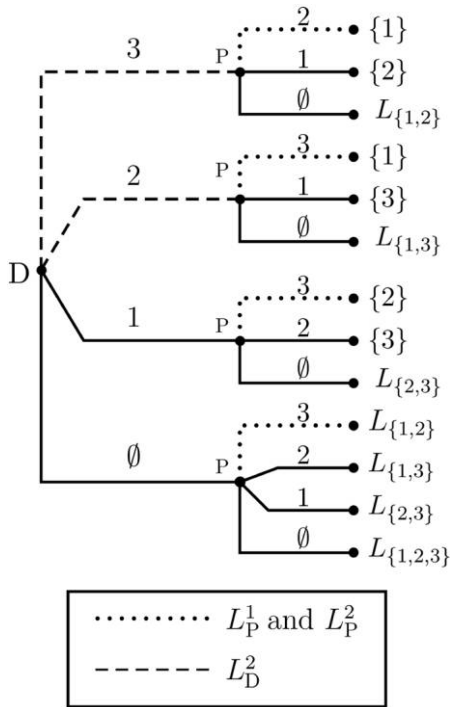


Figure 2. Computing the dominance threshold in example 3

which juror is P's worst juror in order to have a dominant strategy (as opposed to knowing all of P's preferences for jurors; see example 6). See Section 7 for a more detailed discussion of incomplete information.

By proposition 1, because the one-shot_M procedure is an *N*-struck procedure, the dominance threshold of the one-shot_M procedure is at least 2 for every problem. Together with proposition 4, this implies that the dominance threshold of the one-shot_M procedure is never smaller than the dominance threshold of the one-shot_Q procedure, which proves proposition 3.i.

5.2. The One-Shot_M Procedure Is Often Complex: One-Common Profiles

I now show that the one-shot_M procedure is more complex than the one-shot_Q procedure. This is true when the profile is not juror inverse and preferences for jurors satisfy some commonality at the bottom.

5.2.1. Motivating Example

Example 4, Part 1. This example is represented in Table 1. Suppose that $b = c_D = c_P = 1$. Also suppose that the parties' preferences are 1 R_D 2 R_D 3 and

2 R_p 1 R_p 3. Both challenging juror 3 and challenging juror 1 are L_p^0 strategies.¹⁸ Challenging juror 2 is D's best response to P challenging juror 3, and challenging juror 3 is D's best response to P challenging juror 1. Hence, both challenging juror 2 and challenging juror 3 are L_D^1 strategies. But no strategy of P is a best response to both of these L_D^1 strategies. Therefore, the dominance threshold of the one-shot_M procedure is at least 3 for this problem.

In example 4, both parties agree that juror 3 is the worst juror. Therefore, any party playing a best response would challenge juror 3 if her opponent did not. But each party also prefers a situation in which her opponent challenges juror 3 and she challenges her second-worst juror. That is, each party would like to make a credible threat not to challenge juror 3 and free ride on her opponent's challenge of juror 3. But because the procedure is simultaneous, such a credible threat is impossible. As explained in detail in example 4, the impossibility for the parties to commit to leaving juror 3 unchallenged makes the dominance threshold of the one-shot_M procedure larger than 2 for this problem. Together with proposition 4, example 4 therefore proves proposition 3.ii. In fact, the dominance threshold in example 4 is ∞ , which shows just how complex the one-shot_M procedure can become when the profile is not juror inverse.

Example 4, Part 2. Party P's best responses to these two L_D^1 strategies are to challenge juror 3 (P's best response to D challenging juror 2) and to challenge juror 1 (P's best response to D challenging juror 3); see Table 1. Thus, challenging juror 3 and challenging juror 1 are both L_p^2 strategies. But these two L_p^2 strategies are the two L_p^0 strategies considered at the beginning of example 4. The argument therefore extends by induction, which shows that the dominance threshold of the one-shot_M procedure is ∞ for this problem.

As example 4 illustrates, for some profiles that are not juror inverse, the parties' common knowledge of each other's rationality and preferences is not sufficient to provide the parties with dominant strategies and make the game strategically simple. Even for high levels of common knowledge, the game induced by the one-shot_M procedure remains akin to a game of chicken in which each party prefers to swerve (that is, challenge some of her worst jurors) if her opponent stays straight (that is, does not challenge some of her worst jurors) but prefers to stay straight if her opponent swerves.

5.2.2. One-Common Profiles

Profiles for which the dominance threshold of the one-shot_M procedure is ∞ are not rare. Given c_D and c_p , a profile is a one-common profile if a juror w who is among the c_D worst jurors of D is also among the c_p worst jurors of P. Intuitively, the dominance threshold of the one-shot_M procedure is ∞ for one-common profiles because the free-rider problem described in example 4 extends to one-

¹⁸ Challenging juror 2 and not challenging a juror are also L_p^0 strategies. However, it is sufficient to consider the other two L_p^0 strategies.

Table 1
Iterated Best Responses for Juror Challenges in Example 4

	L_p^0 Strategies	L_D^1 Strategies	L_p^2 Strategies	L_D^3 Strategies	...
Strategy 1	3	2	3	2	...
Strategy 2	1	3	1	3	...

common profiles. When the profile is a one-common profile, each party would like to make a credible threat not to challenge juror w and free ride on her opponent's challenge of juror w . But if her opponent does not challenge w , each party prefers to challenge w herself than to leave w unchallenged.

Proposition 5. If the profile is a one-common profile, then the dominance threshold of the one-shot_M procedure is ∞ .

Although proposition 5 relies more directly than proposition 4 on the separability assumption,¹⁹ the intuition behind proposition 5 applies even when separability is relaxed. Regardless of the assumptions about preferences, if for some juror w both parties have best responses that include challenging w , then the dominance threshold is larger than 2 in the one-shot_M procedure. Proposition 5 also extends to situations of incomplete information in which the parties know only that they have a common juror w at the bottom of their ranking of jurors (but do not know each other's complete preferences for jurors).

One-common profiles arise in a number of natural jury selection situations. For example, both parties may dislike a juror whom they view as too unpredictable. Both parties may also dislike devil's advocates or irresolute jurors who are likely to induce a hung jury and to force a retrial of the case. Finally, D may dislike juror j 's position on some charges, while P may dislike juror j 's position on different charges.

Anwar, Bayer, and Hjalmarsson (2014) provide suggestive evidence of overlapping preferences in jury selection.²⁰ In the Online Appendix, I show that one-common profiles are frequent in a different, yet similar, problem: the selection of an arbitrator between unions and employers by the New Jersey Public Employment Relations Commission from 1985 to 1996 (Bloom and Cavanagh 1986; de Clippel, Eliaz, and Knight 2014). I also show that one-common profiles represent a significant proportion of the set of profiles. This is true even when attention is limited to profiles that are close to being juror inverse (in a sense that is made precise in the Online Appendix).

In the Online Appendix, the proportion of one-common profiles is shown to be an increasing function of the number of challenges and a decreasing function of the number of jurors b . In accordance with the objective of reducing strategic complexity, proposition 5 and the results in the Online Appendix therefore

¹⁹ One-common profiles are not well defined without the separability assumption.

²⁰ In 700 felony cases in Florida, Anwar, Bayer, and Hjalmarsson (2014) find that both defendants and plaintiffs are less likely to challenge black jurors. They suggest that this may be due to the parties' shared aversion to being accused of violating *Batson v. Kentucky*.

provide a game-theoretic justification for decreasing the number of peremptory challenges, a measure that has some support among those who defend a reform of the peremptory challenge system (Henley 1996). Procedures in which the number of challenges is high relative to b exist in practice. In the United States, the number of challenges tends to increase with the gravity of the charges. For example, in federal cases for which the death penalty is sought by the prosecution, $b = 12$ and $c_D = c_P = 20$. In this case, the dominance threshold is ∞ for more than 97 percent of the profiles (and more than 15 percent of the profiles that are close to being juror inverse).

Overall, the results in this section contrast with judges' beliefs that blind (that is, simultaneous) procedures leave less room for the parties to strategize than sequential ones (see Shapard and Johnson 1994, p. 6). Contrary to judges' beliefs, the dominance thresholds suggest that the one-shot_Q procedure is strategically simpler than the one-shot_M procedure: by making past actions observable, the one-shot_Q procedure allows the parties to make credible threats about the jurors they challenge, which reduces the amount of guesswork involved in determining an appropriate strategy. Section 6 shows that similar results hold for other N -struck procedures.

6. Alternating and Other N -Struck Procedures

In general, it is unclear how the alternating_M and alternating_Q procedures compare. However, extending the logic of proposition 5, I show that it is possible to obtain a partial comparison for a significant subset of profiles. For this subset of profiles, the dominance threshold of any simultaneous N -struck procedure (including alternating_M) is infinite, whereas the dominance threshold of any sequential N -struck procedure (including alternating_Q) is finite.

If preferences for the outcomes of a sequential N -struck procedure are strict (including preferences regarding lotteries), then the procedure always has a finite dominance threshold. This follows from the fact that, with strict preferences on the set of outcomes, sequential N -struck procedures induce games of perfect information that can be uniquely solved by backward induction.²¹ Then the number of rounds of backward induction required to solve the game is an upper bound for the dominance threshold.

Proposition 6. For any sequential N -struck procedures, if preferences for the outcomes of the procedure are strict, then the dominance threshold is finite and smaller than the depth of the game tree.²²

Again, proposition 6 does not depend on the separability assumption but instead on the assumption that preferences for the outcomes of the procedure are strict.

Recall that the one-shot_M procedure has an infinite dominance threshold when

²¹ More precisely, multiple strategy profiles can survive backward induction, but each of them must yield the same outcome.

²² The depth of a game tree is the length of the longest path from the initial node to a terminal node.

the profile is a one-common profile because each party would like to free ride on her opponent’s challenge of one of the jurors they both dislike (see example 5). This idea generalizes to the class of simultaneous N -struck procedures as a whole. Below I identify for each simultaneous N -struck procedure Γ a set of Γ -one-common profiles. In proposition 7, I show that any Γ -one-common profile induces an infinite dominance threshold in Γ .

Informally, given a simultaneous N -struck procedure Γ , a profile is a Γ -one-common profile if in one of the final subgames of Γ , the set of jurors who remain unchallenged gives rise to the free-rider problem described above. Formally, given Γ , a profile is a Γ -one-common profile if there exists a subgame γ of Γ such that (a) both parties can still challenge jurors in γ (that is, $\ell_i^\gamma \geq 1$ for both $i \in \{D, P\}$), (b) the first round of γ is the final round of γ in which both parties can challenge jurors,²³ and (c) among the unchallenged jurors, one of the ℓ_D^γ worst jurors according to R_D is also one of the ℓ_P^γ worst jurors according to R_P .

Example 5, Part 1. Consider the alternating_M procedure and any problem in which $c_D = c_P = 2$, $b = 1$, and the preferences for jurors are as displayed as follows:

$$\begin{array}{rcccccc} R_D: & 1 & 2 & 3 & 4 & 5 \\ R_P: & 2 & 4 & 5 & 1 & 3 \end{array} \tag{P4}$$

The profile is not a one-common profile because $\{4, 5\} \cap \{1, 3\} = \emptyset$. However, consider the subgame γ^* that follows from D challenging juror 4 and P challenging juror 5 in the first round. Subgame γ^* satisfies conditions a and b in the definition of an alternating_M-one-common profile. In addition, both players have the same worst juror among $\{1, 2, 3\}$, the set of unchallenged jurors at the beginning of γ^* . Hence, condition c is also satisfied, and so this is an alternating_M-one-common profile. Consider one-shot_M-one-common profiles. The only subgame of the one-shot_M procedure is the one-shot_M procedure itself. Hence, in the case of the one-shot_M procedure, conditions a , b , and c boil down to requiring that among N , one of the c_D worst jurors according to R_D is also one of the c_P worst jurors according to R_P , which is the definition of a one-common profile. Because the sets of one-shot_M-one-common and one-common profiles are identical, proposition 7 generalizes proposition 5.

Proposition 7. For any simultaneous N -struck procedure Γ , if the profile is a Γ -one-common profile, then the dominance threshold of Γ is ∞ .

Example 5, Part 2. To see why the dominance threshold is infinite in subgame γ^* , observe that in γ^* , each party wants to free ride on her opponent’s challenge of juror 3. This induces an infinite dominance threshold for the same reasons that the dominance threshold is infinite in example 4.

Propositions 6 and 7 jointly imply that, whenever the profile is an alternating_M-

²³ This could arise because the first round of γ is the terminal round of Γ or because both parties have only one challenge left in γ .

one-common profile (and preferences on outcomes are strict), the dominance threshold of the alternating_Q procedure is smaller than the dominance threshold of the alternating_M procedure. That is, proposition 3 partially extends to alternating procedures. In the Online Appendix, I show that the alternating_M-one-common profiles are a strict superset of the one-common profiles. Hence, the arguments about the prevalence of one-common profiles in Section 5.2 extend to alternating_M-one-common profiles.

7. Extension: Incomplete Information

Because the lowest dominance threshold for reasonable challenge procedures is 2 (proposition 1), the above comparisons implicitly assume that the parties know each other's preferences. As these comparisons show, once preferences are known, some procedures lead to simpler strategic interplays than others.

It is conceivable for the parties to know each other's preferences, especially if their attorneys repeatedly interact with one another. Typically, the parties also share information about the jurors (for example, both parties witness the jurors' answers to all parties' questions), which may help them in forming accurate models of each other's preferences. However, in a number of cases, a party may have only incomplete information about her opponent's preferences. In such cases, dominance thresholds larger than 1 provide little information about the strategic complexity of a mechanism, and the above comparisons are less meaningful.

To deal with incomplete information, one must first model the parties' strategic choices when information about preferences is incomplete but fixed. By "fixed," I mean a situation in which parties do not update their beliefs about the other party's preferences on the basis of the history of the game. In a second step, it may be important to consider the possibility of such updates. In this section, I provide some tools and formal results related to the first step and discuss the second step more succinctly and informally.

A natural extension of the dominance threshold to situations of incomplete information consists in assuming that the parties have set beliefs about one another's preferences rather than (accurate) point beliefs. That is, instead of knowing $-i$'s preference, party i knows only that $-i$'s preference belongs to some subset $\mathcal{S}_{-i}^i \subseteq \mathcal{S}_{-i}$ containing $-i$'s true preference (that is, $R_{-i} \in \mathcal{S}_{-i}^i$). When updating of beliefs is not considered, these sets remain unchanged throughout the game.

Subset \mathcal{S}_{-i}^i represents i 's first-order beliefs about $-i$'s preferences. A general extension of the dominance threshold would require a complete specification of higher-order beliefs as well.²⁴ To avoid such intricacies, it is useful to assume that

²⁴ For example, identifying the set of i 's level-2 models of $-i$ requires knowing how i believes $-i$ would play a best response if $-i$ assumes that i has a level-1 strategy. To do so in general, it is not enough to specify i 's first-order beliefs \mathcal{S}_{-i}^i about $-i$'s preferences. One also needs to determine i 's second-order beliefs about $-i$'s first-order beliefs, say $\mathcal{S}_i^{-i,i}$. Party i then expects $-i$ to model i 's level-1 behavior on the basis of some preference in $\mathcal{S}_i^{-i,i}$. Party i can therefore form a level-2 model of $-i$ playing a best response according to some preference in \mathcal{S}_{-i}^i to a level-1 model of i that is based on some preference in $\mathcal{S}_i^{-i,i}$.

beliefs are consistent and common knowledge. That is, although each party does not know her opponent's preferences, she knows her opponent's best guess about her own preferences. Each also knows that her opponent knows her guess about her opponent's preferences, and so on. Formally, party i believes that $-i$'s preference is some $R_{-i}^i \in \mathcal{S}_{-i}^i$, and $-i$ knows that i believes her preference is R_{-i}^i . In turn, i knows that $-i$ knows that i believes her preference is R_{-i}^i , and so on.

Although simple, this incomplete-information setting departs in a nontrivial way from the complete-information situation studied in the previous sections. In particular, the parties' beliefs $(R_D^P, R_P^D) \in \mathcal{S}_D^P \times \mathcal{S}_P^D$ can differ from P and D's true preferences R_D and R_P . Complete information is of course a special case of this more general model obtained by setting $\mathcal{S}_D^D = \{R_D\}$ and $\mathcal{S}_P^P = \{R_P\}$.

If beliefs are consistent and common knowledge, the iterated elimination of never-best responses unfolds similarly as with complete information but based on some $(R_D^P, R_P^D) \in \mathcal{S}_D^P \times \mathcal{S}_P^D$ possibly different from (R_D, R_P) . One can then conveniently generalize the dominance threshold to make it a function of first-order beliefs only. Leave level-0 strategies unchanged. For any procedure Γ , any first-order beliefs $(\mathcal{S}_D^P, \mathcal{S}_P^D)$, any $i \in \{D, P\}$, and any $k \in \mathbb{N} \setminus \{0\}$, the set of i 's $(\mathcal{S}_D^P, \mathcal{S}_P^D)$ level- k strategies is defined as the union of i 's level- k strategies in all games (Γ, R_D^P, R_P^D) for which $(R_D^P, R_P^D) \in \mathcal{S}_D^P \times \mathcal{S}_P^D$. Observe that the set of $(\mathcal{S}_D^P, \mathcal{S}_P^D)$ level- k strategies of i then corresponds to $-i$'s model of i 's level- k strategies. This set contains—but need not be identical to—the set of i 's actual best responses to $-i$'s $(\mathcal{S}_D^P, \mathcal{S}_P^D)$ level- $(k - 1)$ strategies (where “actual” means according to i 's true preference $R_i \in \mathcal{S}_i^{-i}$).

The $(\mathcal{S}_D^P, \mathcal{S}_P^D)$ dominance threshold of a procedure Γ is the smallest integer r^* such that both players have a strategy that is a best response (according to their true preferences) to every $(\mathcal{S}_D^P, \mathcal{S}_P^D)$ level- $(r^* - 1)$ strategy of their opponent. (Again, the $(\mathcal{S}_D^P, \mathcal{S}_P^D)$ dominance threshold is ∞ if there is no such r^* .) When the threshold is finite, the profile of these best responses is again a subgame-perfect equilibrium of game (Γ, R_D, R_P) because the true profile (R_D, R_P) is included in $\mathcal{S}_D^P \times \mathcal{S}_P^D$.²⁵ For the same reason, the $(\mathcal{S}_D^P, \mathcal{S}_P^D)$ dominance threshold of a procedure is always no smaller than the procedure's (complete-information) dominance threshold.

Example 6. Consider the incomplete-information situation discussed after proposition 4 and based on example 3, which features $b = c_D = c_P = 1$. Suppose that the true preferences are again aligned with 1 R_D 2 R_D 3 and 1 R_P 2 R_P 3 and that D is the first party to challenge in the one-shot_Q procedure. However, unlike under complete information, D's beliefs about P's preference \mathcal{S}_P^D contain both R_P and \hat{R}_P with 2 \hat{R}_P 1 \hat{R}_P 3 (that is, D knows only that juror 3 is P's least-preferred juror). Party P's beliefs \mathcal{S}_D^P can be any subset of \mathcal{S}_D . As in the complete-information case, P has a unique best response to any strategy of D, whereas D's

²⁵ However, the same strategy profile does not need to be a subgame-perfect equilibrium of game (Γ, R_D', R_P') when $(R_D, R_P) \neq (R_D', R_P')$ because in the definition of the $(\mathcal{S}_D^P, \mathcal{S}_P^D)$ dominance threshold, the parties' best responses to all $(\mathcal{S}_D^P, \mathcal{S}_P^D)$ level- $(r^* - 1)$ strategies of their opponent are in terms of their true preferences (R_D, R_P) , and not in terms of (R_D', R_P') .

best response depends on P’s strategy. If D assumes that P has a level-1 strategy, then D anticipates that P will play either of two strategies depending on whether P’s preference is R_p or \hat{R}_p . Yet challenging juror 2 is a best response for D to both of these level-1 strategies of P.²⁶ As a consequence, the $(\hat{\mathcal{S}}_D^p, \hat{\mathcal{S}}_p^D)$ dominance threshold of the one-shot_Q procedure is 2. Observe that challenging juror 2 is the only such strategy for D. Because $\hat{\mathcal{S}}_p^D$ contains preference 2 R_p 1 R_p 3, D considers the possibility that if she acted as if she had a level-1 strategy, P would challenge juror 1 following her challenge of juror 3 (in which case, she would have been better off challenging juror 2, to which P would have responded by challenging juror 3 herself). In contrast, in the complete-information case, both challenging juror 2 and challenging juror 3 are best responses to L_p^1 strategies for D.

For beliefs different from those presented in example 6 (for example, if preference 3 \hat{R}_p 2 \hat{R}_p 1 is added to $\hat{\mathcal{S}}_p^D$; see also example 7), the one-shot_Q procedure can have a larger, even infinite $(\hat{\mathcal{S}}_D^p, \hat{\mathcal{S}}_p^D)$ dominance threshold. This is in spite of the fact that every individual game (one-shot_Q, (R_D^p, R_p^D)) with $(R_D^p, R_p^D) \in \hat{\mathcal{S}}_D^p \times \hat{\mathcal{S}}_p^D$ itself has a dominance threshold of 2 (with respect to preferences (R_p^p, R_p^D)), which shows that the $(\hat{\mathcal{S}}_D^p, \hat{\mathcal{S}}_p^D)$ dominance threshold is a non-trivial generalization of the dominance threshold $\hat{\mathcal{S}}_p^D$. Proposition 5 however generalizes in the sense that the $(\mathcal{S}_D^p, \mathcal{S}_p^D)$ dominance threshold of the one-shot_Q procedure is always no larger (and sometimes smaller) than that of the one-shot_M procedure.

Proposition 8. (i) For every problem and every belief structure $(\mathcal{S}_D^p, \mathcal{S}_p^D)$, the $(\hat{\mathcal{S}}_D^p, \hat{\mathcal{S}}_p^D)$ dominance threshold of the one-shot_M procedure is no smaller than the $(\mathcal{S}_D^p, \mathcal{S}_p^D)$ dominance threshold of the one-shot_Q procedure. (ii) For some problems and some belief structures, the $(\hat{\mathcal{S}}_D^p, \hat{\mathcal{S}}_p^D)$ dominance threshold of the one-shot_M procedure is larger than the $(\mathcal{S}_D^p, \mathcal{S}_p^D)$ dominance threshold of the one-shot_Q procedure.

To what extent the comparison of multiple-rounds procedures in Section 6 also generalizes to incomplete information using the $(\mathcal{S}_D^p, \mathcal{S}_p^D)$ dominance threshold is an open question.

As a transition toward a discussion of updating beliefs, it is interesting to consider the new asymmetry that emerges in the one-shot_Q procedure once incomplete information is taken into account. The one-shot_Q procedure is obviously asymmetric with or without incomplete information because it makes one party challenge before the other. Interestingly, although this asymmetry may be viewed as procedurally unfair compared with the symmetric one-shot_M procedure, judges and state legislators seem to leverage such asymmetries rather than try to alleviate them. In particular, instead of randomizing the order in which the parties challenge in sequential procedures, rules of legal procedures often select a deterministic first mover.²⁷ Judges and state legislators seem to try to take advan-

²⁶ Observe that this would remain true if 1 \hat{R}_p 3 \hat{R}_p 2 was further added to D’s belief set $\hat{\mathcal{S}}_p^D$.

²⁷ For example, New Jersey has the following order: “[T]he State in a criminal case and the plaintiff in a civil case exercis[e] the first challenge” (N.J. R. 1:8-3[e][1]).

tage of a procedure's asymmetries to favor a certain party by giving it the role in the procedure they perceive to be most favorable.

Under complete information, there is a known advantage to being the first to challenge in the alternating_Q procedure, provided that preferences satisfy a mild regularity condition defined in DeGroot and Kadane (1980). The same is true in general in the one-shot_Q procedure. Intuitively, the party who moves first can take advantage of her accurate prediction of her opponent's behavior in the second stage to let her opponent challenge jurors in her place and use some of her challenges optimally on other jurors.

Proposition 9. (i) For every problem, any party weakly prefers the subgame-perfect equilibrium of the one-shot_Q procedure in which she challenges first to the subgame-perfect equilibrium of the one-shot_Q procedure in which she challenges second. (ii) There exist problems for which this preference is strict.

This is not necessarily the case, however, under incomplete information, which introduces a counterbalancing asymmetry between the parties. If the party who moves first plays a best response to a level-1 model of her opponent that is based on the wrong preference, she might have been better off challenging second instead, knowing the set of jurors that the other party first challenged.

Example 7. Consider again the case $b = c_D = c_P = 1$ and suppose that D has preference $1 R_D 2 R_D 3$ and that P has preference set $3 R_P 2 R_P 1$. Whereas P knows D's preferences and $\tilde{S}_D^P = \{R_D\}$, party D is unsure of P's preferences, and \tilde{S}_P^D contains both R_P and \tilde{R}_P defined by $2 \tilde{R}_P 1 \tilde{R}_P 3$. In this case, there is no strategy for D that is a best response to both level-1 models of P based on R_P and \tilde{R}_P . If D challenges first in the one-shot_Q procedure and decides to play a best response to a level-1 model of P based on \tilde{R}_P , D challenges juror 2 hoping that P will then challenge juror 3. But because P's true preference is R_P , following P's challenge of juror 2, P in fact challenges juror 1, which leaves juror 3 as the effective juror. If instead juror P had challenged first in the one-shot_Q procedure, she would have challenged juror 1 knowing that D would then challenge juror 3. This would have resulted in juror 2 becoming the effective juror, an outcome that D prefers to the outcome of the one-shot_Q procedure in which she challenges first on the basis of an erroneous guess about P's preference.

Observe how the possible advantage of moving second in the one-shot_Q procedure relates to the asymmetry in strategic complexity between the parties. The one-shot_Q procedure is strategically more complex for the first party to challenge than for the second. If information is complete enough to predict the behavior of the second party to challenge sufficiently accurately (as in example 6), the first party can take advantage of this strategic complexity to secure a better jury. If information is too incomplete and predictions of the second party's behavior are too fuzzy (as in example 7), the first party might prefer the strategically simpler position of being the second to challenge.

As example 7 illustrates, the usefulness of moving second materializes even in the absence of updating beliefs in the one-shot_Q procedure. In a sense, after the second party observes the first party's behavior, it is irrelevant whether the second party updates her beliefs about the first party's preferences. In multiple-rounds procedures, the disadvantage of moving first can be further reinforced if the parties effectively use past actions to update their beliefs about one another's preferences.²⁸

Consider, for example, the sequential variant of the two-shots procedure described in note 9. In this procedure, the second party to challenge has the opportunity to update her beliefs about the first party's preferences using the first party's action in the first round. This, in turn, allows the second party to better anticipate the first party's action in the second round. This imbalance is somewhat compensated by the fact that the first party then also learns from the second party's action in the first round before choosing her action in the second round. But in sequential procedures in which the order of challenges never changes, the second mover always has one more opportunity to learn from the first mover's action and react to it.²⁹

Informally, it seems that introducing updating beliefs could have two counterbalancing effects on the comparison of sequential and simultaneous procedures. On the one hand, having more accurate beliefs about an opponent's preferences makes a game simpler. This could favor sequential procedures because they give more opportunities for the parties to observe each other's actions and update their beliefs. On the other hand, the parties would then have to worry about the information they reveal through their actions, which can make a game more complex. This could favor simultaneous procedures if simultaneity means that the parties less often have to worry about the information their actions reveal.

How these two effects play out (if at all) and whether moving first really harms the first party by giving her fewer opportunities to react to updated beliefs are open questions that would require a full-fledged model of the parties' updating process. Such a model is beyond the scope of the present paper, but better understanding incomplete information and updating of beliefs in selection procedures (beyond propositions 8 and 9) are interesting directions for future research.

8. Conclusion

This paper shows how jury selection procedures can be compared in terms of their strategic complexity by computing their dominance thresholds, that is, the required number of rounds to eliminate strategies that are never-best responses

²⁸ It may be in this sense that judges view blind procedures as less prone to strategizing than alternating ones.

²⁹ Recall that asymmetries such as this one can be useful if officials find it desirable to favor one party over the other. If, however, symmetry is viewed as desirable, it may be beneficial to alternate the order of challenges (with one party challenging first, the second party challenging twice, followed by the first party challenging twice again, and so on), as has been proposed, for example, to make penalty shootouts fairer (Brams and Ismail 2018).




Figure 3. Arbitrary hierarchical model for a given preference profile

for the parties to have a dominant strategy. The dominance threshold offers a new method to compare the strategic complexity of mechanisms. Unlike previous methods in the literature (Pathak and Sönmez 2013; de Clippel, Eliaz, and Knight 2014; Arribillaga and Massó 2016), it allows for comparisons even when the mechanisms at stake are indirect or induce games of imperfect information.

More generally, the dominance threshold shows how hierarchical models can be used to compare the strategic complexity of mechanisms. As shown in Figure 3, for any profile (R_D, R_P) , a hierarchical model specifies a pair $(\{S_D^0, \dots, S_D^m\}, \{S_P^0, \dots, S_P^m\})$ of collections of nested strategy sets; that is, $S_i \subseteq S_i^0 \subseteq \dots \subseteq S_i^m$ (m could be infinite). As k increases, the sets S_i^k represent increasingly restrictive models of the strategies that i could potentially play.

This paper examines the level- k hierarchical model $(\{L_D^0, \dots, L_D^m\}, \{L_P^0, \dots, L_P^m\})$ defined in Section 4. Given a profile (R_D, R_P) , I define the dominance threshold as the smallest hierarchical level r^* for which each party i has a strategy $s_i^* \in S_i$ that is a best response to every strategy in L_{-i}^{r-1} . I then use the dominance threshold as a measure of strategy complexity.

Clearly, this logic is not specific to the level- k hierarchical model. A natural alternative would be to use the undominated hierarchical model UD defined by the process of iterated elimination of dominated strategies. One could then define an alternative UD dominance threshold. In general, there is no logical relation between the UD dominance threshold and the level- k dominance threshold of a game. However, some of the results in this paper also apply when the UD dominance threshold is used instead.

First, it is not hard to see that the UD dominance threshold of the one-shot_Q procedure is 2. Second, it can be shown that for every problem, the UD dominance threshold of the one-shot_M procedure is at least as large as the level- k dom-

inance threshold of the one-shot_M procedure.³⁰ Hence, the results in Section 5.2 also apply using the UD dominance threshold. In particular, the UD dominance threshold of the one-shot_M procedure is larger than 2 for a significant set of problems. Whether the results of Section 6 also extend to the case of a UD dominance threshold is left as an open question.

References

- Abreu, Dilip, and Hitoshi Matsushima. 1992. Virtual Implementation in Iteratively Undominated Strategies: Complete Information. *Econometrica* 60:993–1008.
- Alpern, Steve, and Shmuel Gal. 2009. Analysis and Design of Selection Committees: A Game Theoretic Secretary Problem. *International Journal of Game Theory* 38:377–94.
- Alpern, Steve, Shmuel Gal, and Eilon Solan. 2010. A Sequential Selection Game with Vetoes. *Games and Economic Behavior* 68:1–14.
- Anwar, Shamena, Patrick Bayer, and Randi Hjalmarrsson. 2014. The Role of Age in Jury Selection and Trial Outcomes. *Journal of Law and Economics* 57:1001–30.
- Arribillaga, R. Pablo, and Jordi Massó. 2016. Comparing Generalized Median Voter Schemes according to Their Manipulability. *Theoretical Economics* 11:547–86.
- Barberà, Salvador, and Danilo Coelho. 2008. How to Choose a Non-controversial List with k Names. *Social Choice and Welfare* 31:79–96.
- Barberà, Salvador, Jordi Massó, and Alejandro Neme. 2005. Voting by Committees under Constraints. *Journal of Economic Theory* 122:185–205.
- Barberà, Salvador, Hugo Sonnenschein, and Lin Zhou. 1991. Voting by Committees. *Econometrica* 59:595–609.
- Bermant, Gordon. 1982. *Jury Selection Procedures in United States District Courts*. Washington, DC: Federal Judicial Center.
- Bermant, Gordon, and John Shapard. 1981. The Voir Dire Examination, Juror Challenges, and Adversary Advocacy. Pp. 69–114 in *The Trial Process*, edited by Bruce Dennis Sales. New York: Plenum Press.
- Bloom, David E., and Christopher L. Cavanagh. 1986. An Analysis of the Selection of Arbitrators. *American Economic Review* 76:408–22.
- Börger, Tilman, and Jiangtao Li. 2017. Strategically Simple Mechanisms. Unpublished manuscript. University of Michigan, Department of Economics, Ann Arbor.
- Brams, Steven J., and Morton D. Davis. 1978. Optimal Jury Selection: A Game-Theoretic Model for the Exercise of Peremptory Challenges. *Operations Research* 26:966–91.
- Brams, Steven J., and Mehmet S. Ismail. 2018. Making the Rules of Sports Fairer. *SIAM Review* 60:181–202.
- Caditz, David M. 2014. On the Application of Game Theory in Jury Selection. *Jury Expert*, August 20. <http://www.thejuryexpert.com/2014/08/on-the-application-of-game-theory-in-jury-selection/>.
- . 2015. Selection under Veto with Limited Foresight. Unpublished manuscript. December. Sausalito, CA.
- Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberry. 2013. Structural

³⁰ By the assumption that preferences for jurors are strict, best responses are unique in the one-shot_M procedure. Thus, every best response is also an undominated strategy, and the set of strategies that survive k rounds of iterated elimination of never-best responses is a subset of the set of strategies that survive k rounds of iterated undominated strategies.

- Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications. *Journal of Economic Literature* 51:5–62.
- Daly, Meghan. 2016. *Foster v. Chatman*: Clarifying the *Batson* Test for Discriminatory Peremptory Strikes. *Duke Journal of Constitutional Law and Public Policy Sidebar* 11:149–62.
- de Clippel, Geoffroy, Kfir Eliaz, and Brian Knight. 2014. On the Selection of Arbitrators. *American Economic Review* 104:3434–58.
- DeGroot, Morris H., and Joseph B. Kadane. 1980. Optimal Challenges for Selection. *Operations Research* 28:952–68.
- Flanagan, Francis X. 2015. Peremptory Challenges and Jury Selection. *Journal of Law and Economics* 58:385–416.
- Futterman, Roy. 2014. Comment on David Caditz’s “On the Application of Game Theory in Jury Selection.” *Jury Expert*, August 20. <http://www.thejuryexpert.com/2014/08/on-the-application-of-game-theory-in-jury-selection/>.
- Henley, Patricia. 1996. Improving the Jury System: Peremptory Challenges. UC Hastings College of the Law, Public Law Research Institute, San Francisco.
- Ho, Teck-Hua, Colin Camerer, and Keith Weigelt. 1998. Iterated Dominance and Iterated Best Response in Experimental “*p*-Beauty Contests.” *American Economic Review* 88:947–69.
- Kadane, Joseph B., Christopher A. Stone, and Garrick Wallstrom. 1999. The Donation Paradox for Peremptory Challenges. *Theory and Decision* 47:139–55.
- LaFave, Wayne R., Jerold H. Israel, Nancy J. King, and Orin S. Kerr. 2009. *Criminal Procedure*. 5th ed. St. Paul, MN: West Academic.
- Leibold, Jill. 2015. Who Is the Ideal Juror to Look for during Voir Dire? *Jury Expert*, February. <http://www.thejuryexpert.com/2015/02/who-is-the-ideal-juror-to-look-for-during-voir-dire/>.
- Li, Shengwu. 2015. Obviously Strategy-Proof Mechanisms. Working paper. Stanford University, Department of Economics, Stanford, CA.
- Marder, Nancy S. 2012. *Batson* Revisited. *Iowa Law Review* 97:1585–1612.
- Pathak, Parag A., and Tayfun Sönmez. 2013. School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation. *American Economic Review* 103:80–106.
- Roth, Arthur, Joseph B. Kadane, and Morris H. DeGroot. 1977. Optimal Peremptory Challenges in Trials by Juries: A Bilateral Sequential Process. *Operations Research* 25:901–19.
- Shapard, John, and Molly Johnson. 1994. Memorandum on Survey concerning Voir Dire to the Advisory Committee on Civil Rules and the Advisory Committee on Criminal Rules. Federal Judicial Center, Research Division, Washington, DC.
- Van der Linden, Martin. 2017. Impossibilities for Strategy-Proof Committee Selection Mechanisms with Vetoers. *Journal of Mathematical Economics* 73:111–21.