Advance Access publication June 8, 2005

Political Analysis (2006) 14:101–120
doi:10.1093/pan/mpi017

Comparing Models of Strategic Choice: The Role of Uncertainty and Signaling

Jonathan Wand
Department of Political Science, Encina Hall West,
Stanford University, Stanford, CA 94305
e-mail: wand@stanford.edu

Testing the fit of competing equilibrium solutions to extensive form games crucially depends on assumptions about the distribution of player types. To illustrate the importance of these assumptions for differentiating standard statistical models of strategic choice, I draw on a game previously analyzed by Lewis and Schultz (2003). The differences that they highlight between a pair of perfect Bayesian equilibrium and quantal response equilibrium models are not produced by signaling and updating dynamics as claimed, but are instead produced by different assumptions about the distribution of player types. The method of analysis developed and the issues raised are applicable to a broad range of structural models of conflict and bargaining.

1 Introduction

Numerous statistical models have recently been proposed that are derived from equilibrium solutions to noncooperative extensive form games, each positing a different theory of information and uncertainty about player types. The quantal response equilibrium (QRE) model is based on individuals (or their "agents") making random perceptual errors each time they face a decision, and the model was originally applied to experiments where participants knew the fixed payoffs to the game but nonetheless made stochastic choices (McKelvey and Palfrey 1998). Quinn and Westveld (2004) have provided a method for relaxing parametric assumptions about the distribution of errors in QRE models. Signorino (2003) proposes a pair of models. The first is an alternative agent-like theory where individuals do not make perceptual errors, but rather information is revealed to individuals when they face a decision. The second is a regressor error model based on a game where players know each others' type, but limited information about player types is available to a researcher observing the game. Lewis and Schultz (2003) allow for each player in a game to have private information about their own type and derive a perfect Bayesian equilibrium (PBE) model. Determining and understanding observable implications of these models is essential to comparing and testing the competing theories.

Author's note: I thank Jas Sekhon and Alex Tahk for providing extensive comments. I also thank Robert Anderson, Tim Büthe, Alberto Diaz-Cayeros, Jim Fearon, Simon Jackman, Paul Sniderman, and Mike Tomz for helpful comments, and the PhD students of my "Mathematical and Statistical Models of Choice" course in the spring of 2004 who went through earlier drafts. I appreciate Ken Schultz and Jeff Lewis reading and commenting on the paper. I am solely responsible for errors and omissions.

I analyze the effects of uncertainty over player types and the effects of signaling in determining substantively interesting features of strategic choice models for a standard crisis bargaining game. Expanding upon previous analysis of QRE and PBE models for this extensive form game, I provide new insights and correct past errors regarding differences in key implications of these models. The following theoretical analysis of the mechanisms that drive the complex and substantively interesting features of these models highlights the need for subject-specific empirical research on the distribution of player types to supplement and guide the implementation of these types of structural models.

This work extends existing analysis of the substantive implications of the distribution of player types. The QRE literature grounded in experimental data explicitly addresses the effect of stochastic variability on comparative statics (McKelvey and Palfrey 1998). Indeed, in the application of the QRE to experiments where the payoffs are fixed and known by the players and researchers, it is the variability of perceptual errors of the players that is estimated. In contrast to an analysis of experimental data, the following results relate to studies where average utilities would be estimated using observational data. The logic of my analysis also pertains to applications with experimental data, highlighting the relevance of testing the common restriction of setting the variance of perceptual errors to be equal at all decision nodes.[1] Signorino (2003) considered the effects of stochastic variability in the QRE and two other strategic models, providing examples of the bias of outcome probabilities and associated parameter estimates under misspecification. This article provides an analysis of the mechanisms underlying differences in observable implications of the models. Most other theoretical investigations and empirical applications employ a single strategic choice model and only make comparisons with nonstrategic models (e.g., Signorino and Yilmaz 2003; Carson 2003, 2005; Leblang 2005).

In their comparison of QRE- and PBE-based models for a crisis bargaining game, Lewis and Schultz (2003) argue that the differences between the two models follow from the signaling and updating dynamics in the PBE model. In particular, the authors make two strong claims about substantive differences between the QRE and PBE solutions. First, they claim that a decrease in the cost of War for a (potential) challenger would lead a researcher using a QRE solution to predict a monotonic increase in the probability of war, while a researcher using a PBE solution would predict a nonmonotonic (first increasing, then decreasing) relationship (Lewis and Schultz 2003, p. 353). Second, they claim that an increased rate of defenders conceding, rather than resisting, an initial challenge would be interpreted in the context of a PBE solution as following from the challenge being more credible, as measured by an increase in audience costs relative to the expected cost of war. They argue that a QRE solution, in contrast, would interpret the same increase in conceding as reflecting a joint improvement in the value of war and audience costs, without changing the relative value of each (Lewis and Schultz 2003, p. 358). The authors argue that the features of the PBE solution premised on asymmetric information, but not the QRE, are consistent with the existing literature and the logic of costly signaling.

However, the differences between the QRE and PBE that have been highlighted are not produced by differences in signaling and updating dynamics. Indeed there is negligible learning by the players in the context of the previous analysis of the PBE. Rather, the differences are due to the fact that the PBE was derived using one distribution of player types and the QRE was derived using another distribution of types. The conclusion that "although both models incorporate uncertainty, it is clear that whether or not the solution

[1] I conjecture that perceptual errors are generally greater at earlier nodes of complex games.

concept captures signaling and updating dynamics has fundamentally important implications for conclusions that would be drawn from the same set of observational data" (Lewis and Schultz 2003, p. 362) is not supported by the evidence the authors present. I also correct other erroneous explanations for some properties of the QRE and PBE models previously examined.

Although the particular illustration that follows uses terminology for actions and outcomes characterizing a stylized military conflict, the insights are relevant to a large class of conflict and bargaining situations. The political science literature applying similar structural models of strategic choice includes studies of candidate entry in electoral competition (e.g., Aragones and Palfrey 2004; Carson 2003, 2005), exchange rate policy (Leblang 2005), and of course interstate conflict (e.g., Signorino 1999). The particular game that I consider could also be used to describe numerous other situations, including litigation, lobbying for legislation, and negotiating with your significant other over what to do on a Friday night. For example, Robert Anderson has proposed the following litigation game: (a) the plaintiff chooses whether to make a threat by filing a lawsuit; (b) if threatened, the defendant may settle immediately or refuse; and (c) the (potential) final stage involves the plaintiff deciding whether to go to trial. My use of particular labels is intended to assist in building intuition and in relating the study to previous work, but they are not intrinsic to my results.

This article proceeds as follows. In Section 2, I describe the simple crisis bargaining game that is the basis of subsequent analysis. In particular, I enumerate the assumptions and restrictions imposed in particular special cases of the game, two of which have been previously considered in the literature. I also present an alternative PBE solution that approximates the assumptions about player types used in deriving the QRE and has essentially the same amount of learning as the previously considered PBE. I elucidate the relationship between these special cases by analytically comparing the distribution of types. In Section 3, I examine a number of comparative statics of the equilibrium for each special case. I show that changing arbitrarily chosen values for the uncertainty about player types can change qualitative features of the QRE and PBE highlighted in previous work. Intuitions based on comparative statics can radically change depending on the amount of variability assumed. Along the way, some previous explanations for properties of the particular PBE and QRE are reexamined and corrected. In the appendix, I present the derivations for the equilibrium of the crisis bargain game using the distribution of types described in each of the special cases from Section 2.

2 A Crisis Bargaining Game

I begin by describing the general features of a crisis bargaining game and will subsequently describe specifics about the stochastic components of the game. My aim is to make clear the assumptions and restrictions that are imposed in the context of different solution concepts and particular special cases.

The sequential game has two players, State A and State B, each desiring to possess a good currently held by State B. The sequence of moves and available actions are as follows. At the initial information set h_0^N, Nature first defines the types of players by drawing from a multivariate distribution a random vector of parameters that will affect payoffs, $\varepsilon = (\varepsilon_A, \varepsilon_B) = (\varepsilon_{A1n}, \varepsilon_{A1c:c}, \varepsilon_{A3b}, \varepsilon_{A3f}, \varepsilon_{B2c}, \varepsilon_{B2r:f}, \varepsilon_{B2r:b})$. At h_1^A, State A then chooses action $a_1 \in$ (Challenge, Not Challenge). State B observes a_1. If $a_1 =$ Challenge then at h_2^B State B must choose $a_2 \in$ (Concede, Resist). If $a_1 =$ Challenge, then State A observes a_2. If $a_2 =$ Resist then at h_3^A State A must choose $a_3 \in$ (Fight, Back Down). The game is represented in Fig. 1.

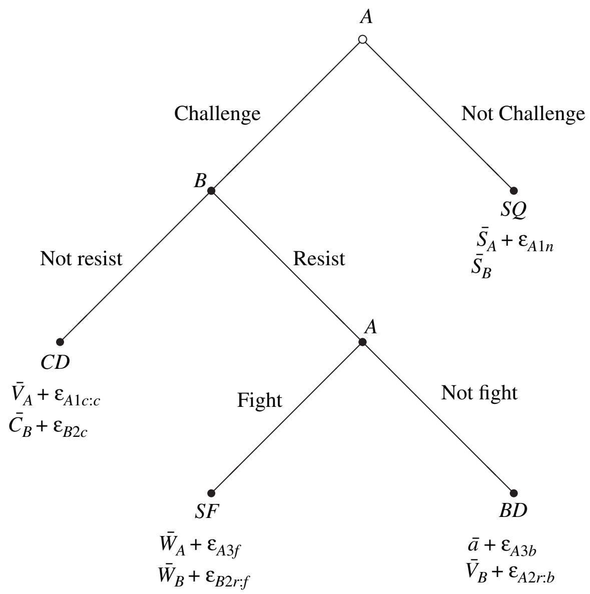

Fig. 1 Crisis bargaining game.

The payoff for each outcome is determined by a linearly additive utility function that combines a fixed payoff (denoted by a roman letter with a bar, e.g., $\bar{S}$), and stochastic payoff (denoted by an element of ε, introduced above). If a_1 = Not Challenge, then players receive $(\bar{S}_A + \varepsilon_{A1n}, \bar{S}_B)$. Since $\bar{S}_B$ plays no role in any player's decisions and therefore can take on an arbitrary value, a stochastic term is omitted to simplify presentation without loss of generality. If a_1 = Challenge and a_2 = Concede, then players receive $(\bar{V}_A + \varepsilon_{A1c:c}, \bar{C}_B + \varepsilon_{B2c})$. If a_1 = Challenge and a_2 = Resist and a_3 = Back Down, then players receive $(\bar{a} + \varepsilon_{A3b}, \bar{V} + \varepsilon_{B2r:b})$. If a_1 = Challenge and a_2 = Resist and a_3 = Fight, then players receive $(\bar{W}_A + \varepsilon_{A3f}, \bar{W}_B + \varepsilon_{B2r:f})$. The $a : a'$ notation specifies information associated with the stochastic payoff to action a, which may be realized depending upon another player's subsequent choice a'. For example, at h_2^B State B may learn information about the stochastic payoffs associated with a = Concede, as well as the potential stochastic payoffs that follow from a = Resist; the choice by A of a' = Fight is associated with $\varepsilon_{B2r:f}$, and a' = Back Down is associated with $\varepsilon_{B2r:b}$. The payoffs for each player are also shown in Figure 1.

Basic information available to the players is as follows. Both players know the structure of the game. Both players know the fixed payoff values $(\bar{S}_A, \bar{S}_B, \bar{V}_A, \bar{V}_B, \bar{a}, \bar{C}_B, \bar{W}_A, \bar{W}_B)$ and the multivariate probability distribution F of player types from which Nature draws the values of ε. They perfectly observe and recall all previous moves by all players, and players will partially observe the initial move by Nature in the case of the PBE.

In the context of the PBE, we assume that before State A chooses a_1, the values of ε_A are privately revealed to State A and those of ε_B are privately revealed to state B. This is not true in the context of the QRE. Instead, following the "agent with perceptual error" version of the QRE (McKelvey and Palfrey 1998), each information set h_i^j is independently played by an agent of the relevant State. Although the payoffs are assumed to be the fixed value (without any stochastic term), an agent makes a perceptual error about the value of each possible action, simply adding a stochastic term to the value of each available branch, and chooses the one with the greatest (perceived) total utility. The agent can only calculate expected utilities for outcomes from future information sets based on the commonly known fixed payoffs and the known and independent distribution of stochastic errors. All agents are assumed to be uninformed about the perceptual errors of other agents in the future; therefore players and their agents can be treated as having symmetric information. Despite their different motivations in the QRE and PBE models, the stochastic components in both models can be fruitfully compared in terms of the distributions of real or perceived total utilities; this comparison is undertaken at the end of this section.

Before considering the PBE and QRE solutions, the distribution of ε needs to be specified in order to define special cases of the crisis bargaining game. The derivations of the equilibrium solutions for these special cases are presented in the appendix.

2.1 *Distribution Assumptions Used in the Original PBE (PBE1)*

In this special case, the utility to both players of fighting a war is stochastic, while the utility of the Back Down outcome is stochastic only for State A. All other payoffs are known with certainty by the players. Formally, the stochastic payoff components of numerous outcomes are set identically to zero: State A's stochastic value of Status Quo ($\varepsilon_{A1n} = 0$) and having State B Concede ($\varepsilon_{A1c:c} = 0$), as well as State B's stochastic value of Conceding ($\varepsilon_{B2c} = 0$) and having State A Back Down ($\varepsilon_{B2r:b} = 0$). Both players know that these stochastic terms are set to zero and that the other stochastic components ($\varepsilon_{B2r:f}$, ε_{A3b}, ε_{A3f}) are independently and identically normally distributed, with mean zero and variance σ_1^2. These are the same distributional assumptions about player types initially described by Lewis and Schultz (2003, Sections 2.1–2.3) and used in their PBE derivations.

2.2 *Alternative Distribution Assumptions for a PBE (PBE2)*

For an alternative PBE solution to the game, let the value to State B of Resist be stochastic ($\varepsilon_{B2c} \neq 0$), but the difference between War and State A Back Down be nonstochastic, $\varepsilon_{B2r} = \varepsilon_{B2r:b} = \varepsilon_{B2r:f}$. Also in contrast to the original PBE, let the values to State A for remaining with Status Quo or having State B Concede be stochastic. Formally, the components (ε_{A1n}, $\varepsilon_{A1c:c}$, ε_{A3b}, ε_{A3f}, ε_{B2r}, ε_{B2c}) are independently and identically normally distributed, with mean zero and variance σ_2^2. These assumptions are designed to have stochastic features similar to those assumed in the QRE, which will be described next.

2.3 *Distribution Assumptions Used in the Original QRE*

As previously discussed, the payoffs of all outcomes are fixed in the QRE, with the stochastic components set to zero ($\varepsilon_{A1n} = \varepsilon_{A1c:c} = \varepsilon_{B2c} = \varepsilon_{B2r:b} = \varepsilon_{B2r:f} = \varepsilon_{A3b} = \varepsilon_{A3f} = 0$). Instead, a different set of stochastic terms is added to the utility of each *action* rather than outcome. A standard motivation for these alternative stochastic terms is the existence of a perceptual error of an agent who acts on behalf of a State (McKelvey and Palfrey 1998). For action a,

$$U_{ij}(a) = \bar{U}_{ija} + \nu_{ija}, \tag{1}$$

where $\bar{U}_{ija}$ is the expected payoff to player i at node j for action a and ν_{ija} is the associated perceptual error. The action a' with the highest total perceived value is chosen,

$$U_{ij}(a') > U_{ij}(a'') \quad \forall a'' \neq a' \quad \text{where } a', a'' \in a. \tag{2}$$

The perceived value of each action in the current crisis bargaining is

$$\begin{aligned}
U_{A3}(\text{Fight}) &= \bar{W}_A + \nu_{A3f} \\
U_{A3}(\text{Back Down}) &= \bar{a} + \nu_{A3b} \\
U_{B2}(\text{Resist}) &= P_F\bar{W}_B + (1-P_F)\bar{V}_B + \nu_{B2r} \\
U_{B2}(\text{Concede}) &= \bar{V}_B + \nu_{B2c} \\
U_{A1}(\text{Challenge}) &= P_R(P_F\bar{W}_A + (1-P_F)\bar{a}) + (1-P_R)\bar{V}_A + \nu_{A1c} \\
U_{A1}(\text{Not Challenge}) &= \bar{S}_A + \nu_{A1n},
\end{aligned}$$

where P_R and P_F are the equilibrium probabilities of State B Resist and State A Fight, respectively. The stochastic errors ν_{ija} are independently distributed normal with mean zero and variance σ_Q^2. This set of assumptions was previously used in a QRE solution derived by Lewis and Schultz (2003).

2.4 *Comparing Distributions of Player Types*

Comparing the distribution of player types highlights the similarities between the QRE and PBE2 assumptions and the differences between assumptions used in the QRE and PBE1. Although the equilibrium probabilities of Fight for each of these special cases (P_F, $P_{F|C}$, and $P'_{F|C}$ in the appendix) might appear the most different, this is not because the distribution of preferences for the two final alternatives is different in each of the versions. Indeed, the distribution of State A's preferences between Back Down and Fight is identical for all three versions if $\sigma_1 = \sigma_Q = \sigma_2$,

$$\Pr\left(\bar{W}_A + \varepsilon_{A3f} > \bar{a} + \varepsilon_{A3b}\right) = \Pr\left(\bar{W}_A + \nu_{A3f} > \bar{a} + \nu_{A3b}\right). \tag{3}$$

This equality also makes clear why the prior probability distribution for State A choosing to Fight in PBE1 and PBE2 is the same as the equilibrium probability of the QRE. Before State B observes the initial choice by State A to Challenge or not, assuming that all σ_k are set to be equal, State B has prior beliefs in each PBE that are no different from a similar player facing a QRE setup. Claims that signaling drives the behavior of a PBE comparative static result can be evaluated by considering the difference between the prior and posterior, elaborated in appendix Eqns. (A10) and (A11). Such differences will be considered in the next section.

The equality in the distribution of types across all three special cases for the relative values of Fight and Not Fight does not hold for any other part of the game. However, by design, the stochastic assumptions underlying QRE and PBE2 produce the following equivalence for State B for $\sigma_Q = \sigma_2$ and a given probability of Fight, P_F^*:

$$\begin{aligned}
&\Pr\left(P_F^*\bar{W}_B + (1-P_F^*)\bar{V}_B + \nu_{B2r} > \bar{C}_B + \nu_{B2c}\right) \\
&\quad = \Pr\left(P_F^*\bar{W}_B + (1-P_F^*)\bar{V}_B + \varepsilon_{B2r} > \bar{C}_B + \varepsilon_{B2c}\right).
\end{aligned} \tag{4}$$

Note that the distribution of preferences for State B in PBE1 differs from PBE2 by setting $\varepsilon_{B2c} - (1-P_F^*)\varepsilon_{B2r} = 0$. A more limited equivalence between the QRE and PBE2 exists with respect to State A's distribution of preferences over the status quo and having State B Concede, assuming $\sigma_Q = \sigma_2$ and a given probability of Resist, P_R^*:

$$\Pr\left(\bar{S}_A + \varepsilon_{A1n} > P_R^*(\bar{V}_A + \varepsilon_{A1c:c})\right) = \Pr\left(\bar{S}_A + \nu_{A1n} > P_R^*(\bar{V}_A + \nu_{A1c:c})\right). \tag{5}$$

The assumptions underlying PBE1 treat State A's preference over Status Quo and Concede as deterministic instead of probabilistic, since it imposes the restriction that

$\varepsilon_{A1c:c} = \varepsilon_{A1n} = 0$. In these ways, the alternative PBE2 solution uses assumptions that are more comparable to the assumptions used in the QRE.

3 PBE versus QRE: Comparative Statics

In this section I consider comparative statics based on the equilibrium solutions for the three special cases of the game described in the previous section. Details of the derivations for the solutions are presented in the appendix. In the context of comparative statics previously considered in the literature, I show that there is negligible updating in the original PBE and that the alternative PBE has an amount of updating similar to that in the original. I then show that a main claim about the differences between the QRE and PBE in the monotonicity of the probability of war is an artifact of arbitrarily chosen values for the amount of variability in player types. I also correct Lewis and Schultz's explanation for why the nonmonotonicity was observed in their original analysis of the PBE. Finally, I reconsider what could be inferred from data on the distribution of outcomes. In this context, I again show that the alternative PBE (PBE2) has essentially the same amount of updating as the original PBE (PBE1), as well as the same qualitative features as the QRE.

To ensure that the subsequent analysis is comparable with previous work, I adopt the identification restrictions on the payoffs, setting the average value of possessing the good as one ($\bar{V}_A = \bar{V}_B = 1$) and the value of not possessing the good as zero ($\bar{S}_A = \bar{C}_B = 0$) (Lewis and Schultz 2003). The values of the remaining free payoff parameters ($\bar{W}_A$, $\bar{W}_B$, $\bar{a}$) will be interpreted relative to the difference between having and not having the contested good for each State, $\bar{V}_B - \bar{C}_B = 1$ and $\bar{V}_A - \bar{S}_A = 1$. For example, an average audience cost of zero ($\bar{a} = 0$) would be equivalent to the average value of not challenging and not obtaining the good ($\bar{S}_A = 0$).

The variances in each version of the game will also be fixed for identification. Previous work has set $\sigma_1 = 1$ and $\sigma_Q = 1$, despite there being many more stochastic terms in the special case used by the QRE than in the original PBE. Lewis and Schultz note that given the fixed payoffs that set the utility scale for each player, the standard deviation σ is a substantive choice: "In setting $\sigma = 1$, we impose the assumption that the standard deviation of the payoff shocks is equal (in utility terms) to the difference between having and not having the good" (p. 356). But it is not simply an issue of the metric, as I will show that this choice has substantive implications for the comparative statics as well.

Adding or restricting stochastic terms can in some cases simply be equivalent to increasing or decreasing the value of σ. For example, setting $\varepsilon_{B2c} = 0$ in the alternative PBE2 would simply reduce the standard deviation by which the payoffs are normalized in the probability of State B resisting a challenge from $\sigma_2\sqrt{2}$ to σ_2 (see the appendix for details). Except where explicitly noted, I will use $\sigma_1 = \sigma_Q = \sigma_2 = 1$ for the following comparative statics to achieve comparability with previous work. It is worth emphasizing that Lewis and Schultz originally set $\sigma_1 = 1$ and $\sigma_Q = 1$ without empirical or theoretical justification, and no greater claim to reflecting real conflicts is made here. Given the arbitrary baseline choices of σ, I also provide some illustrations of how interpretations can change simply by choosing other values for the standard deviations.

3.1 *The Negligible Role of Signaling and Updating*

In the context of the current simple crisis bargaining setup, State B would like to know whether State A will Fight or Back Down if faced with the choice. Using the QRE, State A and State B are equally (un)informed and only the prior probability distribution of $a_3 \in$ (Fight, Back Down) is known until the final node is reached. In the context of the PBE

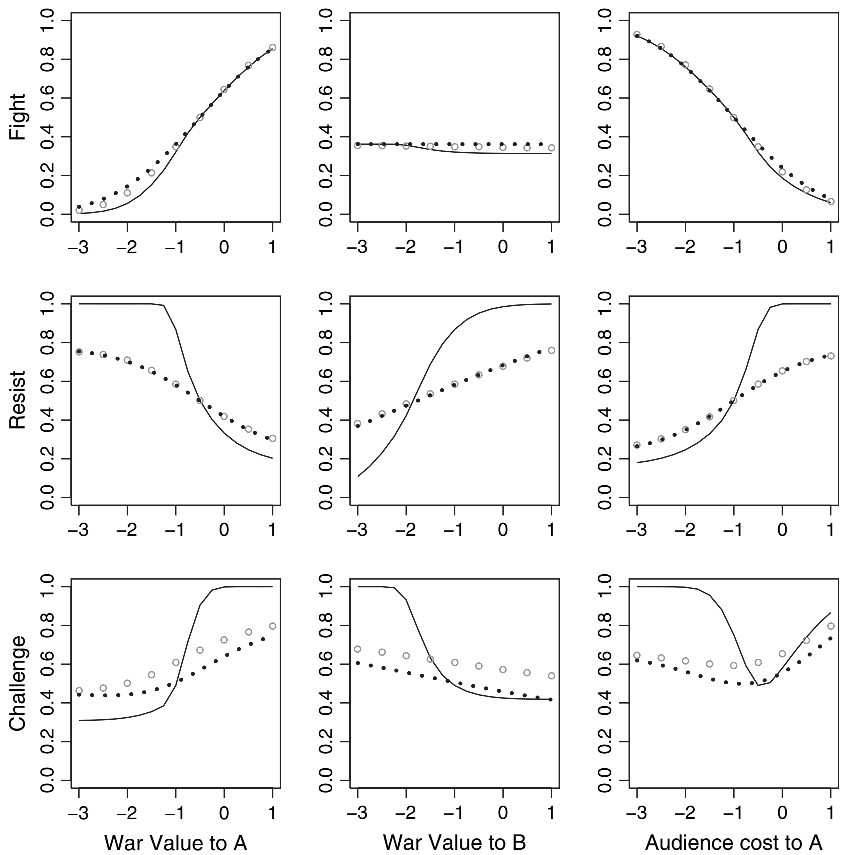

Fig. 2 Effects of changing various game payoffs on the probability of each action. Each payoff is varied from −3 to 1, with a baseline payoff vector of ($\bar{W}_A = -1, \bar{W}_B = -1, \bar{a} = -1/2$). Three curves are shown: QRE (••••••), PBE1 (——), and PBE2 (○○○○○○). For Fight, the QRE values are also the prior probabilities for the PBE solutions.

solutions (PBE1 and PBE2), State A knows with certainty its type before making its initial choice about Challenging, and therefore which action at the final node would be selected. In the presence of this information asymmetry, State B can calculate the posterior probability of State A's type conditional on observing the initial Challenge rather than rely on the prior probability distribution. Analytical results are presented in appendix Eqns. (A1) and (A3). It has been argued that the "signaling and updating dynamics" have fundamentally important implications for the claims that will be discussed in detail in the next two subsections.

How much does State B learn about the willingness of State A to fight from observing the initial Challenge? The role of learning and the importance of signaling can be quantified by the difference between the prior and posterior distributions of the types. If the posterior is not different from the prior, then observing the initial move contains no informative signal and updating beliefs would not be a distinguishing feature of the strategic interaction.

Figure 2 shows the probability of each action being taken at various payoffs. To ensure comparability with previous work and other comparative statics to follow, the baselines payoffs are $\bar{W}_A = -1$, $\bar{W}_B = -1$, and $\bar{a} = -0.5$. One of these payoffs is varied at a time from −3 to 1, holding the others at the baseline values. The solid lines are the original PBE1 equilibrium probabilities, the dots are the QRE, and the circles are the alternative PBE2 probabilities.

In the graphs of the probability of Fight, the QRE values (dots) are unconditional probabilities, while the two PBE values (line and circles) are conditional posterior probabilities. Since $\sigma_1 = \sigma_Q = \sigma_2 = 1$ in this case, the unconditional QRE probabilities were shown in the previous section to be the same as the PBE prior probabilities. For both PBEs the prior and posterior probabilities of Fight are essentially indistinguishable, and both are negligibly different from the prior (the dots).

The main differences between the models occur in the probabilities of Resist and Challenge. In contrast to Fight, players do not update their beliefs about the probability of these actions through the course of the game. Relative to the differences in the distribution of types between the models, signaling and updating dynamics play a negligible role in explaining the patterns of these comparative statics and the patterns of comparative statics of outcome probabilities built from their products, to which I now turn.

3.2 *Monotonicity to Nonmonotonicity and Vice Versa*

Building on the equilibrium probabilities of actions, I consider comparative statics for the resulting equilibrium probabilities of outcomes. The outcome probabilities are defined as

$$\begin{aligned} P_{SQ} &= 1 - P_C && \text{Status Quo} \\ P_{CD} &= P_C(1 - P_R) && \text{Concede} \\ P_{BD} &= P_C P_R (1 - P_F) && \text{Back Down} \\ P_{SF} &= P_C P_R P_F && \text{Stand Firm,} \end{aligned}$$

where the notations $P_F = \Pr(a_3 = \text{Fight})$, $P_R = \Pr(a_2 = \text{Resist})$, and $P_C = \Pr(a_1 = \text{Challenge})$ are for the QRE solution. The equilibrium action probabilities from the PBEs would be multiplied in the same way to obtain outcome probabilities for the other two special cases.

For each solution concept and distribution of types, Fig. 3 shows the effect of changing various game payoffs on the probability with which each terminal node is reached. One of the nonnormalized mean payoff values is varied while the others are fixed at baseline mean values. Again to ensure comparability with previous work, the baselines payoffs are $\bar{W}_A = -1$, $\bar{W}_B = -1$, and $\bar{a} = -0.5$.

With respect to the QRE and the original PBE1, Lewis and Schultz comment that "even when the general direction of the relationship is the same, the functional forms can be quite different" (p. 353). The striking qualitative differences between QRE and PBE1 solutions are not an intrinsic feature of the presence or absence of information asymmetry, but rather mainly a product of choices about the distribution of types in the PBE1. For the alternative PBE2 (the circles), the curves of the equilibrium probabilities are qualitatively similar to the QRE solutions over the range of values for the mean payoffs shown in Fig. 3. For some outcomes and payoffs considered, the QRE and the PBE2 are essentially indistinguishable.

One of Lewis and Schultz's main claims about the substantive differences between the QRE and original PBE1 solutions is that "while the QRE predicts a monotonic increase in the probability of war as $\bar{W}_A$ increases, the PBE predicts a non-monotonic relationship—with the probability of war first increasing then decreasing" (Lewis and Schultz 2003, p. 353). However, the presence or absence of the nonmonotonicity in the probability of

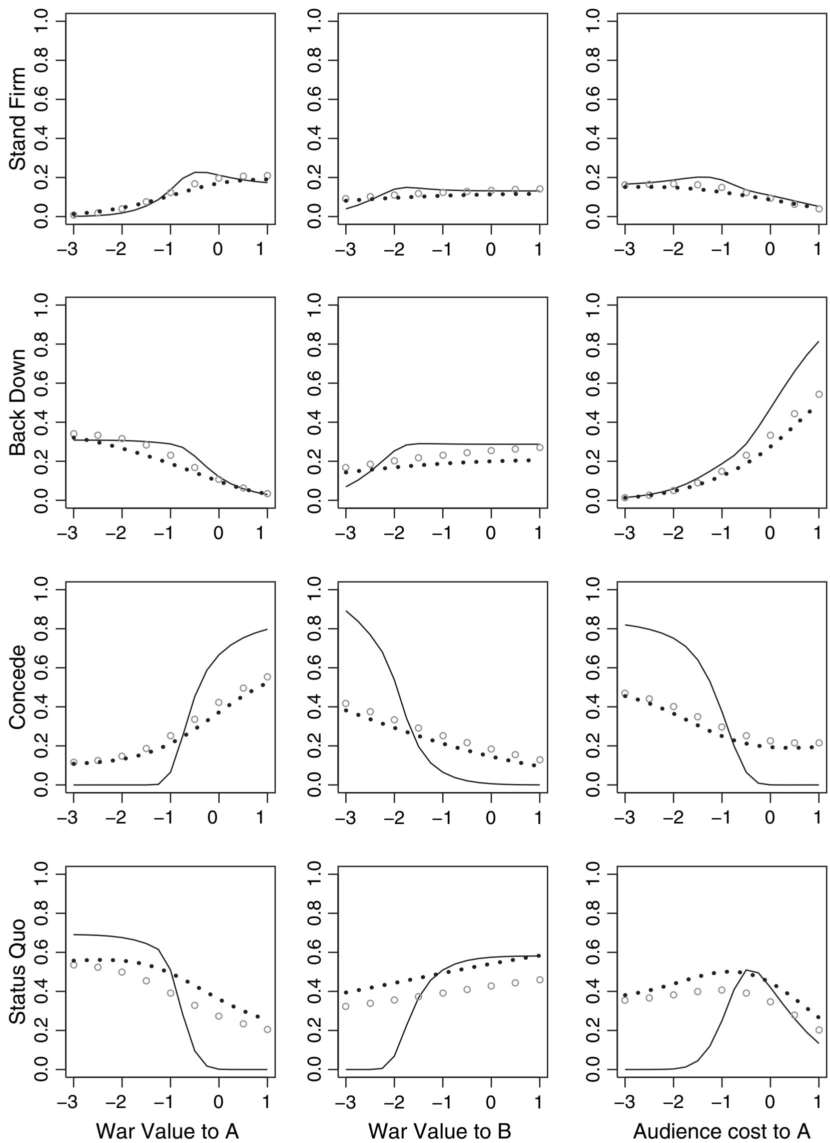

Fig. 3 Effects of changing various game payoffs on the fraction of times each terminal node is reached. Each payoff is varied from −3 to 1, with a baseline payoff vector of ($\bar{W}_A = -1, \bar{W}_B = -1, \bar{a} = -1/2$). Three curves are shown: PBE1 (——), QRE (••••••), and PBE2 (○○○○○○).

Stand Firm with respect to the cost of War for State A is neither a product of signaling and updating, nor even a intrinsic feature of the original PBE1. Nonmonotonicity can be removed in the case of the PBE1 by increasing the variability in types, and can be induced in the case of the QRE by reducing the randomness of each choice. The effect of changing the amount of variability on the relationship between the cost of War ($\bar{W}_A$) and equilibrium probability of Stand Firm is illustrated in Fig. 4. This example demonstrates how the choice of standard deviation of stochastic terms can fundamentally reshape the qualitative features of the comparative statics.

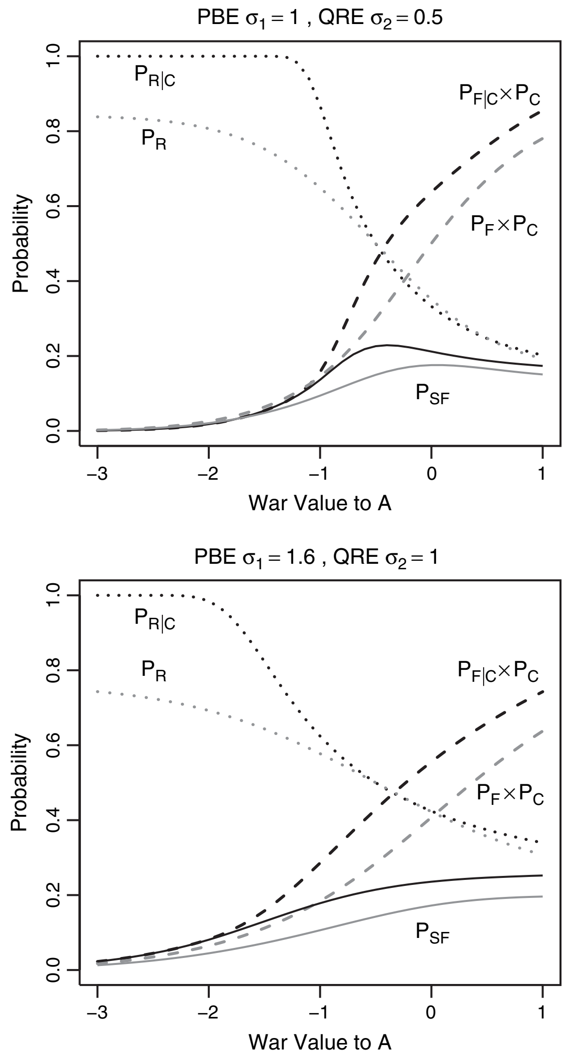

Fig. 4 Decomposing the probability of a Stand Firm outcome, $P_{SF} = P_C \times P_R \times P_F$. Effects of changing the standard deviation of stochastic component on the relationship between State A's game payoffs for War on the probability choices (State B Resist; State A Challenge and Fight) and outcome (State A Stand Firm). W_A is is varied from -3 to 1, with other payoffs being fixed ($\bar{W}_B = -1$, $\bar{a} = -1/2$). Curves are shown for the PBE1 (black) and QRE (gray).

Since nonmonotonicity is neither intrinsic nor unique to the original PBE, it is necessary to reconsider Lewis and Schultz's explanation for why the PBE they considered is nonmonotonic. Their explanation for this pattern is simple and intuitively very appealing: "The maximum probability of war occurs when the mean value of $\bar{W}_A$ is near the mean value of $\bar{a}$, in which case there is maximum uncertainty about whether or not A will fight. In this range, the potential for [State] B to 'mistakenly' resist genuine threats is the highest" (Lewis and Schultz 2003, p. 353). Though intuitively appealing, the statement is not an accurate characterization of the complex interaction between the sequence of equilibrium probabilities that produces the surge in the probability they originally plotted for the PBE.

Given maximum uncertainty about State A's true preferences as the basis for explaining the observed nonmonotonicity in the original PBE, it might seem odd that given other values for the baseline of $\bar{a}$, the location of the maximum value P_{SF} over a range of $\bar{W}_A$ can move even further from $\bar{W}_A = \bar{a}$. As the variance of the stochastic terms increases, we have also just seen that the maximum probability also moves far away from this point that Lewis and Schultz single out for their explanation. The proximity of the peak to $\bar{W}_A = \bar{a}$ in these particular comparative statics for the original PBE is somewhat special.

To understand the theoretical process of arriving at war more generally, and the special features of this particular case, it is useful to illustrate how the variability of preferences filters through the probabilities of choices at individual nodes and into the final outcomes. Figure 4 also has the curves for P_R and $P_C \times P_F$ for the QRE and the equivalent curves for the PBE. Nonmonotonicity in the outcome probability of State A Standing Firm (P_{SF}), of the type highlighted in previous work, depends on the slope of the $P_{R|C}$ and $\tilde{P}_C \times P_{F|C}$ curves being sufficiently symmetric and steep. In the equilibrium probabilities considered so far, σ normalizes the effect of the average payoffs. As σ is made larger, a given change in any particular payoff induces less of a change in the probability of a particular outcome, which is a verbose way of saying the slope decreases. For this particular case, when the standard deviation is increased to $\sigma_1 = 1.6$, the slopes are decreased sufficiently to no longer induce nonmonotonicity over the range of values considered. In each of the two panels shown here, we see that the QRE and PBE curves for the probability of Resist and the combined probability of Challenge and Fight can be similar near the point at which they cross, producing qualitatively similar comparative statics on the outcome Stand Firm.

3.3 *More on Limited Learning, by Players and Researchers*

Given a particular distribution of outcome frequencies, what could be learned about the unrestricted game parameters representing audience costs and the utilities of war? Determining these utilities given information outcomes is simply a matter of inverting the probabilities functions and solving for $\bar{a}$, $\bar{W}_A$, and $\bar{W}_B$. In this context, I again show that the striking differences between the original PBE1 and the QRE are a result of the different distributional assumptions rather than the dynamics of updating beliefs. The alternative PBE2 has the same key qualitative features as the QRE, and there is again negligible updating.

These next comparative statics will again follow previous work for comparability and use a baseline frequency distribution of outcomes: $P^{\#}_{SF} = 0.10$, $P^{\#}_{BD} = .25$, $P^{\#}_{CD} = .15$, $P^{\#}_{SQ} = .50$ (Lewis and Schultz 2003). In each column, the named outcome fraction is varied from 0.1 to 0.6 while holding the relative baseline proportion of the other outcomes fixed. Again, in keeping with previous work, when varying one probability, the relative size of the other probabilities are kept fixed. Focusing on the case of varying CD, increasing the frequency of CD results in a decrease in the frequency of all the other outcomes SQ, BD, and SF. Consider for a given value P^{*}_{CD},

$$\begin{aligned} P^{*}_{SQ} &= (1 - P^{*}_{CD})/(P^{\#}_{SQ} + P^{\#}_{BD} + P^{\#}_{SF}) \times (1 - P^{\#}_{C}) \\ P^{*}_{BD} &= (1 - P^{*}_{CD})/(P^{\#}_{SQ} + P^{\#}_{BD} + P^{\#}_{SF}) \times P^{\#}_{C} \times P^{\#}_{R} \times (1 - P^{\#}_{F}) \\ P^{*}_{SF} &= (1 - P^{*}_{CD})/(P^{\#}_{SQ} + P^{\#}_{BD} + P^{\#}_{SF}) \times P^{\#}_{C} \times P^{\#}_{R} \times P^{\#}_{F}, \end{aligned} \tag{6}$$

where $P^{\#}_{j}$ is a baseline probability, and the other P^{*}_{j} are rescaled probabilities given the alternative value P^{*}_{CD}. Since the relative size of SF and BD is determined solely by the probability of Fighting ($P^*_{BD}/P^*_{SF} = [1 - P^{\#}_F]/P^{\#}_F$), the probability of Fighting is constrained by design not to change when altering CD, and instead all movement is induced in P^*_R and P^*_C. Increasing P^*_{CD} leads to more of an increase ($1 - P^*_R$) than in P^*_C and no change in P^*_F.

In Fig. 5, I plot the utilities that would be inferred from a particular set of outcome frequencies. The solid line, which is sometimes going in the opposite direction or distant from the other curves, is the original PBE1. The dots are the QRE equilibrium values. These two curves replicate the curves that are the subject of Lewis and Schultz's analysis. On the basis of these curves, they note that although there are some similarities, "at the same time, there are some very dramatic differences" (p. 356). In Fig. 5 the alternative PBE is again represented by circles. Unlike the original PBE, no "dramatic differences" appear between the alternative PBE and original QRE solutions.

Lewis and Schultz make their second main claim in the context of analyzing their figure of inferred payoffs given a particular distribution of outcome frequencies. I quote at length the passage since it is the most important substantive claim about differences between the PBE and QRE solutions.

> The most instructive of these differences is in how the estimators react to an increase in the rate of concessions by [State] B. When CD [State B Concedes] becomes more frequent, both estimators assume that war must have become less attractive for B; however, they disagree about what happened to [State] A's payoffs. The PBE estimator concludes that the costs of backing down must have increased (i.e., $\bar{a}$ decreased), thereby making A's threats more credible. Such an inference follows naturally from the logic of costly signaling: the more costly it is for the coercer to back down from a threat, the more likely it is that threat will be carried out (especially, Fearon 1994; Schultz 1999). The QRE estimator, on the other hand, concludes that $\bar{W}_A$ and $\bar{a}$ increased in tandem. Such changes increase the probability that [State] A will make a challenge in the first place—thereby giving [State] B many more opportunities to concede—without changing the relative probability of the SF and BD outcomes. *These contrasting reactions underscore the differences in the behavioral and informational assumptions underlying these two equilibrium concepts.* (Lewis and Schultz 2003, p. 358, emphasis added)

At the most basic level of demonstration, Fig. 5 illustrates that these claims are incorrect about the general properties of the PBE implied by the final sentence of the quote. The alternative PBE, with similar amounts of learning as the original PBE, can have the same qualitative features as the QRE. Instead of inferring the audience costs become increasingly negative when the frequency of State B conceding becomes very large, the direction of change for the alternative PBE is the same as the QRE.

Figure 6 illustrates that for both PBEs the prior and posterior probabilities of Fight are again essentially indistinguishable, and both are negligibly different from the prior (the dots). Since the marginal outcome frequencies are fixed by design, the curves of QRE probabilities and the PBE posterior probabilities are equivalent. The priors of the PBEs are the induced values consistent with the postulated posteriors.

The basic explanation in previous work for the qualitative features of the inferences that would be drawn from the QRE is also misplaced for a large range of CD and simply wrong for a further subset of the range. Along the way to showing this, I first will make some additional comments about the component probabilities and how they interact, since the process generating these comparative statics is complicated and benefits from further unpacking.

First, consider the inference that can be drawn about the payoff to State B of going to War ($\bar{W}_B$). For all three solutions (PBE1, PBE2, QRE), an increase in the probability of Concede by State B ($1 - P^*_R$) can be effected only by a decrease in $\bar{W}_B$, since the other

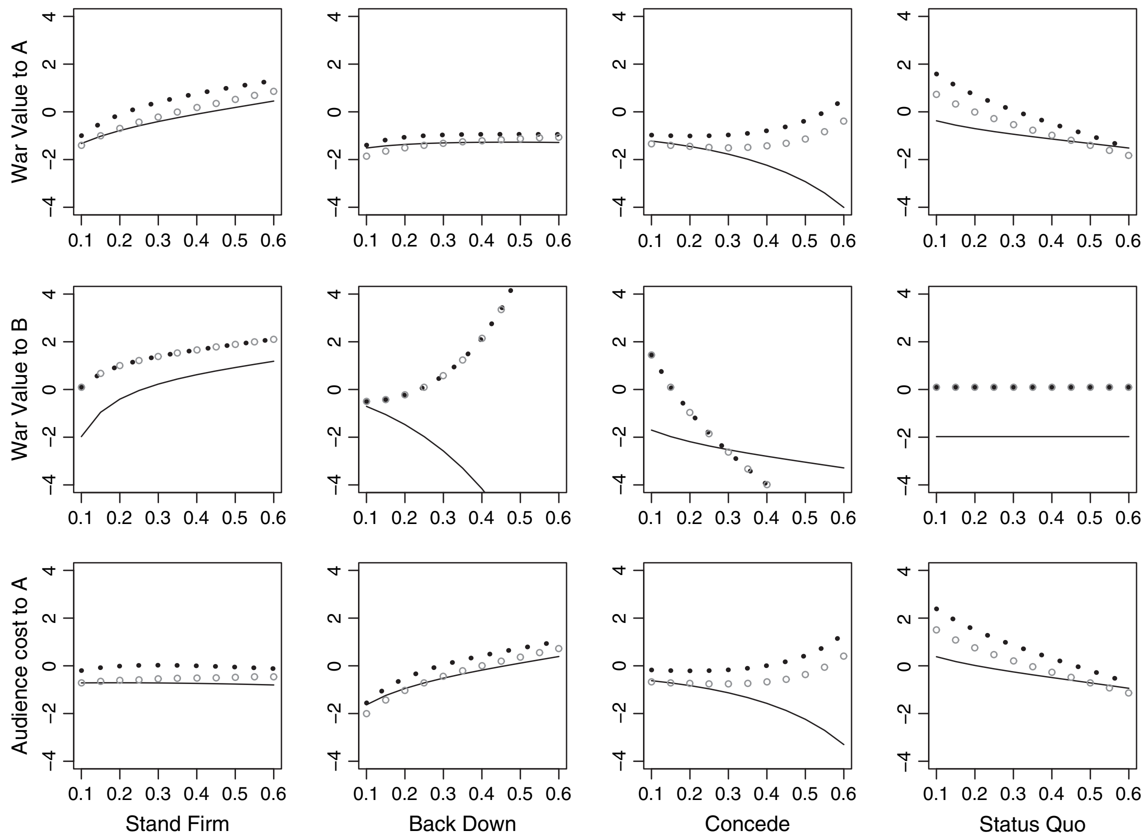

Fig. 5 Inference drawn about payoffs given particular distribution of observed outcome frequencies. The baseline distribution of outcomes is ($P_{SF} = 0.10$, $P_{BD} = .25$, $P_{CD} = .15$, $P_{SQ} = .50$). In each column, the named outcome fraction is varied from 0.1 to 0.6 while holding the relative baseline proportion of the other outcomes fixed. Three curves are shown: PBE1 (——), QRE (••••••), and PBE2 (○○○○○○).

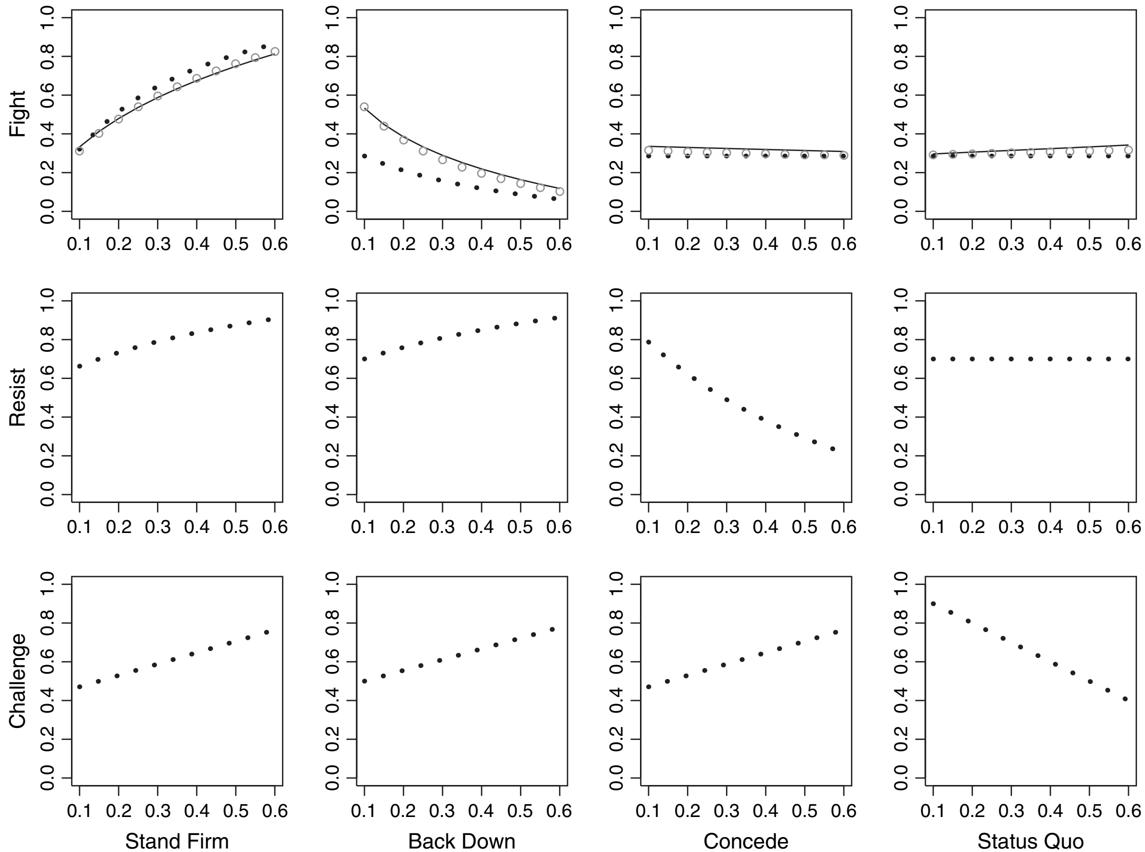

Fig. 6 Prior and posterior probabilities of actions. The baseline distribution of outcomes is ($P_{SF} = 0.10$, $P_{BD} = .25$, $P_{CD} = .15$, $P_{SQ} = .50$). In each column, the named outcome fraction is varied from 0.1 to 0.6 while holding the relative baseline proportion of the other outcomes fixed. Three curves are shown: QRE (••••••), PBE1 prior (——), and PBE2 prior (○○○○○○). For Fight, the QRE values are also the posterior probabilities for the PBE solutions.

parameters are normalized to a fixed value ($\bar{V}_B$, $\bar{C}_B$, σ_Q) or are not changing under the design of the comparative static (since $P_F^{\#} = P_F^*$ is held constant). This can be seen in the QRE solution,

$$1 - P_R^* = 1 - \Phi\left(\frac{P_F^*\bar{W}_B + (1 - P_F^*)\bar{V}_B - \bar{C}_B}{\sigma_Q\sqrt{2}}\right), \tag{7}$$

which has the same form as the PBE2 solution (compare Eqs. [A5] and [A6]. In the PBE1 solution, however, the effect of $\bar{W}_B$ is not attenuated by the multiplication of $P_F^{\#}$:

$$1 - P_{R|C}^* = 1 - \Phi\left(\frac{\bar{W}_B}{\sigma_1} + \frac{(1 - P_F^*)\bar{V}_B - \bar{C}_B}{P_F^*\sigma_1}\right), \tag{8}$$

which is simply a rearrangement of Eq. (A4) in the appendix. This difference explains why larger changes in $\bar{W}_B$ are necessary in the QRE and alternative PBE2, compared to the original PBE1, for equivalent changes in P_{CD}^*. Since $P_F^{\#} = P_F^*$ is fixed, the value of $\bar{W}_A$ and $\bar{a}$ do not otherwise affect State B's comparative statics.

For the QRE, the interesting thing about P_C^* is that there are actually two distinct ways that it can be increased in the equation:

$$P_C^* = \Phi\left(\frac{P_R^*(P_F^*\bar{W}_A + (1 - P_F^*)\bar{a}) + (1 - P_R^*)\bar{V}_a - \bar{S}_A}{\sqrt{2}\sigma_Q}\right). \tag{9}$$

Recall that $P_F^* = P_F^{\#}$ is fixed by design, and $\bar{V}_a$ and $\bar{S}_A$ are normalized constants. One way to increase P_C^*, highlighted by Lewis and Schultz, is to increase $\bar{W}_A$ and $\bar{a}$ in tandem. The two parameters must be changed in tandem because as in any probit such as Eq. (A2), keeping $P_F^* = P_F^{\#}$ fixed requires that the difference $\bar{W}_A - \bar{a}$ remain constant.

Another previously overlooked way of increasing the P_C^* in the QRE is by increasing $(1 - P_R^*)$ if $(P_F^{\#}\bar{W}_A + (1 - P_F^{\#})\bar{a}) < 1$. At the baseline values $(P_F^{\#}\bar{W}_A + (1 - P_F^{\#})\bar{a}) = -0.9$. Those who were puzzled as to why $\bar{W}_A$ and $\bar{a}$ curves are *decreasing* initially in the QRE but are essentially flat over most of the plausible values when varying P_{CD}^* may be reassured to find that this mechanism based on P_R^* exists.

In the QRE, what leads to the initial decrease and then upturn in $\bar{W}_A$ and $\bar{a}$ at the extreme frequency of observing State B choosing to Concede? To understand this, first also note the mapping of outcome frequencies back to the probabilities of choices.

$$\begin{aligned} P_C^* &= 1 - P_{SQ}^* && \text{State A Challenge} \\ 1 - P_R^* &= P_{CD}^*/P_C^* && \text{State B Concede} \\ P_F^* &= 1 - P_{BD}^*/(P_C^* \times P_R^*) = P_{SF}^*/(P_C^* \times P_R^*) && \text{State A Fight} \end{aligned} \tag{10}$$

While P_{SQ}^* (Eq. 6) and thus P_C^* (Eq. 10) are constrained to change linearly with changes in P_{CD}^*, the rate of increase in $1 - P_R^*$ is reduced by this same increasing value of P_C^*. Thus, while maintaining a linear rate of change in P_C^* induced by the design of the comparative static, the rate of change in P_R^* may not be adequate to account for this change in P_C^*. At extreme levels, $\bar{W}_A$ and $\bar{a}$ must jointly increase to account for the diminishing changes in P_R^*. But at low levels of P_{CD}^*, changes in P_R^* are not only enough to mostly account for changes in the frequency of State A Challenging, but indeed $\bar{W}_A$ and $\bar{a}$ must be reduced slightly. This is the opposite of Lewis and Schultz's description, and for some reasonably sized values of P^*_{CD} the negative direction of change in these parameters (though negligible in magnitude) is the same as for the original PBE.

4 Conclusions

The primary focus of this analysis has been on the distribution of player types in competing models of strategic choice. Key observable differences between the models considered have been shown to be the result of choices about the placement and variability of stochastic components, rather than being intrinsic consequences of signaling and updating dynamics. I have also illustrated how competing models can mimic some comparative statics of one another through the manipulation of the amount of variability.

The sensitivity of these models to changes in the distribution of types raises the issue of stochastic terms differing at various stages of a game. It is unlikely that in applications, with either experimental or observational data, all choices or outcomes have equal uncertainty. Continuing the example of military conflict, are preferences over the status quo better understood by players than the value of challenging and obtaining a contested good or the cost of fighting a war? Resolving such questions will be specific to particular sets of empirical cases and will likely require the collection of evidence supplementary to the strategic model itself. More generally, supplementary studies are needed to inform the placement of stochastic terms and narrow the range of plausible values to be considered in the context of testing or theorizing about these types of strategic choice model.

The opportunities of the parametric strategic choice are tempered by their greater complexity. Even in the context of this very simple model, correctly characterizing the interaction between parameters is difficult and laborious, at best. Small changes to the parameterization or identification restrictions, which are often treated as ancillary assumptions, can significantly alter what inferences would be drawn. These are not arguments against the use of these types of models, but rather arguments for more detailed study.

Appendix: Analytical Derivations of PBE and QRE Solutions

In this appendix, I derive the alternative PBE (PBE2) and contrast this solution with the original PBE (PBE1) and the QRE solutions derived in previous work. The PBE solutions are represented by the posterior probabilities $P_{F|C} = \Pr(a_3 = \text{Fight} \mid a_1 = \text{Challenge})$ and $P_{R|C} = \Pr(a_2 = \text{Resist} \mid a_1 = \text{Challenge})$, as well as $\tilde{P}_C = \Pr(a_1 = \text{Challenge})$. The QRE solutions are characterized by the probabilities $P_F = \Pr(a_3 = \text{Fight})$, $P_R = \Pr(a_2 = \text{Resist})$, and $P_C = \Pr(a_1 = \text{Challenge})$. I distinguish between the original PBE1 and alternative PBE2 solutions by adding a prime ($'$) to notation in the latter case.

Beginning at the last node h^A_3, the PBE1 and QRE solutions for the probability of State A choosing to Fight are, respectively,

$$\begin{aligned} P_{F|C} &= \Pr(\bar{W}_A + \varepsilon_{A3f} > \bar{a} + \varepsilon_{A3b} \mid a_1 = \text{Challenge}) \\ &= \Phi_2\left(\frac{\bar{W}_A - \bar{a}}{\sigma_1\sqrt{2}}, \frac{\bar{W}_A - C^*}{\sigma_1}, \frac{1}{\sqrt{2}}\right) \Big/ \tilde{P}_C \end{aligned} \tag{A1}$$

$$P_F = \Pr(\bar{W}_A + \nu_{A3f} > \bar{a} + \nu_{A3b}) = \Phi\left(\frac{\bar{W}_A - a}{\sigma_Q\sqrt{2}}\right), \tag{A2}$$

where $C^* = \frac{\bar{S}_A-(1-P_{R|C})\bar{V}_A}{P_{R|C}}$ and $\Phi_2(\theta_1, \theta_2, \rho)$ is the standard bivariate normal cumulative distribution function with correlation ρ evaluated at (θ_1, θ_2).

The PBE2 probability of fighting, $P'_{F|C}$, is different from (A1) due to the allowance for a stochastic component in State A's payoff to the Status Quo (ε_{A1n}) and to having State B Concede ($\varepsilon_{A1c:c}$). A simple justification of this addition could follow from a desire to make the amount of variability more comparable with the amount of stochastic variation assumed in the QRE solution. Substantively, it could be argued that State A has private information about their own future value for Status Quo and State B Conceding. The equilibrium probability of State A choosing to Fight is

$$\begin{aligned} P'_{F|C} &= \Pr(\bar{W}_A+\varepsilon_{A3f}>\bar{a}+\varepsilon_{A3b}\,|\,a_1=\text{Challenge}) \\ &= \Pr\left(\bar{W}_A+\varepsilon_{A3f}>\bar{a}+\varepsilon_{A3b}\,|\,\bar{W}_A+\varepsilon_{A3f}\vee\bar{a}+\varepsilon_{A3b}>\frac{S_A+\varepsilon_{A1n}-(1-P'_{R|C})(\bar{V}_A+\varepsilon_{A1c:c})}{P'_{R|C}}\right) \\ &= \Pr\left(\bar{W}_A-\bar{a}>\varepsilon_{A3b}-\varepsilon_{A3f},\bar{W}_A-C'^*>\frac{(1-P'_{R|C})\varepsilon_{A1c:c}+\varepsilon_{A1n}}{P'_{R|C}}-\varepsilon_{A3f}\right)\Big/\tilde{P}'_C \qquad \text{(A3)} \\ &= \Phi_2\left(\frac{\bar{W}_A-\bar{a}}{\sigma_2\sqrt{2}},\frac{\bar{W}_A-C'^*}{\sigma_2\sqrt{1+\omega/P'^2_{R|C}}},\frac{1}{\sqrt{2[1+\omega/P'^2_{R|C}]}}\right)\Big/\tilde{P}'_C, \end{aligned}$$

where $\omega = ((1-P'_{R|C})^2+1)$. The derivation of the bivariate normal distributions in Eqs. (A3) and (A1) follows immediately from Theorem 3.2.1 of Tong (1990). $P'_{R|C}$ now appears not only in C'^*, but also in the variance and covariance terms of the bivariate distribution.

As discussed in Section 2, the prior probability for both the original PBE1 and alternative PBE2 has the same value as the equilibrium probability for the QRE, if the standard deviations of the stochastic terms are set to be the same.

The PBE1 and the QRE probability of State B choosing to Resist an initial challenge by State A have been derived as, respectively,

$$\begin{aligned} P_{R|C} &= \Pr(P_{F|C}(\bar{W}_B+\varepsilon_{B2r:f})+(1-P_{F|C})\bar{V}_B>\bar{C}_B) \\ &= \Phi\left(\frac{P_{F|C}\bar{W}_B+(1-P_{F|C})\bar{V}_B-\bar{C}_B}{P_{F|C}\sigma_1}\right) \end{aligned} \qquad \text{(A4)}$$

$$\begin{aligned} P_R &= \Pr(P_F\bar{W}_B+(1-P_F)\bar{V}_B+\nu_{B2r}>\bar{C}_B+\nu_{B2c}) \\ &= \Phi\left(\frac{P_F\bar{W}_B+(1-P_F)\bar{V}_B-\bar{C}_B}{\sigma_Q\sqrt{2}}\right). \end{aligned} \qquad \text{(A5)}$$

The alternative PBE2 stochastic assumption leads to an equilibrium probability of State B Resisting,

$$\begin{aligned} P'_{R|C} &= \Pr(P'_{F|C}(\bar{W}_B+\varepsilon_{B2r})+(1-P'_{F|C})(\bar{V}_B+\varepsilon_{B2r})>\bar{C}_B+\varepsilon_{B2c}) \\ &= \Pr(P'_{F|C}\bar{W}_B+(1-P'_{F|C})\bar{V}_B-\bar{C}_B>\varepsilon_{B2c}-\varepsilon_{B2r}) \\ &= \Phi\left(\frac{P'_{F|C}\bar{W}_B+(1-P'_{F|C})\bar{V}_B-\bar{C}_B}{\sigma_2\sqrt{2}}\right). \end{aligned} \qquad \text{(A6)}$$

Note that in the alternative PBE2, $\sqrt{2}$ appears instead of $P'_{F|C}$ in the denominator of the equilibrium solution for the probability of Resist (A6), which is the same form as the QRE solution (A5) though the value is equivalent only if $P_F = P'_{F|C}$ and $\sigma_Q = \sigma_2$. This follows from a key feature PBE2: the equality $\varepsilon_{B2r} = \varepsilon_{B2r:b} = \varepsilon_{B2r:f}$. This assumption implies that the difference in value for State B between War and State A Back Down is nonstochastic, which is the same assumption made in the QRE.

Consider now the first move by State A, the probability of the initial Challenge. Again, the PBE1 and QRE equilibrium probabilities are

$$\begin{aligned}\tilde{P}_C &= \Pr(P_{R|C}(\bar{W}_A + \varepsilon_{A3f} \vee \bar{a} + \varepsilon_{A3b}) + (1 - P_{R|C})\bar{V}_a > \bar{S}_A) \\ &= 1 - \Phi_2\left(\frac{C^* - \bar{W}_A}{\sigma_1}, \frac{C^* - \bar{a}}{\sigma_1}, 0\right) = 1 - \Phi\left(\frac{C^* - \bar{W}_A}{\sigma_1}\right)\Phi\left(\frac{C^* - \bar{a}}{\sigma_1}\right)\end{aligned} \tag{A7}$$

$$\begin{aligned}P_C &= \Pr(P_R(P_F\bar{W}_A + (1 - P_F)\bar{a}) + (1 - P_R)\bar{V}_a + \nu_{A1c:c} > \bar{S}_A + \nu_{A1s}) \\ &= \Phi\left(\frac{P_R(P_F\bar{W}_A + (1 - P_F)\bar{a}) + (1 - P_R)\bar{V}_a - \bar{S}_A}{\sqrt{2}\sigma_Q}\right).\end{aligned} \tag{A8}$$

Another important difference in the alternative PBE2 is again due to ε_{A1n} and $\varepsilon_{A1c:c}$ not being constrained to equal zero—which induces correlation where in the original PBE1 there is independence:

$$\begin{aligned}\tilde{P}'_C &= \Pr(P'_{R|C}(\bar{W}_A + \varepsilon_{A3f} \vee \bar{a} + \varepsilon_{A3b}) + (1 - P'_{R|C})(\bar{V}_a + \varepsilon_{A1c:c}) > \bar{S}_A + \varepsilon_{A1n}) \\ &= 1 - \Pr\left(C'^* - \bar{W}_A > \varepsilon_{A3f} + \left(\frac{(1 - P'_{R|C})\varepsilon_{A1c:c} - \varepsilon_{A1n}}{P'_{R|C}}\right),\right. \\ &\qquad\qquad \left. C'^* - \bar{a} > \varepsilon_{A3b} + \left(\frac{(1 - P'_{R|C})\varepsilon_{A1c:c} - \varepsilon_{A1n}}{P'_{R|C}}\right)\right) \\ &= 1 - \Phi_2\left(\frac{C'^* - \bar{W}_A}{\sigma_2\sqrt{1 + \omega^2/P'^2_{R|C}}}, \frac{C'^* - \bar{a}}{\sigma_2\sqrt{1 + \omega^2/P'^2_{R|C}}}, \frac{\omega}{\omega + P'^2_{R|C}}\right).\end{aligned} \tag{A9}$$

Again, this derivation follows immediately from Theorem 3.2.1 of Tong (1990).

The difference between the prior and posterior probabilities measures the amount of updating,

$$\begin{aligned}&P_{F|C} - P_F \\ &= \Phi_2\left(\frac{\bar{W}_A - \bar{a}}{\sigma_1\sqrt{2}}, \frac{\bar{W}_A - C^*}{\sigma_1}, \frac{1}{\sqrt{2}}\right)\Big/\tilde{P}_C - \Phi\left(\frac{\bar{W}_A - a}{\sigma_1\sqrt{2}}\right)\end{aligned} \tag{A10}$$

$$\begin{aligned}&P'_{F|C} - P_F \\ &= \Phi_2\left(\frac{\bar{W}_A - \bar{a}}{\sigma_2\sqrt{2}}, \frac{\bar{W}_A - C'^*}{\sigma_2\sqrt{1 + \omega/P'^2_{R|C}}}, \frac{1}{\sqrt{2[1 + \omega/P'^2_{R|C}]}}\right)\Big/\tilde{P}'_C - \Phi\left(\frac{\bar{W}_A - a}{\sigma_2\sqrt{2}}\right).\end{aligned} \tag{A11}$$

References

Aragones, Enriqueta, and Thomas R. Palfrey. 2004. "The Effect of Candidate Quality on Electoral Equilibrium: An Experimental Study." *American Political Science Review* 98:77–90.

Carson, Jamie L. 2003. "Strategic Interaction and Candidate Competition in U.S. House Elections: Empirical Applications of Probit and Strategic Probit Models." *Political Analysis* 11:368–80.

Carson, Jamie L. 2005. "Strategy, Selection, and Candidate Competition in U.S. House and Senate Elections." *Journal of Politics* 67:1–28.

Leblang, David. 2005. "To Devalue or to Defend? The Political Economy of Exchange Rate Policy." *International Studies Quarterly* 47:533–559.

Lewis, Jeffrey, and Kenneth Schultz. 2003. "Revealing Preferences: Empirical Estimation of a Crisis Bargaining Game with Incomplete Information." *Political Analysis* 11:345–67.

McKelvey, Richard, and Thomas Palfrey. 1998. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics* 1:9–41.

Quinn, Kevin, and Anton Westveld. 2004. "Bayesian Inference for Semi-parametric Quantal Response Equilibrium Models." Paper presented at the 2004 Midwest Political Association meetings, Chicago, IL.

Signorino, Curtis S. 1999. "Strategic Interaction and the Statistical Analysis of International Conflict." *American Political Science Review* 93:279–297.

Signorino, Curtis S. 2003. "Structure and Uncertainty in Discrete Choice Models." *Political Analysis* 11: 316–344.

Signorino, Curtis S., and Kuzey Yilmaz. 2003. "Strategic Misspecification in Discrete Choice Models." *American Journal of Political Science* 47:551–66.

Tong, Y. L. 1990. *The Multivariate Normal Distribution.* New York: Springer-Verlag.