Mutual knowledge rationality among international policy elites*

Brad L. LeVeck
Department of Political Science, University of California, Merced, CA, USA
correspondence: bleveck@ucmerced.edu

In game theoretic experiments, the predictions made by Nash equilibrium sometimes fail, either because players choose irrational strategies or lack mutual knowledge of each other's rationality. However, when game theoretic experiments are used to study international interactions, researchers worry that decision-making differs between international policymakers and the convenience samples typically used in behavioral experiments. If so, the results of experiments using convenience samples may not generalize to decision-making by elite policymakers. In this research note, I examine whether international policy elites exhibit behavior that is consistent with both rationality and mutual knowledge rationality across a set of games. I find that, unlike convenience sample subjects, a high percentage of elites consistently act as if they are rational. However, very few elites consistently act as if they believe that all other elites are rational as well. So, when experimental results are driven by a lack of mutual knowledge rationality, the results may also apply to decision-making by international policy elites.

*Prepared for the 2018 Western Political Science Association Meeting in San Francisco, CA. Presented on March 30, 2018.

In empirical applications of game theory, the Nash equilibrium solution concept (along with its various refinements) remains the most widely used tool for deriving hypotheses about decision-makers' behavior. In essence, this represents a methodological bet that decision-makers' behavior will correspond to at least one of a game's Nash equilibria; and, the widespread use of the Nash equilibrium solution concept may stem from the fact that this bet can be justified on a number of different grounds. For instance, it may be justified by making specific assumptions about decision-makers' reasoning and beliefs (Aumann and Brandenburger, 1995; Crawford, 2016). This includes the well-known assumption that players are rational (Aumann and Brandenburger, 1995; Aumann, 1987; Bernheim, 1984). However, additional assumptions about players' knowledge and beliefs are usually needed as well (Aumann and Brandenburger, 1995; Aumann, 1987; Crawford, 2016). Alternatively, even if players are not rational, the Nash equilibrium may be justified via various assumptions about learning (Fudenberg and Levine, 1998; Sandholm, 2010). That is, if players repeatedly play the same game over time and receive feedback, they may learn by adjusting their strategy in the direction of higher payoffs (Traulsen et al., 2010; Hoffman et al., 2015). The Nash equilibria of a game then correspond to scenarios where there is no more learning to be done, as no individual player can improve their payoff via further adjustment (Sandholm, 2010; Fudenberg and Levine, 1998; Weibull, 1997).

However, like all bets, the Nash equilibrium solution concept does not always pay off, and may make the wrong behavioral prediction (Camerer, 2003; Crawford, 2016; Crawford, Costa-Gomes and Iriberri, 2013; Camerer and Fehr, 2006; Camerer, Ho and Chong, 2004; Healy, 2017). This remains true in international relations, where subjects in game theoretic experiments often fail to choose Nash equilibrium strategies (e.g. Tingley and Walter, 2011*a*,*b*; Tingley and Wang, 2010; LeVeck et al., 2014; Quek, 2016).

In such cases, researchers may wish to understand *why* the Nash equilibrium solution concept fails to predict behavior. This is particularly true when there is some reason to doubt that experimental results will hold outside of the lab. For example, many experiments in international relations use decision-making by undergraduates as a model system for studying decision-making by elite policymakers or states (e.g. Tingley and Walter, 2011*a*,*b*; Tingley and Wang, 2010; Quek, 2016; LeVeck and Narang, 2017; Renshon, Lee and Tingley, 2017; Tingley, 2011; Tingley and Wang, 2010; Rathbun, Kertzer and Paradis, 2017; Quek, 2017). This then raises the question: Do students (or other convenience samples) play non-equilibrium strategies for same reasons as international policymakers? If not, failures of Nash equilibrium in the lab may say very little about strategic decision-making by actors in international relations (Hafner-Burton, Hughes and Victor, 2013; Hafner-Burton et al., 2014; LeVeck et al., 2014; Mintz, Redd and Vedlitz, 2006). International relations scholars should therefore be interested in determining which causes of non-equilibrium play (if any) are conserved across convenience samples (such as undergraduates) and international policymakers.

In this research note, I focus on empirical applications where researchers use the Nash equilibrium solution concept because they are willing to make particular assumptions about players' reasoning and beliefs (i.e. they are not justifying the Nash equilibrium via assumptions about adaptive learning over time, a point I return to at the end of this research note). I then ask two questions: First, in cases where players' behavior does not correspond to a Nash equilibrium strategy profile of a particular game, which aspect of players' reasoning or beliefs is most responsible? Second, does this change when the players are international policymakers, rather than the populations used in most behavioral experiments, such as undergraduates or Amazon Mechanical Turk workers?

With regards to the first question, it should be noted that the definition of Nash equilibrium says nothing directly about players' reasoning or beliefs. It simply requires that each player's strategy maximizes her own payoff given the strategy of every other player (Nash, 1950). However, the Nash equilibrium can be justified by making specific assumptions about players' reasoning and beliefs. In particular, Aumann and Brandenburger (1995) reinterpret the Nash equilibrium as arising from a collection of beliefs that players hold about other players' chosen strategies. These beliefs, which are called conjectures, are probability distributions over other players' chosen strategy. For two player games, Aumann and Brandenburger show that players' conjectures lead to a Nash equilibrium if the following conditions hold: First, players are rational, meaning they maximize their utility given their beliefs about other players' actions.[1] Second, players' rationality, payoff functions, and conjectures are all mutually known. For n player games, where $n > 2$, players must also have common knowledge of other players' conjectures and have a common prior about other players' strategies.

Aumann and Brandenburger's result is incredibly useful because it defines sufficient conditions for Nash equilibrium. This means that if players are not at a Nash equilibrium of a game, one or more of the assumptions above must not be true. It therefore provides a checklist for explaining failures of Nash equilibrium in terms of players' reasoning and beliefs (Healy, 2017). Furthermore, all of these assumptions above are potentially testable (Healy, 2017). However, I can only test two of them using existing data on international policy elites; and even then, this is only possible via certain (hopefully reasonable) restrictions on players' beliefs. First, I test whether elites consistently act as if they are rational. Second, I test whether elites consistently act as if they believe that all other elites are rational. As shorthand, I refer to this belief as "mutual knowledge rationality."

Previous work suggests that humans regularly violate both the assumption of rationality and the assumption of mutual knowledge rationality. Some studies have found that subjects do not play best-responses to their own elicited beliefs about what other players will do, meaning that they do not act rationally (Costa-Gomes and Weizsäcker, 2008; Heinemann, Nagel and Ockenfels, 2004; Healy, 2017). Meanwhile, other studies have found that players choose a best

[1] This is equivalent to the concept of "rationalizability" (Bernheim, 1984), where players are not required to have correct beliefs about the strategies of other players.

response to their beliefs about other players strategies. However, these players' beliefs often underestimate the rationality of their opponents, meaning that they do not "know" other players are rational (Healy, 2017; Weizsäcker, 2003).

Yet, to date, all of these previous studies have been conducted on convenience samples of undergraduates or members of the general public. It therefore remains unknown whether deficits of rationality and mutual knowledge rationality similarly affect the strategic decision-making of international policymakers—who arguably make the majority of decisions in international relations between states.

Method and Results

To study the question of why elites sometimes fail to play Nash equilibrium strategies, I reanalyze data from LeVeck, Hughes, Fowler, Hafner-Burton and Victor (2014). This previous study used a unique sample of 102 international policy elites with an average of 21 years experience, largely drawn from the area of trade policy. This included 67 respondents with "high-level experience in government, including former members of the US House of Representatives, US Department of State, Treasury, and other agencies of government," 27 respondents with experience as "strategists within firms, frequently tasked with implementing the provisions of regulatory policy," as well as 8 subjects from "policy think tanks and nongovernmental organizations tasked with consulting government on trade and energy policy" (LeVeck et al., 2014, pg 18537). In sum, this was a set of decision-makers whose jobs led them to deal with various aspects of international policy on a regular basis.

As part of their study, LeVeck et al. had 95 of these elites play a series of p-beauty contest games, which went as follows: five players guess an integer number from 0 to 100. A target number t is then constructed by taking the average of all players' guesses and multiplying that average by a parameter p (in the most well-known variant of the game, $p = 2/3$). The player with a guess closest to the target number t earns a payoff of 100 monetary units (mu). All other players earn 0 mu. If multiple players choose the same winning number, those players split 100 mu equally. All elites knew they were playing other elites who had substantial professional experience with international policy-making.

LeVeck et al. (2014) also had two other convenience samples play the same p-beauty contest games. One sample consisted of 132 undergraduates at the University of California, San Diego. The other sample used 1007 subjects from the online labor market Amazon Mechanical Turk (AMT). These samples allow me to compare elite behavior with the behavior of subjects that are typically used in behavioral experiments.

Rationality

Three of the p-beauty contest games used by LeVeck et al. (2014) can be reanalyzed to measure whether players' act as if they are consistently rational. To fix ideas, consider the game in Figure 1A, where the multiplier p is 2/3. If all player choose the highest number, 100, the target number t cannot be higher than $66\frac{2}{3}$. Therefore, all strategies > 67 are "weakly dominated," meaning that there is always another number less than or equal to 67 that earns the same payoff or better. So, no matter what a player believes about other players' strategies, they have no incentive to play any number above 67. It therefore seems intuitive that rational players—who best respond to their beliefs—will always avoid these weakly dominated strategies. However, this intuition is not quite correct because there are some beliefs that will leave players perfectly indifferent between playing a weakly dominated strategy and playing a non weakly dominated strategy.

To definitively rule out weakly dominated strategies in this game, we need one or more additional restrictions on players' beliefs. While there are a number of restrictions that might work, I choose to assume that a player's initial belief (prior) about other players' strategies is a non-degenerate probability distribution. By "non-degenerate," I mean that a player's belief places at

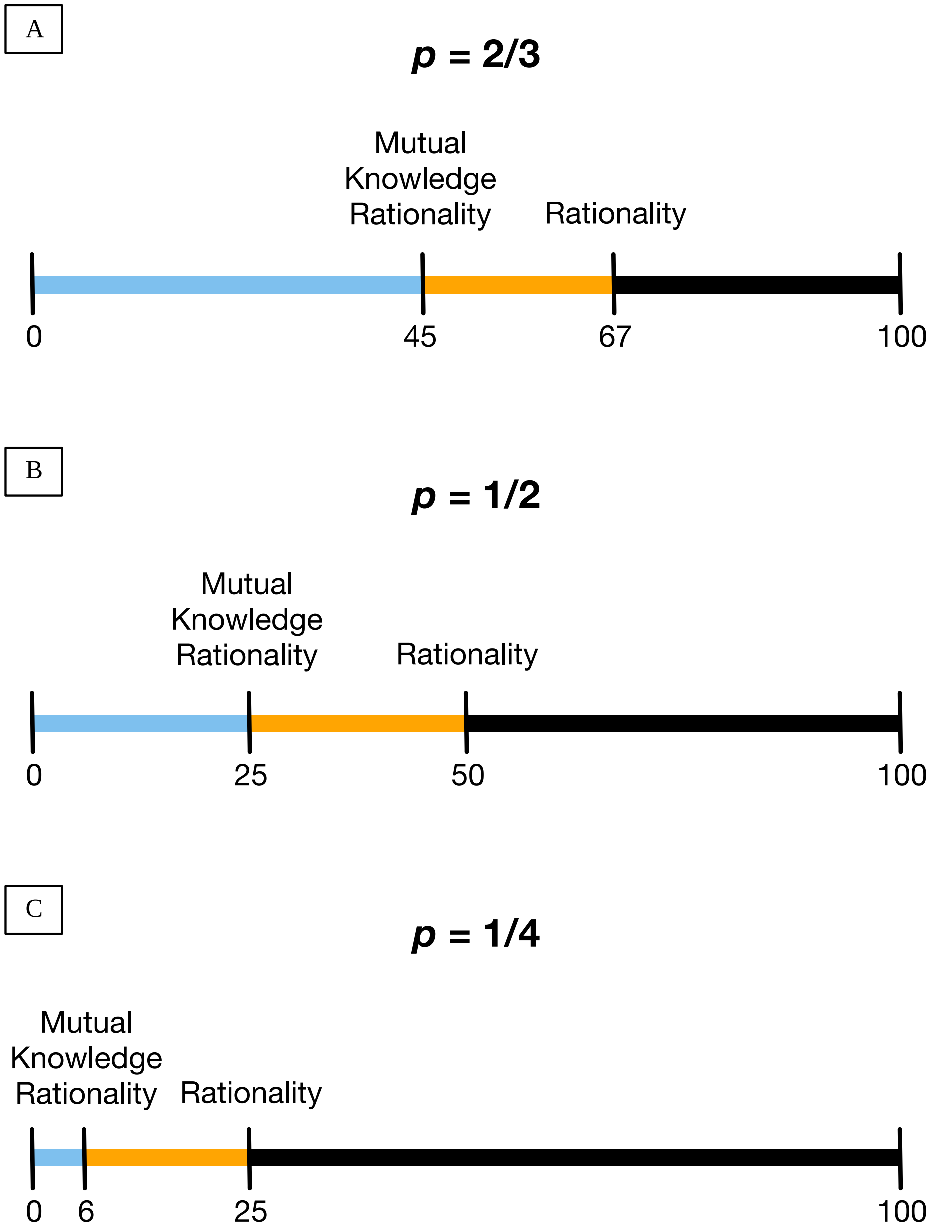

Figure 1: The three p-beauty contest games from LeVeck et al. (2014) use different multipliers p. This changes which choices are consistent with rationality (orange + blue), and mutual knowledge of other players' rationality (blue). **(A)** When $p = 2/3$, choosing any number ≤ 67 is consistent with rationality, while choosing any number ≤ 45 is consistent with mutual knowledge of other players' rationality. **(B)** When $p = 1/2$, choosing any number ≤ 50 is consistent with rationality, while choosing any number ≤ 25 is consistent with mutual knowledge of other players' rationality. **(C)** When $p = 1/4$, choosing any number ≤ 25 is consistent with rationality, while choosing any number ≤ 6 is consistent with mutual knowledge of other players' rationality.

least *some* non-zero probability on other players choosing any number in their strategy set.[2] This restriction on players' beliefs is fairly mild, as it still allows for a wide variety of distributions over other players' strategies. However, the assumption of non-degenerate priors implies that rational players will always try to pick the integer closest to the target number t.[3] Therefore, no rational player will play a weakly dominated strategy, as t cannot be in the set of weakly dominated strategies. In the Supplementary Information Appendix, I report an experiment that shows the assumption of non-degenerate priors plausibly applies to many individuals.

LeVeck et al. (2014) used three p-beauty contests that had sizable ranges of weakly dominated strategies, shown in Figure 1. I am therefore able to measure how often a player avoids choosing an irrational strategy across each of these three games, which considerably reduces the probability that a player is mistakenly categorized as consistently avoiding irrational strategies. For example, a player that chose numbers uniformly at random would only make a potentially rational choice in all three games $\sim 8\%$ of the time.

Figure 2A-C shows the results of this analysis for the three samples used by LeVeck et al. (2014): elites, undergrads, and AMT workers. Among both undergraduates and AMT workers, 36% and 43% of subjects played a potentially rational strategy across all three games. This result is similar to previous studies of rationality in laboratory games (Agranov, Caplin and Tergiman, 2015; Costa-Gomes and Weizsäcker, 2008; Nagel, 1995; Bosch-Domènech et al., 2002; Coricelli and Nagel, 2009). By contrast, 69% of elites played a potentially rational strategy in all three games. This result suggests that it is a reasonable bet to assume that international policy elites act rationally most of time.

[2] Another possible restriction is to require players to place > 0 probability on other players picking any number ≤ 67.

[3] See the Supplementary Information Appendix for a proof of this point.

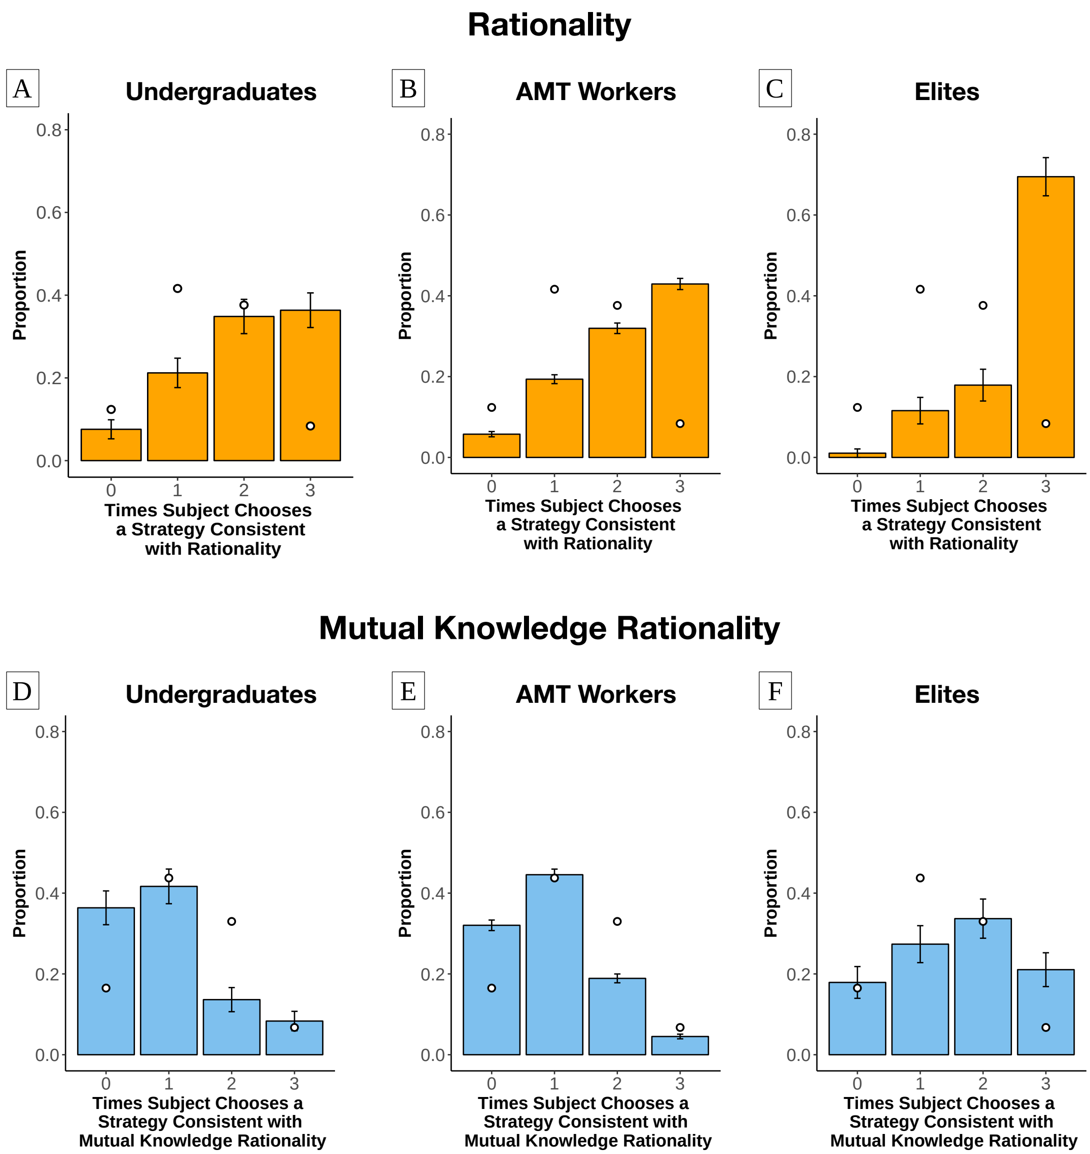

Figure 2: The proportion of players who—across three p-beauty contest games—make 0, 1, 2, or 3 choices that are consistent with rationality (A-C), and mutual knowledge of other players' rationality (D-F). Error bars correspond to ± 1 SEM. Open circles show the expected proportion of choices if players choose numbers uniformly at random. **(A-B)** Undergraduates and AMT workers choose a potentially rational strategy across all three games 36% and 43% of the time. **(C)** Meanwhile, elites choose a potentially rational strategy substantially more often, 69% percent of the time. **(F)** Elites choices are consistent with mutual knowledge rationality across all three games 21% of the time, which is more than twice as often as **(D)** undergraduates (8%) and **(E)** AMT workers (5%). However, this result still indicates that few elites consistently act as if they believe other elites are rational.

Mutual Knowledge Rationality

If rational players avoid weakly dominated strategies in the p-beauty contest game and assume other rational players do the same,[4] then we can also identify behavior that is consistent with mutual knowledge rationality. This is because players who obey mutual knowledge rationality would then rule out additional strategies, beyond those ruled out by players who are merely rational. For example, consider the game where $p = 2/3$. If no rational player will choose a number greater than 67, then the target number t cannot be greater than $44\frac{2}{3}$. Therefore, all strategies greater than 45 become weakly dominated; and, these numbers will be avoided by players who are both rational and assume other players are rational as well.

Figure 2F shows that 21% of elites behaved consistently with mutual knowledge rationality across all three games. Meanwhile, Figure 2D-E shows that only 8% of undergraduates and 5% of AMT workers behaved consistently with mutual knowledge rationality across all three games. These latter two proportions (5% and 8%) are not significantly higher than what one would expect if players chose numbers uniformly at random (7%, as shown in Figure 2E-F). So, compared to the types of convenience samples used in behavioral experiments, international policy elites are substantially more likely to play strategies consistent with mutual knowledge rationality. However, it is also true that 21% is not a particularly high number in absolute terms. Therefore, this result casts some doubt on any theory or hypothesis that requires international decision-makers to have mutual knowledge of each other's rationality.

Discussion

The above findings have a number of implications for how researchers in international relations interpret game theoretic experiments. Furthermore, they may inform how researchers choose

[4] To make this assumption, one must assume that any restrictions placed on rational players' beliefs (to rule out weakly dominated strategies) are also mutually known.

alternatives to the Nash equilibrium solution concept. After discussing both of these points, I end with three potential caveats.

The above results suggest that experimental researchers should exercise some caution when using convenience samples of undergraduates or online workers as a model system for international policymakers. Compared to international policy elites, these convenience samples are much more likely to choose irrational strategies; and, this may cause predictions derived via Nash equilibrium to fail. By contrast, irrationality should cause predictions to fail much less often among international policynakers.

On the other hand, the general lack of mutual knowledge rationality among elites suggests that many failures of Nash equilibrium in convenience samples should not be dismissed. Take, for example, the study by Tingley and Walter (2011*a*), who study cheap talk in entry-deterrence games played by undergraduates. Using a refinement of the Nash equilibrium (Sequential Equilibrium), Tingley and Walter hypothesize that players will ignore cheap talk from their opponent, but find the opposite. Players are highly influenced by cheap talk signals. This finding is plausibly driven by a lack of mutual knowledge rationality because responding to cheap talk is sometimes rational if you are uncertain about whether your opponent is rational. For example, Crawford (2003) presents a model where rational players respond to cheap talk because they believe their opponent irrationally reveals the truth some of the time. If Tingley and Walter's result is, in fact, driven by a lack of mutual knowledge rationality, the same result may also hold among international policymakers.

Furthermore, if researchers think that behavior in a previous experiment is primarily driven by a lack of mutual knowledge rationality, they can test this out with further experiments. While it can be difficult to change whether subjects act rationally in the lab, it may be relatively possible to change subjects' beliefs about the rationality of others. For example, researchers may provide subjects with suggestive evidence that their opponent is calculating and rational, as in Agranov, Potamites, Schotter and Tergiman (2012). Alternatively, subjects may play against robots who are programmed explicitly to rule out irrational strategies, as in Enemark, McCubbins and Turner (2016). Again, if Nash equilibrium primarily fails because players do not believe all other players are rational, there is little reason to believe the result would necessarily differ among international policymakers.

The results of this study may also inform how researchers choose alternatives to the Nash equilibrium solution concept. In recent decades, researchers in experimental economics and political science have formulated a number of different solution concepts, which may be psychologically more plausible. However, these alternative concepts often differ substantially in their assumptions. For example, Quantal Response Equilibrium (QRE), relaxes the assumption that players' are always rational (McKelvey and Palfrey, 1995, 1998). It instead assumes that players probabilistically choose strategies that are not a best responses to their beliefs. However, QRE retains the assumption that players' beliefs about other players' strategies are correct and mutually consistent. Meanwhile, concepts such as Level-*k* (LK) or Cognitive Hierarchy (CH) maintain the assumption that players' rationally play best responses to their beliefs, but relax the assumption that players' beliefs are correct—allowing players to systematically underestimate the rationality of others (Camerer, Ho and Chong, 2003; Crawford, Costa-Gomes and Iriberri, 2013; Camerer, Ho and Chong, 2003).

In recent years, QRE has been used by a number of applied researchers in international relations (Esarey, Mukherjee and Moore, 2008; Signorino, 1999; Signorino and Yilmaz, 2003; Gent, 2007; Tingley and Wang, 2010). However, the results here suggest that QRE may not solve the main problem with strategic decision-making by international policymakers. QRE relaxes the assumption of rational play, but international policy elites in this study choose potentially rational strategies most of the time. Meanwhile, QRE keeps the assumption that players have correct and mutually consistent beliefs about other players' strategies. Yet, many elites in this study act as if they have incorrect and mutually inconsistent beliefs about other players strategies. This seems to be the result of many elites systematically underestimating the rationality of other elites. Therefore, "non-Nash" behavior by elites may be better explained by concepts like CH, LK, or even risk dominance, where players best respond to uniformly random play by their opponent.[5]

Before concluding, I will note three important caveats to this paper's findings. First, without additional experiments, it is impossible to know for certain whether the findings would change if subjects played different games. That is, if elites played games other than *p*-beauty contests, the distributions of rational and mutually rational play might change. However, previous studies have examined whether undergraduates play rational strategies in games other than the *p*-beauty contest. So, I can at least check whether the results for undergraduates in my study differ substantially from these previous studies. Costa-Gomes and Weizsäcker (2008) used a number of 3×3 normal form games, some of which were not dominance solvable. They find that about 1/2 of players do not best respond to their own elicited beliefs. Meanwhile, across a series of global games Heinemann, Nagel and Ockenfels (2009) find that only 37% of subjects best respond to their elicited beliefs. In this paper, undergraduates play potentially rational strategies across all three games 36% of the time (95% ci 32-44), while AMT workers do so 43% of the time (95% ci 40-46). Therefore, my results for undergraduates and AMT workers appear to be within the range of prior studies that used different games.

A second caveat is that this study compares populations of individual human decision-makers. However, many theories in international relations concern themselves with the decisions of states; and, state decisions are often made by aggregating the decisions of individual policy-makers (Powell, 2017). Theoretically, it is uncertain how such aggregation would affect rationality and mutual knowledge rationality. Aggregation can sometimes lead to groups that make

[5]See Section 3 of the Supplementary Information Appendix, for a further discussion of why the data in this study are generally incompatible with QRE.

fewer decision-making errors (Page, 2008; Hong and Page, 2012). So, aggregated decisions may be more rational. On the other hand, there is no theoretical guarantee that aggregating the decisions of even perfectly rational individuals will lead to rational collective decisions (Arrow, 1950; McKelvey, 1976).

Empirically, however, there is some work on how aggregation affects choices in strategic games, which may be informative. Most relevant to this study is work by Sutter (2005), who examines a variant of the *p*-beauty contest game where decisions are aggregated via deliberation. Sutter find that groups of four play somewhat lower numbers than individuals. However, this result is mostly driven by excluding weakly dominated strategies. Therefore, aggregation may be more likely to solve problems related to rationality, rather than mutual knowledge of other players' rationality. This would certainly be consistent with a "wisdom of crowds" logic, where aggregation helps random errors in strategic decision-making cancel out (LeVeck and Narang, 2017; Kugler, Kausel and Kocher, 2012).

A third caveat is that the results above mostly apply to experiments where Nash equilibrium is justified on the basis of "strategic thinking" — where researchers place sufficient restrictions on players' reasoning and beliefs to justify Nash equilibrium behavior (Crawford, 2016). As noted at the beginning of this research note, Nash equilibrium may also be justified on other grounds. For example, it could be justified by assuming that players iteratively learn from their own experience over time (Sandholm, 2010). And, there are many forms of iterative learning that will eventually lead to a Nash equilibrium strategy profile, regardless of players' rationality or belief in other players' rationality (Sandholm 2010). Therefore, if Nash equilibrium is justified on theses alternative grounds, the results of this research note may not apply.

In general, if Nash equilibrium is justified on the basis of iterative learning over time (rather than strategic thinking), experiments should be conducted differently. Namely, experiments should have participants play the game multiple times with feedback. This will give participants the opportunity to learn. Researchers should then analyze whether behavior converges to a Nash equilibrium strategy profile across these multiple repetitions. An example of this can be seen in Tingley (2011).

However, such iterative learning cannot be used to justify the Nash equilibrium concept in all games. As Richard Thaler noted in a recent address to the American Economics Association, many of the most important decisions in life afford few opportunities for learning (2016). For example, most people will (hopefully) save for their retirement only once (Thaler, 2016). One could imagine a number of similarly important, but rare choices in international relations. For example, war and crisis bargaining are rare events (King and Zeng, 2001), and may only occur a few times in the course of any policymaker's career. Furthermore, even if policymakers could go beyond their own experience and learn from the historical record, there may still be very few cases that are sufficiently analogous to the crisis faced by a particular decision-maker (King and Zeng, 2005). This dearth of cases may then leave decision-makers highly uncertain about the best course of action. Also, the lack of applicable cases may leave decision-makers uncertain about the expectations and likely actions of other states. In these scenarios, even rational leaders may lack the mutually consistent expectations needed to play their part of a Nash equilibrium strategy profile.

References

Agranov, Marina, Andrew Caplin and Chloe Tergiman. 2015. "Naive play and the process of choice in guessing games." *Journal of the Economic Science Association* 1(2):146–157.

Agranov, Marina, Elizabeth Potamites, Andrew Schotter and Chloe Tergiman. 2012. "Beliefs and endogenous cognitive levels: An experimental study." *Games and Economic Behavior* 75(2):449–463.

Arrow, Kenneth J. 1950. "A difficulty in the concept of social welfare." *Journal of political economy* 58(4):328–346.

Aumann, Robert and Adam Brandenburger. 1995. “Epistemic Conditions for Nash Equilibrium.” *Econometrica* 63(5):1161–1180.

Aumann, Robert J. 1987. “Correlated Equilibrium as an Expression of Bayesian Rationality.” *Econometrica* 55(1):1.

Bernheim, B Douglas. 1984. “Rationalizable Strategic Behavior.” *Econometrica* 52(4):1007–1028.

Bosch-Domènech, Antoni, José G. Montalvo, Rosemarie Nagel and Albert Satorra. 2002. “One, Two, (Three), Infinity, ... : Newspaper and Lab Beauty-Contest Experiments.” *The American Economic Review* 92(5):1687–1701.

Camerer, Colin. 2003. *Behavioral game theory: experiments in strategic interaction*. Princeton University Press.

Camerer, Colin F. and Ernst Fehr. 2006. “When Does "Economic Man" Dominate Social Behavior?” *Science* 311(5757):47–52.

Camerer, Colin F., Teck-Hua Ho and Juin-Kuan Chong. 2004. “A Cognitive Hierarchy Model of Games*.” *Quarterly Journal of Economics* 119(3):861–898.

Camerer, Colin, Teck Ho and Kuan Chong. 2003. “Models of Thinking, Learning, and Teaching in Games.” *The American Economic Review* 93(2):192–195.

Coricelli, Giorgio and Rosemarie Nagel. 2009. “Neural correlates of depth of strategic reasoning in medial prefrontal cortex.” *Proceedings of the National Academy of Sciences* 106(23):9163–9168.

Costa-Gomes, Miguel A. and Georg Weizsäcker. 2008. “Stated Beliefs and Play in Normal-Form Games.” *The Review of Economic Studies* 75(3):729–762.

Crawford, Vincent P. 2003. “Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions.” *The American Economic Review* 93(1):133–149.

Crawford, Vincent P. 2016. “New Directions for Modelling Strategic Behavior: Game-Theoretic Models of Communication, Coordination, and Cooperation in Economic Relationships.” *Journal of Economic Perspectives* 30(4):131–50.

Crawford, Vincent P., Miguel A. Costa-Gomes and Nagore Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." *Journal of Economic Literature* 51(1):5–62.

Enemark, Daniel P, Mathew D. McCubbins and Mark B Turner. 2016. "Nashbots: How Political Scientists Have Underestimated Human Rationality, and How to Fix It.".

Esarey, Justin, Bumba Mukherjee and Will H. Moore. 2008. "Strategic Interaction and Interstate Crises: A Bayesian Quantal Response Estimator for Incomplete Information Games." *Political Analysis* 16(3):250–273.

Fudenberg, Drew and David K. Levine. 1998. *The Theory of Learning in Games*. The MIT Press.

Gent, Stephen E. 2007. "Strange bedfellows: The strategic dynamics of major power military interventions." *The Journal of Politics* 69(4):1089–1102.

Hafner-Burton, Emilie M, Brad L LeVeck, David G Victor and James H. Fowler. 2014. "Decision Maker Preferences for International Legal Cooperation." *International Organization* 68(04):845–876.

Hafner-Burton, Emilie M, D Alex Hughes and David G Victor. 2013. "The Cognitive Revolution and the Political Psychology of Elite Decision Making." *Perspectives on Politics* 11(02):368–386.

Healy, Paul J. 2017. "Epistemic experiments: Utilities, beliefs, and irrational play." *Working Paper* .

Heinemann, Frank, Rosemarie Nagel and Peter Ockenfels. 2004. "The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information." *Econometrica* 72(5):1583–1599.

Heinemann, Frank, Rosemarie Nagel and Peter Ockenfels. 2009. "Measuring Strategic Uncertainty in Coordination Games." *Review of Economic Studies* 76(1):181–221.

Hoffman, Moshe, Sigrid Suetens, Uri Gneezy and Martin A. Nowak. 2015. "An experimental investigation of evolutionary dynamics in the Rock-Paper-Scissors game." *Scientific Reports* 5:8817.

Hong, Lu and Scott E. Page. 2012. Some Micro-foundations of Collective Wisdom. In *Collective Wisdom: Principles and Mechanisms*, ed. Hélène Landemore and Jon Elster. Cambridge University Press.

King, Gary and Langche Zeng. 2001. "Logistic Regression in Rare Events Data." *Political Analysis* 9(2):137–163.

King, Gary and Langche Zeng. 2005. "The Dangers of Extreme Counterfactuals." *Political Analysis* 14(2):131–159.

Kugler, Tamar, Edgar E Kausel and Martin G. Kocher. 2012. "Are groups more rational than individuals? A review of interactive decision making in groups." *Wiley Interdisciplinary Reviews: Cognitive Science* 3(4):471–482.

LeVeck, Brad L, D Alex Hughes, James H. Fowler, Emilie Hafner-Burton and David G Victor. 2014. "The role of self-interest in elite bargaining." *Proceedings of the National Academy of Sciences of the United States of America* 111(52):18536–18541.

LeVeck, Brad L and Neil Narang. 2017. "The Democratic Peace and the Wisdom of Crowds." *International Studies Quarterly* 61(4):867–880.

McKelvey, Richard D. 1976. "Intransitivities in multidimensional voting models and some implications for agenda control." *Journal of Economic Theory* 12(3):472 – 482.

McKelvey, Richard D and Thomas R Palfrey. 1995. "Quantal response equilibria for normal form games." *Games and economic behavior* 10(1):6–38.

McKelvey, Richard D and Thomas R Palfrey. 1998. "Quantal response equilibria for extensive form games." *Experimental economics* 1(1):9–41.

Mintz, Alex, Steven B. Redd and Arnold Vedlitz. 2006. "Can We Generalize from Student Experiments to the Real World in Political Science, Military Affairs, and International Relations?" *Journal of Conflict Resolution* 50(5):757–776.

Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *The American Economic Review* 85(5):1313–1326.

Nash, John F. 1950. "Equilibrium points in n-person games." *Proc Nat Acad Sci USA* .

Page, Scott E. 2008. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies (New Edition)*. 1 ed. Princeton University Press.

Powell, Robert. 2017. "Research Bets and Behavioral IR." *International Organization* 71(S1):S265–S277.

Quek, Kai. 2016. "Are Costly Signals More Credible? Evidence of Sender-Receiver Gaps." *The Journal of Politics* 78(3):925–940.

Quek, Kai. 2017. "Rationalist experiments on war." *Political Science Research and Methods* 5(1):123–142.

Rathbun, Brian C, Joshua D Kertzer and Mark Paradis. 2017. "Homo diplomaticus: Mixed-method evidence of variation in strategic rationality." *International Organization* 71(S1):S33–S60.

Renshon, Jonathan, Julia J Lee and Dustin Tingley. 2017. "Emotions and the micro-foundations of commitment problems." *International Organization* 71(S1):S189–S218.

Sandholm, William H. 2010. *Population Games and Evolutionary Dynamics*. MIT Press.

Signorino, Curtis S. 1999. "Strategic interaction and the statistical analysis of international conflict." *American Political Science Review* 93(2):279–297.

Signorino, Curtis S and Kuzey Yilmaz. 2003. "Strategic misspecification in regression models." *American Journal of Political Science* 47(3):551–566.

Sutter, Matthias. 2005. "Are four heads better than two? An experimental beauty-contest game with teams of different size." *Economics Letters* 88(1):41–46.

Thaler, Richard H. 2016. "Behavioral Economics: Past, Present, and Future." *American Economic Review* 106(7):1577–1600.
URL: *http://www.aeaweb.org/articles?id=10.1257/aer.106.7.1577*

Tingley, Dustin. 2011. "The Dark Side of the Future: An Experimental Test of Commitment Problems in Bargaining." *International Studies Quarterly* 55:521–544.

Tingley, Dustin and Barbara Walter. 2011*a*. "Can Cheap Talk Deter? An Experimental Analysis." *Journal of Conflict Resolution* 55(6).

Tingley, Dustin and Barbara Walter. 2011*b*. “Reputation Building in International Relations: An Experimental Approach.” *International Organization* 65:343–365.

Tingley, Dustin H and Stephanie W. Wang. 2010. “Belief Updating in Sequential Games of Two-Sided Incomplete Information: An Experimental Study of a Crisis Bargaining Model.” *Quarterly Journal of Political Science* 5(3):243–255.

Traulsen, Arne, Dirk Semmann, Ralf D. Sommerfeld, Hans-Jürgen Krambeck and Manfred Milinski. 2010. “Human strategy updating in evolutionary games.” *Proceedings of the National Academy of Sciences* 107(7):2962–2966.

Weibull, Jörgen W. 1997. *Evolutionary game theory*. MIT press.

Weizsäcker, Georg. 2003. “Ignoring the rationality of others: evidence from experimental normal-form games.” *Games and Economic Behavior* 44(1):145–171.

Supplementary Information Appendix for: Mutual knowledge rationality among international policy elites

Brad L. LeVeck[1]*
[1]Department of Political Science, University of California, Merced, CA, USA
*correspondence: bleveck@ucmerced.edu

Contents

1 Proof that non-degenerate priors rule out weakly dominated strategies

Claim: If player i's prior belief places > 0 probability on other players $-i$ choosing any number in their strategy set, she will avoid numbers $> 100p$ rounded to the nearest integer.

Proof: Denote μ as player i's belief about the average number played by $n-1$ other players. Player i then expects the target number t to be

$$[\mu(n-1) + s_i]/n \times p \tag{1}$$

, where s_i is player i's chosen number.

If a player also believes she knows each other player's chosen number with certainty, then this belief can be described as a degenerate probability distribution over each other player's strategy. In this situation, the player may be indifferent between many different strategies. For instance, take the game where there are 5 players and $p = 2/3$. If player i believes that all other players choose 100 with certainty, she expects to win the game by playing any number from 99 to 10. Via equation 1, all of these numbers will leave player i closer to the target number t than the other four players choosing 100. It is therefore rational for player i to play any of those numbers.

However, if a player's belief is a non-degenerate probability distribution over other players' strategies, then playing the closest integer to t is the only way to guarantee a payoff > 0. All other strategies must have a lower expected payoff, as they cannot yield a higher payoff than picking the closest integer to t, and carry some risk that another player $-i$ picks a number closer to t (in which case, player i's payoff is 0). Therefore, a rational player would always play the integer that is closest to her expectation for t. Therefore, a rational player i will not pick strategies $> 100p$ rounded to the nearest integer, as this set of strategies cannot include t. ■

A similar logic would also apply to some other, weaker restrictions. For example, one might only require players to put > 0 probability on strategies $\leq 100p$ rounded to the nearest integer. This still creates a risk that some player $-i$ will play the strategy closest to t. Therefore, rational players will always wish to play a strategy closest to their expectation for t.

Also, if a player is trying to minimize her distance from t, she can do so by minimizing the expression $|s_i - [\mu(n-1) + s_i]/n \times p|$, rounded to the nearest integer. Generally, this will lead rational players to rule out additional strategies when n is finite, as player i's strategy affects the target number t. For, example in a game where $p = 2/3$ and $n = 5$, rational player trying to minimize her distance from t would never play a number above 62. This is because $|s_i - [100(4) + s_i]/5 \times 2/3|$ is minimized at $s_i \approx 61.54$. However, I ignore this complication, as it does not substantially change any of the empirical results in the paper.

2 An experiment on whether people have degenerate priors in the *p*-beauty contest

To test whether many players plausibly have non-degenerate beliefs over other players' strategies, I ran an experiment: 200 subjects from the online labor market Amazon Mechanical Turk (AMT) read the *p*-beauty game instructions from LeVeck et al. (2014).[1] They were then told (truthfully) that 1000 AMT workers had participated in a previous study where they played the *p*-beauty contest game. Subjects, were then given the following options: Option A, subjects could earn a 1 USD bonus. Option B, subjects could pick a number 0 to 100. If the number they picked was not chosen by any one of the participants in the previous study, the subject would earn 1.1 USD. If the number was chosen by at least one of the previous participants, the subject would earn only 0.1 USD.

If a subject had a degenerate prior over strategies, and placed 0 probability on players choosing a particular number, then they should always take Option B and pick a number they believe no one will play. In fact, risk neutral subjects should choose Option B if they place less than 10% probability on players picking a particular number.[2] So, the experiment was stacked in favor of finding evidence consistent with degenerate priors.

Of the 200 subjects, only 18 (9%) took Option B. In other words, the vast majority subjects did not even act as if they *might* have degenerate priors over other players' strategy.

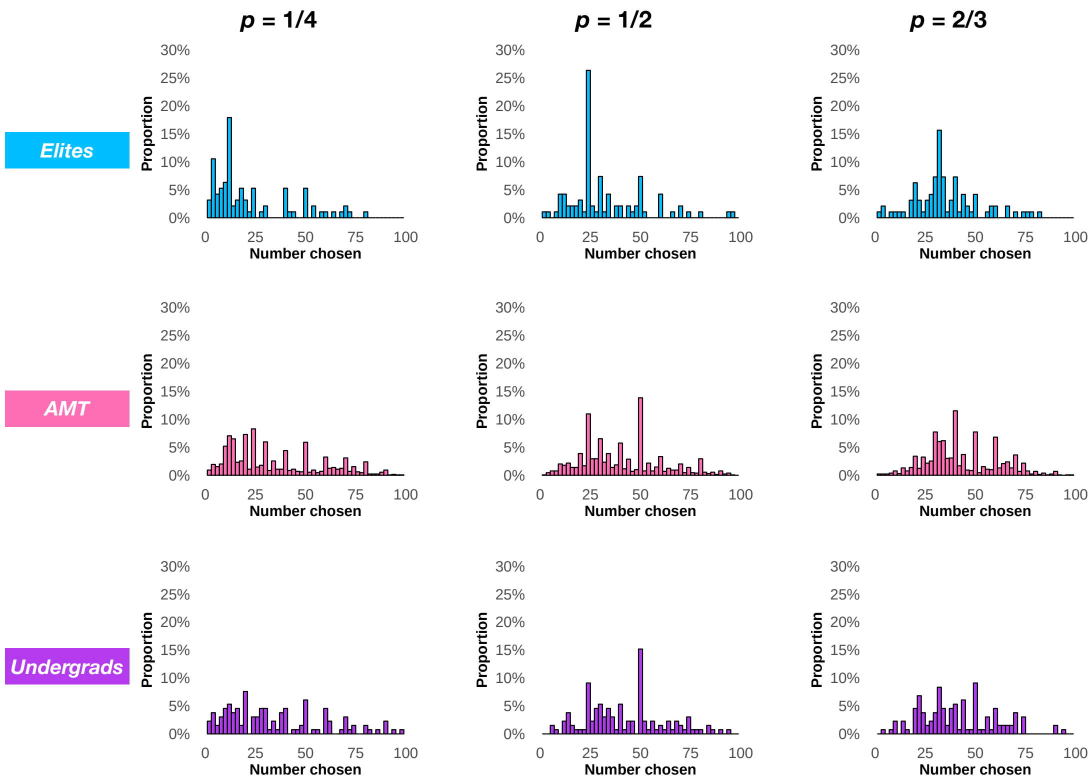

Figure A1: Histograms of player strategies in the p-beauty contest games from LeVeck et al. (2014). Columns correspond to one of three games, each with a different multiplier p. Rows correspond to one of three different populations: elites, AMT workers, and undergraduates.

3 Histograms of player strategies

Figure A1 shows histograms of each population's strategy for each of the games reported in the main manuscript. I present these histograms because some readers will likely be interested in seeing what happens in the raw data.

The one commonality across all three populations is that few players play the Nash equilibrium strategy of 0. There are also visible "spikes" in the data across all games. Many other studies of the p-beauty contest have found similarly spikey data (Nagel, 1995; Camerer, 2003;

[1]They were given the instructions for the game where the multiplier $p = 2/3$.

[2]Risk preferences should have zero effect on whether players with degenerate priors take Option B, as degeneracy implies 100% certainty that other players do not pick a set of numbers.

Bosch-Domènech et al., 2002). The spikey nature of the data may be explained by structural theories of non-equililbrium reasoning such as cognitive hierarchy, or Level-k reasoning, where players play a finite number of best-responses to an irrational player who chooses numbers at random (Nagel, 1995; Camerer, 2003; Crawford, Costa-Gomes and Iriberri, 2013). Spikey data may also be well explained by players who iteratively delete dominated strategies for a limited number of iterations (Crawford, Costa-Gomes and Iriberri, 2013). The spikes cannot, however, be well explained by equilibrium concepts, such as QRE or Nash equilibrium + noise, as these theories predict a smooth distribution centered around lower numbers (Crawford, Costa-Gomes and Iriberri, 2013).

For elites, first consider the game where $p = 1/4$. Spikes around 6 and 12 are consistent with a mixture of reasoning by players. Some players may iteratively delete dominated strategies twice, leading them to play strategies ≈ 6. For shorthand, label these player D2. Others players may anchor their beliefs at 50 and play a best response, leading the player to choose numbers around 12 or 13. This, second type of player then acts as a L1 player in Level-k models of strategic reasoning (Nagel, 1995; Crawford, Costa-Gomes and Iriberri, 2013). In Level-k models, L1 players best respond to an irrational L0 type that uniformly picks numbers at random (playing 50 on average). These L1 players would be categorized by my measure as rational, but their strategy would not be categorized as being consistent with mutual knowledge rationality. This seems appropriate since these L1 players are assumed to be best responding to a random opponent, who may irrationally choose dominated strategies. Meanwhile the D2 players would be categorized as obeying mutual knowledge rationality according to my measure.

For elites, the histogram where $p = 1/2$ is also consistent with the idea that many elites act as a mixture of D2 and L1 players. The large spike at 25 likely exists because this will be chosen by both D2 and L1 players. D2 players eliminate all strategies above 50, and then eliminate all strategies above 25. L1 players best respond to 50 by choosing 25. However, this mixture of D2 and L1 strategies is less apparent in the histogram where $p = 2/3$. For elites, there is a very pronounced spike around 33, which L1 players would pick if they best responded to 50. However, there is no pronounced spike around 47, which D2 players would pick.

Compared to elites, AMT workers and undergraduates appear to be characterized by a different mixture of player types. Some players appear to engage in one round of eliminating dominated strategies. This is especially clear in the AMT data where there is a cluster around 25 in the game where $p = 1/4$, a spike around 50 in the game where $p = 1/2$, and a spike around 67 in the game where $p = 2/3$. Meanwhile, other players seem to act more as L1 types. Again, this is clearest in the AMT data where there is a cluster around 12 in the game where $p = 1/4$, a spike around 25 in the game where $p = 1/2$, and a spike around 33 in the game where $p = 2/3$.

4 Level-k models' relation to rationality and mutual knowledge rationality

LeVeck et al. (2014) studied behavior in this games by fitting Level-k models to subjects' behavior. According to these models, players begin with with a specific conjecture about other players' strategies. The most common assumption is that players anchor their beliefs in a L0 player who chooses their strategy uniformly at random. Players are then assumed to recursively think through a limited number of k best-responses.

While these Level-k models are useful for many purposes, the fitted parameters from these models cannot clearly distinguish how rationality (a reasoning process) and mutual knowledge rationality (a belief about the reasoning of others) separately affect players' strategies. This is because a fitted parameter in these models may correspond to many different combinations of reasoning and beliefs about other players' reasoning (Agranov et al., 2012).

Furthermore, Level-k models require researchers to make highly specific assumptions about players' initial beliefs in a particular game. This limits the extent to which results from a particular experiment are portable to other games where players may start with different initial conjectures (Georganas, Healy and Weber, 2015). By contrast, one must make much weaker assumptions about players' beliefs in order to justify the current paper's measure of rationality. As shown in Section 1 above, treating weakly dominated strategies as irrational can be justified by many different belief distributions. This should give us some confidence that the assumption holds for a wider number of decision-makers. This, in turn, should ensure more confidence in my measure of mutual knowledge rationality, which requires rational players to place a similar restriction on other rational players' beliefs.

References

Agranov, Marina, Elizabeth Potamites, Andrew Schotter and Chloe Tergiman. 2012. "Beliefs and endogenous cognitive levels: An experimental study." *Games and Economic Behavior* 75(2):449–463.

Bosch-Domènech, Antoni, José G. Montalvo, Rosemarie Nagel and Albert Satorra. 2002. "One, Two, (Three), Infinity, ... : Newspaper and Lab Beauty-Contest Experiments." *The American Economic Review* 92(5):1687–1701.

Camerer, Colin. 2003. *Behavioral game theory: experiments in strategic interaction*. Princeton University Press.

Crawford, Vincent P., Miguel A. Costa-Gomes and Nagore Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." *Journal of Economic Literature* 51(1):5–62.

Georganas, Sotiris, Paul J Healy and Roberto A Weber. 2015. "On the persistence of strategic sophistication." *Journal of Economic Theory* 159:369–400.

LeVeck, Brad L, D Alex Hughes, James H. Fowler, Emilie Hafner-Burton and David G Victor. 2014. "The role of self-interest in elite bargaining." *Proceedings of the National Academy of Sciences of the United States of America* 111(52):18536–18541.

Nagel, Rosemarie. 1995. “Unraveling in Guessing Games: An Experimental Study.” *The American Economic Review* 85(5):1313–1326.