

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Kwon, Suehyun

Working Paper Behavioral Players in a Game

CESifo Working Paper, No. 7504

Provided in Cooperation with: Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Kwon, Suehyun (2019) : Behavioral Players in a Game, CESifo Working Paper, No. 7504, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at: https://hdl.handle.net/10419/198864

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU



Behavioral Players in a Game

Suehyun Kwon



Impressum:

CESifo Working Papers ISSN 2364-1428 (electronic version) Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute Poschingerstr. 5, 81679 Munich, Germany Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de Editor: Clemens Fuest www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- · from the SSRN website: <u>www.SSRN.com</u>
- from the RePEc website: <u>www.RePEc.org</u>
- from the CESifo website: <u>www.CESifo-group.org/wp</u>

Behavioral Players in a Game

Abstract

This paper points out issues with having behavioral players together with fully rational players in a game. One example of behavioral players is naive or sophisticated players; one can study higher-order beliefs when sophistication is the first-order belief, but the paper also considers alternative ways of modelling the type space and non-Bayesian updating. The paper shows that players must have heterogeneous priors and this type of heterogeneous priors cannot be justified by acquiring private information from the common prior. Furthermore, equilibrium definitions need to be modified for games with behavioral players.

Keywords: naivete, misspecified beliefs, heterogeneous priors, higher-order beliefs, equilibrium definition, Harsanyi doctrine.

Suehyun Kwon Department of Economics University College London 30 Gordon Street United Kingdom – London, WC1H 0AX suehyun.kwon@gmail.com

February 4, 2019 First draft: March 3, 2017

I thank Joel Sobel for comments.

1 Introduction

Earliest examples of behavioral types as "crazy" types who don't change their behavior and only play the same action every period go back to early folk theorem papers with incomplete information or reputation papers, see works of Fudenberg and Maskin in the 1980s. Bargaining papers also allow for behavioral types who only take the same action, e.g., Abreu and Gul (2000).

On the other hand, evolutionary game theory has taken a different view on behavioral types; these papers often focus on the evolution within a whole population when a proportion follows a rule of thumb or a behavioral rule. See for example works of Peyton Young.

The common feature of these papers is that behavioral types are characterized by either their action or their "rule of thumb", that even if the action itself is not fixed every single period, the choice of action in each period is fixed and known to rational players. On the other hand, mechanism design papers with behavioral consumers or agents also only focus on their actions and not their higher-order beliefs, see for example, Galperti (2015) and Farhi-Gabaix (2017).

We know from Rubinstein's email game (1986) that there is a discrepancy between finite higher-order beliefs and infinite higher-order beliefs, i.e., commom knowledge. For papers on level k, Crawford (2016) has a paper on mechanism design with level-k agents.

If one were to consider different types of "behavioral" players in different branches of economics, some focus on outcomes of having a few players whose action or decision rule is known; some focus on designing a mechanism when faced with such agents. And there are papers focusing on decision makers with such behavior, e.g., Heidhues et al (2018). When behavioral players are summarized by their actions or decision rules, they are auxiliary players in a sense that the focus of the paper is on what rational players or mechanism designers should do when faced with such players. With decision makers, the focus is on the implications of such behavior.

Now when one wants to ask seriously, how often do we observe these be-

havioral types in reality, one can ask whether these are heursitics that work with the law of large numbers and whether having a few such players in a finite game justifies such simplification. There is also a question of whether these behavioral players end up playing optimal strategies in a game, or due to their "fictitious play,"¹ rational players have to change their behavior in any equilibrium.

As a remark, I do not know of any paper in economics arguing something similar, but in physics, it has been known for a long time that one can only measure one of the following two, location and momentum.² This is genuinely a personal view that might turn out to be extremely unpopular, but in macroeconomics, or when one aggregates over a population of several millions if not billions, one could potentially think about rough heuristics that work on the basis of the law of large numbers. In recent years, particulary over the past 5 years or so, the number of macro papers with more microfoundation has multipled, but at the same time, the number of behavioral economics papers with better neuroscience foundation has also increased. Better microfoundation need not be purely for the purpose of characterizing better aggregate behavior; however, better microfoundation doesn't necessarily lead to a different aggregate behavior eventually, nor one can apply those heuristic behavior to a game with a handful of players.

Compared to the "behavioral" types mentioned so far, sophisticated players with time inconsistency are more than just the action. In this case, they are utility maximizers, just not the standard expected-utility maximizers; they even know their own utility function. However, if one were to start using the vocabulary "know," we need to discuss higher-order beliefs, common knowledge, level k, and so forth. One can think of the difference between the aforementioned behavioral types and time-inconsistent players as the difference between automata and players with something different from expected-utility maximizers with common knowledge.

 $^{^1\}mathrm{It}$ means something more specific in evolutionary game theory than playing the same action as I'm referring to here.

 $^{^2{\}rm This}$ is the Heisenberg uncertainty principle in quantum mechanics.

Previous two paragraphs point to the question of how to interpret "behavioral" players. Do we want them to be automata? Do we want them to match the law of large numbers? Do we want to explain certain behavior with psychological evidences? Why do we ignore game-theoretic foundations with these "behavioral" players? Particularly, in a game when one cannot just assume a continuum of players and the law of large numbers, do we only want to focus on rational players faced with the "crazy" types? But the argument doesn't work so well with naive or sophisticated players and neither with players with misspecified beliefs.

I do not provide decision-theoretic foundation in this paper. I also do not characterize implications with a large population. I focus on games with a finite number of players, some of which are behavioral. But I assume these players are naive, sophisticated or have misspecified beliefs instead of being automata. The rest of the paper discusses the type space, higher-order beliefs and equilibrium definitions for both the normal-form games and extensiveform games. I do not claim that my formulation or definition is the only way to accomodate these types of players; but the problems I point out are evident. Existing notions do not add up without any modification.

Before I describe the issues, please refer to O'Donoghue and Rabin (1999) and follow-ups for time inconsistency. Equilibrium definitions discussed in sections 3.3 and 3.4 involve belief updating just as with Perfect Bayesian equilibrium or sequential equilibrium and are related to Ortoleva (2012).

Consider a naive player with time inconsistency. A naive player doesn't believe that he is naive. However, if he plays a game with a fully rational player and the fully rational player believes his opponent is naive, the second-order belief of the naive player is "I believe you believe I'm naive." The naive player doesn't believe that he is naive yet at the same time he believes his opponent believes he's naive. This already shows that one cannot have common prior about the primitives of the model and naivete simultaneously, if one were to assume "time inconsistency" is a primitive. If one replaces "believe" with "know" in the first three sentences of this paragraph, it highlights why common knowledge of behavioral players' sophistication doesn't work. Exactly the same type of problem arises when one player believes in a misspecified model. This type of arguments does crucially depend on what is a primitive, what is the first-order belief and higher. I will be more concrete about the type space and (non-)Bayesian updating later.

I believe this is not the best place to explain equilibrium definitions for perfect Bayesian equilibrium, sequential equilibrium and so forth. However, among refinements, some focus on conditions on the strategies, and some focus more on off-the-equilibrium-path beliefs, i.e., beliefs after a probability-0 event. When Bayesian updating doesn't apply, what should be players' beliefs? Obviously, it won't happen on the equilibrium path, but it matters for characterizing what can happen on the equilibrium path. On a related note, issues with updating a joint distribution based on a single observation come up with fully rational players as well, and many issues in this paper can occur with any model with fully rational players if the prior on the set of states doesn't have full support.

When all players are fully rational, the discussion on common knowledge and common prior vs. heterogeneous prior goes back to the Harsanyi doctrine in the 1960s. I will discuss Aumann (1976) later in more detail; Weinstein and Yildiz (2007) discuss rationalizability and the universal type space.

The rest of the paper is organized as follows. Section 2 describes the model, and section 3 presents results. Section 4 concludes.

2 Model

The benchmark model has one behavioral player and one fully rational player. This allows to describe main ideas, but when there are multiple behavioral players and multiple fully rational players, one needs to model the type space and higher-order beliefs more carefully.

Two players are indexed by $I = \{1, 2\}$, and the set of states of the world is $\Omega \subseteq \mathbb{R}$. Players have priors over the states of the world in the beginning of the game which are denoted by F_1 and F_2 . The pdfs are denoted by f_1 and f_2 , respectively. Player 1 is fully rational, and F_1 has a positive density on the true state of the world, ω_0 . Player 2 has a misspecified belief and assigns probability 0 on the true state; $f_2(\omega_0) = 0$. This already assumes that the players don't have common prior in any model with misspecified beliefs. However, this formulation embeds naivete and sophistication in time consistency.

Consider time inconsistency. One could model the state of the world as time inconsistency of player 2, i.e., $\Omega = \{\omega_0, \omega_1\}$ with ω_0 = time inconsistent, ω_1 = fully rational. Player 1 assigns probability 1 on ω_0 , but a naive player 2 assigns probability 1 on ω_1 . If player 2 were sophisticated, he would have assigned probability 1 on ω_0 . This way, one can think of sophistication of a behavioral player as his first-order belief. In existing papers with behavioral agents, firms or the mechanism designer "know" the sophistication of the agent, but they don't consider higher-order beliefs.

This way of modelling also embeds misspecified models as in Heidhues, Koszegi and Strack (2018). They have an agent who assigns probability 0 on the true ability of the agent.

Player *i* has a set of actions $A_i \subseteq \mathbb{R}$, and payoffs are given by the vNM utility function $u_i(a_1, a_2)$. I will discuss both normal-form games and extensiveform games. Belief updating at the end of period in any repeated game or dynamic game will be discussed together with extensive-form games. Ω , timing of the game, sets of actions and payoffs are common knowledge.

The type space, initial priors and higher-order beliefs are discussed in section 3.

3 Results

This section presents main results of the paper. I will discuss the type space, initial priors and higher-order beliefs in turn; section 3.1 discusses why players must have heterogeneous priors and shows that this type of heterogeneous priors cannot be justified by private information obtained after players share the common prior. Section 3.2 discusses Aumann (1976) in this context. Section 3.3 and 3.4 discuss equilibrium definitions for normal-form games and extensive-form games, respectively. Issues with higher-order beliefs are discussed as they become relevant.

All results in this section are obvious with misspecified models; when it comes to time inconsistency, or naivete and sophistication, the modelling in section 2 is not the only way to model Ω . When there is more than one way of modelling the type space, this needs to be discussed together with issues of Bayesian updating, or more importantly, what to do when Bayesian updating doesn't apply. Essentially, one can fix one of a few things at hand, and why we should fix one instead of another needs further discussion. This is in section 3.5.

Until section 3.5, I will use the type space for naivete and sophistication as defined in section 2. Also for the purpose of illustration, I will refer to naive and sophisticated players with time inconsistency throughout the paper, instead of players with misspecified models.

3.1 Heterogeneous Priors

I first consider higher-order beliefs of players about the true state of the world. It is a common problem without full support, and the heart of the problem is that most behavioral models involving naivete or sophistication assign probability 0 on the true state of the world, thereby ruling out full support by assumption.

Whenever player 2 is naive or has a misspecified belief, player 2 assigns 0 probability on the true state of the world ω_0 . Player 1, on the other hand, is fully rational and assigns a strictly positive probability on ω_0 .

The problem arises with the second-order beliefs of players. Consider the time-inconsistent player 2; player 1 has the correct belief. The second-order belief of player 2 is that player 1's first-order belief assigns probability 1 on ω_0 . Since the first-order belief of player 2 assigns probability 1 on ω_1 , these beliefs are clearly at odds. We have the same problem with the second-order belief of player 1.

By assumption, if a behavioral player is naive or has a misspecified belief, two players have heterogeneous priors on the true state of the world. In addition, one cannot justify these first-order beliefs as a result of acquiring private information with some common prior π . If players acquired private information when they initially believed in π , Bayesian updating requires

$$f_1(\omega_0) = \frac{\Pr(s_1|\omega_0)\pi(\omega_0)}{\int \Pr(s_1|\omega)\pi(\omega)d\omega},$$

$$f_2(\omega_0) = \frac{\Pr(s_2|\omega_0)\pi(\omega_0)}{\int \Pr(s_2|\omega)\pi(\omega)d\omega}$$

where the signal structure is given by $s_i : \Omega \to S_i$ and player *i* received signal s_i .

Given that player *i* received s_i , one can assume without loss of generality that the denominators $\int \Pr(s_1|\omega)\pi(\omega)d\omega$, $\int \Pr(s_2|\omega)\pi(\omega)d\omega$ are strictly positive. Given that $f_1(\omega_0) > 0$, we cannot have $\pi(\omega_0) = 0$. However, we have $f_2(\omega_0) = 0$, and this requires $\Pr(s_2|\omega_0) = 0$. Player 2 received signal s_2 that cannot occur in ω_0 even though the true state of the world is ω_0 .

This already shows that any game with a behavioral player assumes that players have heterogeneous beliefs that cannot be justified by acquiring an informative signal from common prior; Harsanyi doctrine doesn't apply to games with behavioral players.

Theorem 1. If a naive player and a fully rational play a game, the fully rational player's first-order belief assigning a strictly positive probability on the sophistication of the naive player implies all of the following:

- 1. Initial priors don't have full support.
- 2. Players have heterogeneous priors.
- 3. Their priors cannot be justified by acquiring private information from the common prior.

3.2 Agreeing to Disagree

A bigger problem is when one considers an argument as in Aumann (1976). In particular, a behavioral player assigns probability 0 on ω_0 , and therefore, his belief will never change by learning the belief of the fully rational player. What about the fully rational player?

One option is that the fully rational player believes that the behavioral player got his belief for some reason that cannot be justified by any private information. If the fully rational player accepts this, then he cannot justify where the first-order belief of his opponent came from, and the second-order belief is in conflict with the first-order belief. However, this approach is in line with the heterogeneous-prior literature. Let me just comment that there is always an issue of whether one can rewrite the model as one with heterogeneous preferences, not heterogeneous priors. This is a long-standing debate, and I won't get into any further discussion here.

Things start to go even more awry when the fully rational player tries to justify the first-order belief of the behavioral player as coming from some type of private information he doesn't have. If at the beginning of the game, player 1 knows the first-order belief of player 2 and believes that it came from a private signal, then player 1 has to believe that he has a misspecified belief; either the common prior or the signal structure has to be wrong.

In any game with fully rational players and common knowledge about the primitives of the model, this type of problem doesn't happen. However, any behavioral player in the same game with a fully rational player requires that one cannot assume common knowledge about the primitives of the model anymore. One must have heterogeneous priors, then what part of the model can still be common knowledge across all players? How much must one relax the assumption of common knowledge to define a game with both behavioral players and fully rational players? And once one starts relaxing the common knowledge assumption, why one relaxes one part of the model and not the others? And what fully rational players should do after observing the action of behavioral player crucially depends on what is common knowledge across all players, and whether fully rational players assume the behavioral player has private information or is just "behavioral."

Furthermore, behavioral players may never update their beliefs about the true state of the world, but from fully rational players' perspective, if they are

willing to update their posterior beliefs, then I did define in section 2 that anything except initial priors and higher-order beliefs are common knowledge, but why must players have common knowledge regarding information structures and so forth? Or more importantly, even if information structures as primitives of the model can be common knowledge, but common knowledge breaks down quickly after the first period if the monitoring technology is imperfect public monitoring or private monitoring.

Theorem 2. In a static game with complete information about the naive player's first-order belief, either players agree to disagree, or the fully rational player knows that his prior is a probability-0 event.

Either with incomplete information or in dynamic games without perfect monitoring, the fully rational player needs to update his belief at the end of each period.

3.3 Normal-form Games

In a normal-form game, given the information structure and the belief hierarchy, each player maximizes their expected utility when they choose their actions. Issues from section 3.2 remain, but once the beliefs are taken as given, the players choose their actions simultaneously, and the game ends. It is not too different from other models of heterogeneous priors.

3.4 Extensive-form Games

This section discusses problems with equilibrium definition for extensive-form games, repeated games and other forms of dynamic games. Precedents of "behavioral types" in dynamic games only take the same action every period and don't choose their actions each period; see early folk theorem papers, reputation papers and bargaining papers. These behavioral types are more of an automata and don't take best responses. The focus of these papers are on strategies of (in most cases) single fully rational player.

However, any behavioral decision maker in the recent literature has a behavioral trait that can be axiomatized, i.e., some axioms are relaxed, but these decision makers still maximize their (expected) utility.

Once behavioral players in a game can choose their strategies, then it matters what these behavioral players optimize when they choose their strategies, and what needs to be taken into account simultaneously is what information these behavioral players have. As pointed out in section 3.2, in order for fully rational players playing with behavioral players to update their beliefs after observing a signal or an action profile, fully rational players need to have beliefs about behavioral players' information, higher-order beliefs etc.

Most papers with time-inconsistent players do not study extensive-form games. However, it is implicitly assumed that being a naive player or being a sophisticated player itself is a primitive of the model; this assumes both (i) sophistication doesn't change and (ii) being naive and being sophisticated are just two different states of the world with no relation to each other. This approach shuts down any role of naivete and sophistication in a game; what distinguishes this type of behavioral players from the rest is the "belief" of player about his own utility function. It need not be the utility function to have naive and sophisticated players. Any primitive regarding a player could form the basis for this distinction. However, if two types of players are treated as being in either of two states of the world, there's no room for studying what happens in extensive-form games when a player has a correct belief or incorrect belief about his utility function.

The rest of this section describes problems with equilibrium definition once some part of the model is no longer common knowledge; section 3.5 discusses different ways of modelling the type space and defining equilibrium together with non-Bayesian updating.

The main problem with any extensive-form games, repeated games or dynamic games is the prpbability-0 events, i.e., when a deviation is detected. With fully rational players, the equilibrium definition stipulates that across all refinements, whenever possible, update the belief on the type or the private information of the other player instead of assuming the opponent has deviated.

Consider the game with the naive player and the fully rational player from

section 2. After a signal or an action profile is observed at the end of the period, Bayesian inference and whether players believe they are on or off the equilibrium path depend not only on what's common knowledge but also on the monitoring technology. Even when all players are fully rational, the latter issue of whether players believe they are still on the equilibrium path does come up in repeated-game models with private monitoring, which I believe is part of the reason many papers have focused on belief-free equilibria.

Loosely speaking, if all players' belief on sophistication of behavioral players is 0 or 1 then we have the usual problems without full support, and there is no learning. Higher-order beliefs regarding sophistication of behavioral players never change from the beginning of the first period, and any probability-0 event will be attributed to strategies or different types of players, not sophistication of a player. If there are multiple types, whether learning will be complete in the limit is a separate question, but if all players' beliefs on sophistication are degenerate, then my model becomes isomorphic to the usual modelling choice of modelling sophistication itself as the state of the world. It still leads to problems, but in this case, it really is all due to the lack of full support and heterogeneous priors.

What if players' beliefs are nondegenerate and get updated over time? This is where my modelling choice makes a difference. If sophistication by itself is a primitive of the model, then players' higher-order beliefs never change. Once we model sophistication as the first-order belief, then behavioral players can learn whether they are naive or sophisticated. This type of learning is not allowed if one were to model naivete and sophistication as two different states of the world. Obviously, once we split sophistication into two levels of beliefs instead of one, one can also ask why time inconsistency has to be a state of the world that never changes. In a bigger picture, this is related to the discussion of heterogeneous preferences vs heterogeneous priors. It is also related to the question of what is really the state of the world that never changes. The following few paragraphs describe issues with learning in this type of games.

Once we allow learning, then any player needs to update the joint distribution on the type of behavioral player and his strategy when there is a deviation. But updating a joint distribution when the piece of information can be generated either by a different type or a deviation in strategy requires some discussion even when it is not a probability-0 event. The usual approach in PBE is to apply Bayes' rule whenever possible. It is the same with sequential equilibrium or any equilibrium definition for extensive-form games. This priority by itself is not a part of Bayesian inference. Equilibrium definitions require that when Bayes' rule applies, apply it first; do not attribute it to a deviation if there can be a type of a player who can lead to the signal or the outcome with a strictly positive probability. Bayesian inference also doesn't dictate any priority in how to update a joint distribution based on a single observation. These are all assumptions that have to be made for specific equilibrium notions. I will discuss Ortoleva (2012) in more detail in section 3.5.

If the fully rational player were to assume that a probability-0 event is due to a type of behavioral player, not a deviation, and update his belief by Bayes' rule, this leads the fully rational player to "learn" from the behavioral player and update his belief to a misspecified one, which is already pointed out in section 3.2. To be more concrete, if the behavioral player assigns probability 0 on the true state of the world, then his posterior never assigns a positive probability on the true state of the world. Depending on the number of players and the monitoring technology, one can ask whether the fully rational player's belief on the true state of the world will converge to 0, and in case there are multiple behavioral types, whether the fully rational player's second-order belief of the behavioral player's first-order belief will assign probability 0 on the true state eventually.

The fully rational player needs to play the best response to the behavioral player, which implies that if the behavioral player never assigns a positive probability on the true state of the world, the fully rational player needs to "accomodate" the behavioral player's belief and therefore his strategy. Regardless of the limit learning, there is a problem with usual equilibrium definitions in any game where a fully rational player plays against a behavioral player, and we need an appropriate notion of equilibrium (in particular, how to update a joint distribution given a single piece of information) in this type of games. However, as pointed out earlier, some of the issues can occur with fully rational players as long as there is no full support, and the distinction between the lack of full support and having naive or sophisticated players also needs to be made more precisely. One can also ask whether only modifying the first-order belief of the naive player is the best way to model sophistication, and in relation to the previous paragraph and this paragraph, can the fully rational player learn differently knowing that the behavioral type just assigns 0 probability on the true state of the world, even if the fully rational player doesn't know the sophistication of behavioral player with probability 1? Can the sophisticated player leverage this learning? Why would the time-inconsistent utility function itself be known, or be a primitive of the model? What happens if the behavioal player needs to learn whether he is naive or sophisticated?

In addition to the issues I have pointed out so far, (i) what if the fully rational player's belief on the sophistication of the behavioral player is degenerate and (ii) what if the belief is non-degenerate that the fully rational player can learn each period, another issue, which is more of a modelling choice, is what part of the primitives should be common knowledge. Compared to games with only fully rational players, the type of games I described in this paper have behavioral players. If one were to literally modify the common-knowledge assumption minimally, we will just assign probability 0 on the naive player's first-order belief on the state of the world and keep everything else the same as with fully rational players. However, once we allow for incomplete information, there is a difference between making an inference after a signal on the type and the strategy of fully rational player and making an inference of those of behavioral player. We already made modifications to the first-order belief of the behavioral player. With degenerate beliefs, nothing much can be done to belief hierarchy, which still leaves the question of modelling choice, but with incomplete information, do we want to only change the first-order belief of the naive player and keep everything else? It is a benchmark, but why only the first-order belief, and when there are different ways of modelling the type space in the first place, what do we want to get from the type space and belief hierarchy?

Furthermore, if the fully rational player assigns a positive probability on the behavioral player putting probability 0 on the true state of the world, in line with discussions of awareness, why does the fully rational player never doubts or reconsiders his information and the belief hierarchy? This is not allowed if we minimally change only the first-order belief of the naive player.

3.5 Remedies

This section discusses the type space and higher-order beliefs together with equilibrium definitions for extensive-form games. I will first describe what has been implicitly suggested in the existing literature albeit not in finite games then compare to the modelling choices I made in section 2.

As mentioned earlier, existing literature on naive or sophisticated players implicitly assume that these are two separate states of the world. The benchmark model in section 2 of this paper takes the utility function of the behavioral player, or time inconsistency itself, as the state of the world; a naive player and a sophisticated player are distinguished by their first-order beliefs. Any player with a misspecified model assigns probability 0 on the true state of the world.

In section 2, I assumed the naive player or the sophisticated player has a degenerate belief about the true state of the world but the fully rational player's beliefs don't have to. When the behavioral player's first-order belief is degenerate, the k-th order belief of the behavioral player and (k + 1)-th (or (k - 1)-th) order belief of the fully rational player can be compared to the implicitly assumed model in the existing literature for (k - 1)-th order belief of the behavioral player and k-th (or (k - 2)-th) order belief of the fully rational player, respectively. How much difference does it make to model sophistication as the first-order belief instead of the state of the world?

Unfortunately, or fortunately, if the behavioral player's first-order belief is degenerate, the difference between two modelling choices is not as big as when it is interior. The key difference here comes from the first-order belief and

the second-order belief of the fully rational player. But given the degenerate first-order belief of the behavioral player, the second-order belief of the fully rational player in the end is analogous to the first-order belief of the existing literature. So the key here is that now the fully rational player can learn whether his opponent is time inconsistent or not. This learning will be separate from learning about his opponent's first-order belief, whether his opponent believes he is naive or sophisticated; but if the fully rational player doesn't know if his opponent is time inconsistent in the first place, now he needs to consider the possibility that his opponent is fully rational as well. Formally, in section 2, ω_0 was being time-inconsistent, and ω_1 was being time-consistent. If the first-order belief of fully rational player assigns probability 1 on ω_1 , then his opponent assigning probability 1 on ω_1 coincides with his opponent being fully rational, up to the second-order belief. On the other hand, if the fully rational player assigns probability 1 on ω_0 , then his opponent is naive, up to the second-order belief. If being naive or sophisticated was the state of the world themselves, this difference due to the first-order belief of the fully rational player cannot happen. The way I modelled sophistication allows the fully rational player's first-order belief to be independent from that of the behavioral player, and therefore, being naive, sophisticated or fully rational can be accomodated all by the same type space. The only difference is literally the first-order beliefs of each player. One might argue this is how the literature views sophistication as well, but as pointed out earlier in this paragraph, it cannot be accomodated if sophistication is the state of the world. Leaving aside allowing for all possibilities, if beliefs are degenerate, they can never be updated by Bayes' rule, and unless one comes up with an equilibrium notion that changes higher-order beliefs after a probability-0 event instead of attributing it to a deviation, there is no learning over time.

Now if one were to allow for interior beliefs, one can consider both the fully rational player and his opponent to have a degenerate first-order belief or an interior first-order belief. This then allows for learning, and one now needs to think about how to update after a probability-0 event.

4 Conclusion

I studied game-theoretic foundations of finite games with behavioral players in this paper. Instead of automata-like "crazy" types, I focus on behavioral players with incorrect beliefs about themselves, or to be more precise, naive and sophisticated players. Same issues come up with players with misspecified models, and one should compare to dynamic games with fully rational players who have heterogeneous priors without full support.

I illustrate the issues with one naive player and one fully rational player, and they can be summarized as (i) initial priors without full support, (ii) heterogeneous priors that cannot be justified by private information, (iii) fully rational players should not try to learn from naive players' higher-order beliefs, (iv) learning in dynamic games.

Section 3.5 discusses the type space, belief hierarchy and non-Bayesian updating for the model in section 2 and compares to the implicitly assumed type space in the existing literature. Since assigning probability 1 on the naive player's first-order belief brings down the problem to the lack of full support with heterogeneous priors, learning issues arise only with interior beliefs.

References

- Abreu, Dilip, and Faruk Gul. 2000. "Bargaining and Reputation," *Econo*metrica, 68(1):85-117.
- [2] Ali, Nageeb. 2011. "Learning Self-Control," Quarterly Journal if Economics, 126:857-893.
- [3] Aumann, Robert, J. 1976. "Agreeing to Disagree," Annals of Statistics, 4(6):1236-1239.
- [4] Farhi, Emmanuel, and Xavier Gabaix. 2017. "Optimal Taxation with Behavioral Agents," working paper.

- [5] Galperti, Simone. 2015. "Commitment, Flexibility, and Optimal Screening of Time Inconsistency," *Econometrica*, 83(4):1425-1465.
- [6] Harsanyi, John, C. 1967. "Games with Incomplete Information Played by "Bayesian" Players, Parts I-III," *Management Science*, 14(3).
- [7] Heidhues, Paul, Botond Koszëgi, and Philipp Strack. 2018. "Unrealistic Expectations and Misguided Learning," conditionally accepted at Econometrica.
- [8] O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later" American Economic Review, 89(1): 103-124.
- [9] Ortoleva, Pietro. 2012. "Modeling the Change of Paradigm: Non-Bayesian Responses to Unexpected News," American Economic Review 102(6):2410-2436.
- [10] Rubinstein, Ariel. 1989. "The Electronic Mail Game: Strategic Behavior under "Almost Common Knowledge."" American Economoc Review, 79(3):385-391.
- [11] Weinstein, Jonathan, and Muhamet Yildiz. 2007. "A Structural Theorem for Rationalizability with Application to Robust Predictions of Refinements," *Econometrica*, 75(2):365-400.