

“Games with Time Inconsistent Players”

Yianis Sarafidis*

February 22, 2006

PRELIMINARY DRAFT

Comments are more than welcome

Abstract

We embed time inconsistent agents (players) in non-cooperative games. To solve such games, we introduce two solution concepts, which we refer to as *equilibrium* and *naive backwards induction*. When all players are sophisticated time inconsistent, these solution concepts are equivalent and coincide with the standard notions of SPNE and backwards induction (in finite games of perfect information). When some players are naive time inconsistent, however, these solution concepts (may, and most likely will) lead to different predictions. We apply these solution concepts on two well-known economic applications, the alternating-offers bargaining game and the durable good monopoly, and we explore their relative strengths and weaknesses.

*Yale University (visiting), Department of Economics, 37 Hillhouse Av., New Haven, CT 06520, e-mail:yianis.sarafidis@yale.edu

1 Introduction

The study of time inconsistent preferences is a classical topic in economics, dating back to the seminal work of Strotz (1956). The existing literature has analyzed time inconsistent preferences primarily in the context of (single-agent) decision problems. This is the case for both the early contributions of Strotz (1956), Phelps (1968) and Phelps and Pollak (1968), as well as for the more recent work of Laibson (1997), O'Donoghue and Rabin (1999; 2001), Carillio and Marriotti (2001) and others. In this paper we will embed time inconsistent agents (players) in games.¹

While bringing time inconsistent preferences into game theory is a reasonable theoretical exercise, there is an applied motivation. Time inconsistent agents do not only make decisions in isolation, but they also participate in markets, organizations, teams and other institutions together with other time (in)consistent agents. In other words, they play games. Consider, for example, a durable good monopolist who screens time inconsistent consumers through inter-temporal price discrimination. How does the presence of consumers with changing preferences affect prices, profits and consumer welfare in this classical Coasian (1972) problem? Or, consider Rubinstein's (1982) alternating-offers bargaining game. What agreement, if any, will time inconsistent players reach?

As we shall see, when all players are sophisticated time inconsistent there is no need to introduce any new game theory tools.² However, when some players are naive time inconsistent, it is not evident what the appropriate solution concept is. We will propose two solution concepts, which we will refer to as *equilibrium* and *naive backwards induction*. We have chosen these terms to emphasize the fact that the former is similar in spirit to subgame perfect Nash equilibrium (SPNE), whereas the latter relies on backwards

¹To prevent any misunderstandings, we would like to make two clarifications upfront. First, in a recent paper Della Vigna and Malmendier (2004) also discuss a game between a rational monopolist and a time inconsistent consumer. However, in their game the time inconsistent player moves once, after she has observed the move of the monopolist. Hence, from her point of view, this is still a decision problem. Second, single agent decision problems with sophisticated time inconsistent agents are modelled as games between the agent's current and future selves. As we shall see, games between two, or more, time inconsistent players give rise to a whole new set of issues that have not been explored before. These issues arise primarily in games with *naive* time inconsistent players.

²See Chade, Prokopovych and Smith (2005) for an analysis of repeated games with sophisticated time inconsistent players.

induction. In fact, when all players are sophisticated the two concepts are equivalent and coincide with the standard notions of SPNE and backwards induction (in finite games of perfect information). When some players are naive, however, these solution concepts (may, and most likely will) lead to different predictions.³

We also apply these solution concepts on two well-known economic applications, namely Rubinstein's alternating-offers bargaining game and the durable good monopoly. We show how our solution concepts give rise to different behavior and we discuss their pros and cons in the context of these two applications. We will argue that the social and economic context is important and each solution concept may be more or less appropriate depending on the particular application.

Finally, we also hope that the paper makes a more general contribution in exploring some issues that arise when one relaxes the standard common knowledge assumption. Chung and Ely (2005) refer to this general research program as the *Wilson Doctrine*, which is articulated by Wilson (1987) as follows:

“Game theory has a great advantage in explicitly analyzing the consequences of trading rules that presumably are really common knowledge; it is deficient to the extent that it assumes other features to be common knowledge such as one agent’s assessment about another’s preferences or information. [...] I foresee the progress of game theory as depending on the successive reduction in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeating weakening of common knowledge assumptions will the theory approximate reality.”

A game between a sophisticated and a naive time inconsistent player (see, for example, the game discussed in section 2) will give rise to the type of situation that Wilson alludes to. This is because in such a game the two players have different beliefs (including higher-order beliefs) about each other's preferences. We believe that the non-equivalence between subgame perfection and backwards induction that arises in games with naive time inconsistent players is likely to apply in other dynamic games where the assumption of common knowledge is relaxed.

³It will first be necessary to generalize the notions of sophistication and naivete to account for whether a player understands the changing nature of *other* players' preferences (and all other higher-order beliefs). In other words, we need to define *higher order sophistication and naivete*.

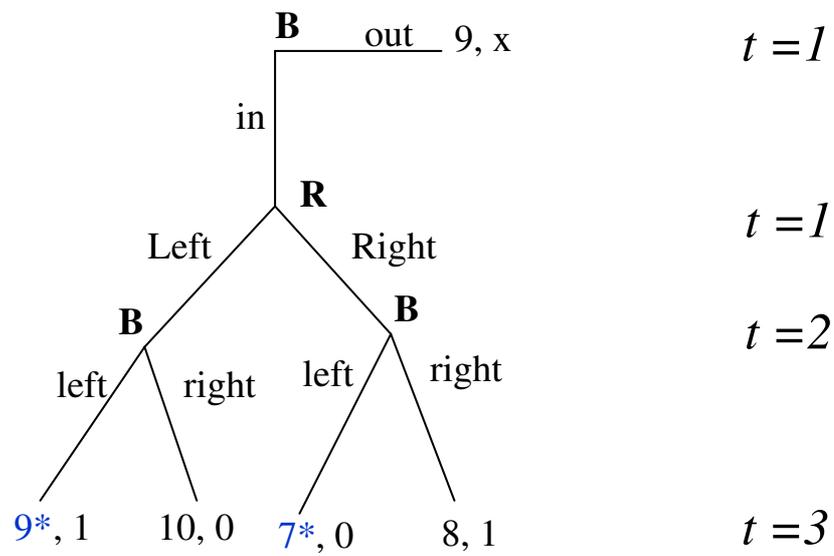
The rest of the paper is organized as follows. Section 2 presents the main ideas in the context of a stylized example. In section 3 we formalize the two solution concepts. In section 4 we apply the two solutions concepts on two textbook application and we discuss the pros and cons of each concept. Section 5 concludes.

2 A motivating example

We will motivate the main ideas in the context of a stylized example. Consider the following game, whose extensive form is depicted in figure 1. There are two players, referred to as player *B* (*B*oundedly rational and female) and player *R* (*R*ational and male). The game spans three periods. Player *B* moves first in period 1 and she can play *in* or *out*. If she plays *in*, then player *R* gets to move *Left* or *Right*, still in period 1. Player *B* observes the move of player *R* and moves again in period 2, *left* or *right*. Payoffs (even those that arise when player *B* plays *out*) are realized in period 3. This is unless player *B* moves *left* at either of her nodes, in which case she can enjoy her payoff (9 or 7 respectively) immediately, i.e. in period 2.

Furthermore, suppose that player *B* has time inconsistent preferences. In period 1 she prefers a payoff of 10 (resp. 8) in period 3 to a payoff of 9 (resp. 7) in period 2. Once period 2 rolls around, however, her preferences are reversed. These preference reversals could arise, for example, if player *B* used a β - δ discount function with $\beta = \frac{1}{2}$ and $\delta = 1$. Player *B* is also assumed to be *naive time inconsistent*: her period-1 self erroneously believes that her period-2 self will have the same preferences as her period-1 self. For example, in the context of the $\beta - \delta$ preferences, her period-1 self erroneously believes that her period-2 self will use the discount function $\{1, \delta\}$, whereas she will actually use $\{1, \beta\delta\}$. Player *B* is also naive in a different sense: she believes that player *R* shares her naive expectations about her future patience. Again in the context of $\beta - \delta$ preferences, this means that in period 1 player *B* thinks that player *R* thinks that her period-2 self uses the discount function $\{1, \delta\}$.

Player *R* knows all the information in the preceding paragraph. In a nutshell, player *R* knows everything the modeler does, whereas player *B* thinks that player *R* thinks about her beliefs and her preferences what she, herself, thinks. In this example we can also safely assume that player *R* does not discount the future at all.



All payoffs are realized in period $t = 3$, unless otherwise noted. An asterisk denotes a payoff realized in period 2.

Figure 1. A motivating example

How would this game be played out? A reasonable prediction is that in period 1 player B will reason by backwards induction. In period 1, when she carries out the backwards induction reasoning in her head, she thinks that in period 2 she will play *right* at both future nodes. Then, by putting herself in the shoes of player R , she deduces that player R will play *Right*. Thus, in period 1 she thinks that playing *in* will eventually give her a payoff of 8 in period 3. Playing *out*, instead, increases her payoff to 9 in period 3. Therefore, according to this prediction player B stays *out*.

We will refer to this solution as *naive backwards induction* (NBI) and formalize it shortly. We have added the qualifier “naive” in order to highlight the fact that player B ’s reasoning is naive in two different senses. First, in period 1 she naively thinks that her period-2 self will have the same preferences as she does. Second, in period 1 she erroneously believes that player R shares her naive expectations about her future patience.

There is, however, an alternative prediction where player B moves *in* instead. Suppose that player B expects player R to play *Left*. Player B does not care why player R will play this way and takes this as a given. It could be, for example, that before the game starts an outside observer announces that player R will play *Left*. Then, in period 1 player B reasons as follows: “If I play *in*, then he will play *Left*. I will then play *right* to get a payoff of 10 in period 3. This is more than the payoff of 9 in period 3 that I get by playing *out*.” Hence, if player B expects player R to play *Left*, then in period 1 she moves *in* (and plans to move *right* at both future nodes). And vice versa, if player B moves *in*, player R should play *Left*, just as player B expects him to play. In other words, player B ’s expectation that player R plays *Left* is self-confirming. To sum up, according to this prediction we have: a) player R plays *Left* and b) in period 1 player B *plans* to play (*in, right, right*), but *actually* plays (*in, left, left*).⁴ Given these strategies (and what players know at each of their information sets) neither player has a unilateral incentive to deviate. Moreover, each player’s strategy at each point in time is sequentially rational given this player’s current preferences.

We will refer to this prediction as *equilibrium* to emphasize its similarity to subgame perfect Nash equilibrium (SPNE). As in a Nash equilibrium, both players are endowed with beliefs about how the other plays and choose a best response to these beliefs. In equilibrium, each player’s beliefs are confirmed

⁴The convention we use here is that the second (resp. third) action in the triple refers to what player B plays on the information set on the left-hand (resp. right-hand) side.

by how the other player actually plays. Also, as in a SPNE, players' strategies are sequentially rational.

One can identify some key differences between our two solutions. Under NBI player B can rationalize her playing *out*. If we were to ask her in period 1 why she moves *out*, she can defend her choice with the following speech: "I play *out*, because I expect him to play *Right*, because he expects me to play *right* at both future nodes." Under equilibrium in period 1 player B can not rationalize why player B would play *Left*. However, she can defend her choice of playing *in* with a different speech: "I play *in*, because I expect her to play *Left*. In fact, she does plays *Left*, so my theory holds water." This argument is particularly attractive if we were to interpret equilibrium behavior as a *self-confirming* state of affairs or as the *steady-state* of social interaction as is often done in game theory textbooks.

Furthermore, notice that under equilibrium player B will be surprised by how her future selves play. For example, she will discover that in period 2 she plays *left* at both nodes, rather than *right* as she planned. This is in contrast to what happens under NBI where player B is surprised by how *both* she *and* player R actually play. For example, if she deviated and played *in* instead, she would be surprised to discover that player R plays *Left* and that she plays *left* at both future nodes.⁵

This last observation hides the key difference between our two solution concepts. Recall that in a single-agent decision problem a naive time inconsistent agent, by definition, cannot correctly predict her own future behavior. Then, when we embed such an agent into a game, we are confronted with two plausible possibilities for a solution concept. Under the first one (NBI), the inability of the naive agent to predict her own future behavior also leads her to incorrectly predict how other players will behave. Under the second one (equilibrium), the naive agent takes the behavior of others as given and simply chooses a best response. As we will argue, the appropriateness of using one solution concept over the other will ultimately depend on the economic application that one has in mind.

To summarize, in this section we discussed two sensible solutions in the context of our stylized example. Next, we formalize these two solution concepts.

⁵In our example such surprises can occur only off the NBI path. This is an artifact of this specific game. In general player B may be surprised even if she stays on the NBI path.

3 Formalizing the solution concepts

We cast the solution concepts in the context of finite extensive form games of perfect information. We will later extend to infinite horizon games (see alternating-offers bargaining game) and games of imperfect information (see durable goods monopoly).

Because it will be necessary to have access to some formal notation, we begin with a formal definition of a finite extensive form game. The following definition is adapted from Mas-Colell, Whinston and Green (1995).

A finite extensive form game can be described by the following ingredients:

1. A finite set of nodes \mathcal{X} , a finite set of actions \mathcal{A} , a finite set of players $\{1, 2, \dots, I\}$ and a finite set of periods $\{1, 2, \dots, T\}$.
2. A function $p : \mathcal{X} \rightarrow \{\mathcal{X} \cup \emptyset\}$, specifying the (unique) *immediate predecessor* of node x . The function $p(x)$ is non-empty for all nodes $x \in \mathcal{X}$, except for one, referred to as the *initial node* x_0 . A node $x' \in \mathcal{X}$, such that $p(x') = x$, will be referred to as an *immediate successor* of node x . Iteratively, we can obtain *all* successors and *all* predecessors of node x . We assume that the sets of all successors and all predecessors of a node x are disjoint, which implies that we have a *tree*. A node is referred to as a *terminal node*, if it has no successor nodes.
3. A function $\alpha : \mathcal{X} \setminus \{x_0\} \rightarrow \mathcal{A}$, specifying the action that one would need to take to go to a node x from its immediate predecessor $p(x)$. The function $\alpha(\cdot)$ satisfies the following property: Let x' and x'' be two distinct immediate successors of node x . Then, $\alpha(x') \neq \alpha(x'')$. The *set of choices available at a node x* is defined as $c(x) = \{a \in \mathcal{A}, \text{ such that } a = \alpha(x') \text{ for some immediate successor, } x', \text{ of node } x\}$.
4. A function $\iota : \mathcal{X} \rightarrow \{1, 2, \dots, I\}$ and a function $\tau : \mathcal{X} \rightarrow \{1, 2, \dots, T\}$, assigning each node to a player and a period respectively. The function $\tau(\cdot)$ satisfies the following property: If a node x' is a successor of node x , then $\tau(x') \geq \tau(x)$. We define $X_{t,t'}^i = \{x \in \mathcal{X}, \text{ such that } \iota(x) = i \text{ and } t \leq \tau(x) \leq t'\}$ as *the set of nodes that player i controls between periods t and t'* .
5. A collection of rational preference relations \succ_t^i , for each player and for each period, representing the preferences (over terminal nodes) of player i in period t .

Notice that the preceding definition differs from that of Mas-Colell, Whinston and Green (1995) in two respects. First, in the interest of cutting back on notation that is not essential for our purposes, we did not introduce information sets (recall that we are looking at games of perfect information) and we left no room for moves by *nature*. Second, we added a time dimension (periods $\{1, 2, \dots, T\}$) and we allowed players' preferences to depend on time (\succ_t^i).

All items 1 through 4 are common knowledge. Below we describe players' beliefs about their own and about other players' preferences.

3.1 Beliefs about preferences

We first describe players' beliefs about their own preferences.

3.1.1 Beliefs about own preferences

Each player $i \in I$ belongs to one, and only one, of two (disjoint) sets, \mathcal{R} and \mathcal{B} . Players who are members of set \mathcal{R} (\mathcal{R} for *rational*) are referred to as *sophisticated time inconsistent* (male players), whereas those who are members of set \mathcal{B} (\mathcal{B} for *boundedly rational*) are referred to as *naive time inconsistent* (female players). Below, we define the precise meaning of these terms. (Note the difference in the subscript in $\succ_t^{i'}$ and $\succ_{t'}^{i'}$, in the last line of each definition.)

Definition (sophisticated time inconsistent). *In every period t' , a sophisticated time inconsistent player $i' \in \mathcal{R}$ believes (correctly) that his preferences at any future period $t \geq t'$ will be given by $\succ_t^{i'}$.*

Definition (naive time inconsistent). *In every period t' , a naive time inconsistent player $i' \in \mathcal{B}$ believes (erroneously) that her preferences at any future period $t \geq t'$ will be given by $\succ_{t'}^{i'}$.*

In other words, a sophisticated player knows the preferences of all his future selves. On the contrary, a naive player (erroneously) believes that her future selves have the same preferences as her current self.⁶ These definitions are standard in the literature.

⁶A time consistent player can be thought of as a special case of either a sophisticated or a naive player.

3.1.2 Beliefs about other players' preferences

The following statements describe players' beliefs about other players' preferences.

- In each period t' it becomes common knowledge that the preferences of player i , for all $i \in I$, are given by $\succ_{t'}^i$.
- In each period t' all sophisticated time inconsistent players believe (correctly) that the preferences of player i , for all $i \in I$, in all future periods $t \geq t'$, will be given by \succ_t^i .
- In each period t' all naive time inconsistent players believe (erroneously) that the preferences of player i , for all $i \in I$, in all future periods $t \geq t'$, will be given by $\succ_{t'}^i$.

The first statement simply says that in every period the preferences of all *current* selves of all players are revealed. The second statement says that sophisticated time inconsistent players are aware that the future selves of other players may have different preferences from their current selves and they know what these preferences are. In other words, sophisticated time inconsistent players are sophisticated, not only about their own preferences, but about the preferences of others as well. On the contrary, according to the third statement, naive time inconsistent players are naive about their own preferences *and* about the preferences of others.⁷

3.1.3 Higher-order beliefs

The following statements describe players' higher-order beliefs.

- Each naive time inconsistent player, $i \in \mathcal{B}$, believes that all players are members of \mathcal{B} , i.e. she believes that $\mathcal{R} = \emptyset$.
- The previous statement is common knowledge in the game.

⁷To appease any likely objections, let us come clean and say upfront that there are many other beliefs (or combination of beliefs) that one could alternatively impose. For example, we could also have players who are sophisticated about their own preferences but naive about other players' preferences, or players who are naive about their own preferences but sophisticated about others' preferences. We believe, however, that the beliefs we specify are very natural for the applications of the following section.

- Each sophisticated time inconsistent player knows the identities of players that belong to the sets \mathcal{R} and \mathcal{B} respectively.
- All of the above is common knowledge among the sophisticated time inconsistent players, $\in \mathcal{R}$.

In other words, naive players (erroneously) believe that all other players share their (naive) expectations about all players' preferences. Sophisticated time inconsistent players, however, are sophisticated about the naivete or sophistication of other players. In a nutshell, each naive time inconsistent player (erroneously) believes that her beliefs are shared by everyone else in the game. Each sophisticated time inconsistent player knows everything the modeler does.⁸

3.2 Definitions of strategies

We define below a (pure) strategy for player i .

Definition (strategy). *A strategy for player i , s_i , is a family of functions, one for each period t , $s_{i,t} : X_{t,T}^i \rightarrow \mathcal{A}$, such that $s_{i,t}(x) \in c(x)$. Hence, we will write $s_i = (s_{i,1}, s_{i,2}, \dots, s_{i,T})$.*

In other words, the strategy $s_{i,t}$ is the contingency plan of the period- t self of player i . The following analogy may be useful. Think of a strategy for player i as a *notebook* with T pages. Each page has T lines. The period-1 self of player i starts on line 1 of page 1 and on each line $t \geq 1$ she writes down what actions she plans to execute in all of her period- t nodes. The period-2 self of player i starts on line 2 of page 2 and on each line $t \geq 2$ she writes what actions she plans to execute in her period- t nodes. The period-3 self of player i starts on line 3 of page 3 and so on. Keeping this analogy in mind, we can read the notation $s_{i,t}$ as “page t of player i 's notebook”.

⁸The following metaphor may help clarify the type of environment we have in mind. Suppose that in each period t all players appear in front of Zeus, who announces publicly what the preferences of all *current* selves are. Moreover, Zeus assures everyone that these are the preferences that all players will also have in the future. Subsequently, Zeus gathers in a secret meeting all sophisticated players ($i \in \mathcal{R}$). There, he informs them that he lied and that the preferences of some (or all) players may change over time. He also tells them what these future preferences will be. Finally, he assures them that this is a secret meeting and players not present in the meeting ($i \in \mathcal{B}$) have no reason to believe or suspect that such a meeting is taking place.

Also, the notation s will refer to the strategy profile of all players. That is, $s = (s_1, s_2, \dots)$

Given a strategy $s_i = (s_{i,1}, s_{i,2}, \dots, s_{i,T})$ of player i , it will also be convenient to define \tilde{s}_i as the *actual strategy* of player i as follows:

Definition (actual strategy). Let $s_i = (s_{i,1}, s_{i,2}, \dots, s_{i,T})$ be a strategy for player i . Then, the actual strategy \tilde{s}_i is the function $\tilde{s}_i : X_{1,T}^i \rightarrow \mathcal{A}$, such that $\tilde{s}_i(x) = s_{i,\tau(x)}(x)$.

In other words, the actual strategy \tilde{s}_i is the strategy that one would execute if in each period t she played according to the contingency plan of the period- t self of player i . Using our earlier notebook analogy, \tilde{s}_i is the *diagonal* of the notebook, i.e. line 1 of page 1, line 2 of page 2, line 3 of page 3 and so on. Also, the notation \tilde{s}_{-i} will have its usual meaning: the profile of actual strategies of all players, except the one that corresponds to player i . Thus, we write $\tilde{s}_{-i} = (\tilde{s}_1, \dots, \tilde{s}_{i-1}, \tilde{s}_{i+1}, \dots)$.

For future reference, it will also be convenient to define $s_{i,t,t'}$, for $t' \geq t$. This is line t' , page t , of player i 's notebook. Formally, we have:

Definition (line t , page t). Let $t' \geq t$. Then, $s_{i,t,t'} : X_{t',t'}^i \rightarrow \mathcal{A}$, such that $s_{i,t,t'}(x) = s_{i,t}(x)$, for all $x \in X_{t',t'}^i$.

Example 1 Return to our motivating game from figure 1. (Here, $i \in \{B, R\}$). A strategy for player B consists of a pair $(s_{B,1}, s_{B,2})$, such that:

$$\begin{aligned} s_{B,1} &\in \{\{in, out\} \times \{left, right\} \times \{left, right\}\} \\ s_{B,2} &\in \{\{left, right\} \times \{left, right\}\} \end{aligned}$$

An example of a strategy for player B is:

$$\begin{aligned} s_{B,1} &= (in, right, right) \\ s_{B,2} &= (left, left) \end{aligned}$$

Using our notebook analogy, this strategy corresponds to the following “notebook”:

	page 1	page 2
line 1	in	
line 2	right, right	left, left

This strategy specifies that the period-1 self of player B plans to play $(in, right, right)$, whereas the period-2 self of player B plans to play $(left, left)$.

Given this strategy, the actual strategy for player B , \tilde{s}_B , is:

$$\tilde{s}_B = (in, left, left)$$

This is the strategy that a delegate of player B would execute if in each period t she played according to the contingency plan of the period- t self of player B .

Finally, given this strategy (or notebook), we can write:

$$\begin{aligned} s_{B,1,1} &= in \\ s_{B,1,2} &= (right, right) \\ s_{B,2,2} &= (left, left) \end{aligned}$$

We will also utilize two more pieces of notation, $\tilde{s}_{i,-t}$ and $\tilde{s}_{i,\hat{-}t}$. The former is the actual strategy of player i , excluding the t^{th} element. In other words, it is the diagonal of player i 's notebook, excluding the entry that corresponds to page t . The latter is defined only as long as $t > 1$ and is the actual strategy of player i , up to, but not including, the t^{th} element. In other words, it is the diagonal of player i 's notebook, excluding the entries that correspond to pages $t, t + 1, \dots, T$. Returning to our earlier example, we can write: $\tilde{s}_{B,-1} = (left, left)$ and $\tilde{s}_{B,\hat{-}2} = in$.

For the reader's convenience, we summarize all notation pertaining to players' strategies in Table 1, using our "notebook" analogy.

3.3 Solution concepts

We now formalize the solution concepts of *equilibrium* and *naive backwards induction*.

Definition (equilibrium). A strategy profile $s = (s_1, s_2, \dots)$ constitutes an equilibrium if the following statements hold:

1. For all $i \in \mathcal{R}$ and all periods t , $s_{i,t,t}$ is a best response to $(\tilde{s}_{i,-t}, \tilde{s}_{-i})$ under \succ_t^i .
2. For all $i \in \mathcal{B}$ and all periods t , $s_{i,t}$ is a best response to $(\tilde{s}_{i,\hat{-}t}, \tilde{s}_{-i})$ under \succ_t^i .
3. The strategy profile $s = (s_1, s_2, \dots)$ induces behavior that is consistent with the previous statements in all subgames.

Table 1: Summary of notation pertaining to strategies

Notation	Meaning
s_i	the notebook of player i
$s_{i,t}$	page t of player i 's notebook
$s_{i,t,t'}$	line t , page t' of player i 's notebook
\tilde{s}_i	the diagonal of player i 's notebook
$\tilde{s}_{i,-t}$	the diagonal, excluding the element corresponding to page t
$\tilde{s}_{i,\widehat{-t}}$	the diagonal, excluding the elements that correspond to pages $t, t + 1, \dots, T$ (defined as long as $t > 1$)

The first statement says that each sophisticated player plays a best response to others' *actual* strategies. Moreover, the term "others" in the previous sentence refers to both "other actual players" and all "other selves" of the player in question. The second statement says that naive players also play a best response to others' actual strategies. In this case, however, the term "others" refers only to "other actual players". This is because naive players fail to realize that their own future selves will feel about the game differently. The final statement guarantees that players' actions are sequentially rational.

Example 2 *Return again to our motivating example. The unique equilibrium is the strategy profile:*

$$(s_B, s_R) = ((in, right, right), (left, left); (Left))$$

Given player B's actual strategy, $\tilde{s}_B = (in, left, left)$, the strategy $s_{R,1,1} = Left$ is a best response for player R. And vice versa, given player R's actual strategy, $\tilde{s}_R = Left$, the strategy $(in, right, right)$ is a best response for the period-1 self of player B. So is the strategy $(left, left)$ for the period-2 self of player B.

The strategy profile $s = ((out, right, right), (left, left); Left)$ is an example of a strategy profile that satisfies requirements (1) and (2), but fails requirement (3). Given player B's actual strategy, $\tilde{s}_B = (out, right, right)$, the strategy $\tilde{s}_2 = Right$ is a best response for player R. And vice versa, given

player R 's actual strategy, $\tilde{s}_R = \text{Right}$, the strategy of the period-1 self of player B , $(\text{out}, \text{right}, \text{right})$, is a best response. So is the strategy of period-2 self of player B , $(\text{left}, \text{left})$. However, in the sub-game that starts at the information set of player R , the strategy Right is not sequentially rational for player R .

Definition (Naive Backwards Induction). A strategy profile $s = (s_1, s_2, \dots)$ survives naive backwards induction if the following statements hold:

1. For all $i \in \mathcal{R}$ and all periods t , $s_{i,t,t}$ is a best response to $(\tilde{s}_{i,-t}, \tilde{s}_{-i})$ under \succ_t^i .
2. For all $i \in \mathcal{B}$ and all periods t , the strategy $s_{i,t}$ survives backwards induction in the subgame played by the period- t selves of all players.
3. The strategy profile s induces behavior that is consistent with the previous statements in all subgames.

In other words, under NBI sophisticated players play a best response to how others (including their own future selves) will actually play (statement 1, same as for equilibrium). Naive players play a best response to how they *deduce* other players will play (statement 2). Statement 3 guarantees that strategies are also sequentially rational.

Example 3 In our motivating example the strategy profile that survives NBI is:

$$(s_B, s_R) = ((\text{out}, \text{right}, \text{right}), (\text{left}, \text{left}); (\text{Left}))$$

Given player B 's actual strategy, $\tilde{s}_B = (\text{out}, \text{left}, \text{left})$, the strategy Left is a best response for player R (requirement 1). Also, in the game between the period-1 selves of both players, the strategy $(\text{out}, \text{right}, \text{right})$ does survive backwards induction (requirement 2). The strategy profile $(s_B, s_R) = ((\text{out}, \text{right}, \text{right}), (\text{left}, \text{left}); (\text{Right}))$ satisfies statements 1 and 2, but fails the requirement of sequential rationality imposed by statement 3.

Notice, and this follows directly from the preceding definitions, that when all players are sophisticated, the two solution concepts are equivalent.

4 Two economic applications

In this section we apply the solution concepts of equilibrium and NBI on two economic applications, namely Rubinstein's alternating-offers bargaining game and the durable goods monopoly. Our objective is twofold. First, we would like to evaluate how our solution concepts perform in practice. Second, we would like to discuss the relative pros and cons of each solution concept. We believe that in order to address both of these issues, it is important to move away from stylized examples and cast our solution concepts in the context of well-understood games and full-blown economic applications.

We should also add that embedding time inconsistent players in both of these applications raises various interesting economic issues which go beyond the problem of the appropriate solution concept. Here we simply solve these games under both solution concepts in a homework-exercise manner. For a richer discussion of the alternating-offers bargaining game with time inconsistent players see Akin (2005). The durable goods monopoly is analyzed in greater detail in Sarafidis (2004).

4.1 Rubinstein's alternating-offers bargaining game

Two players, A and B , bargain on how to split a divisible prize, say one dollar. The rules of the game are as follows. First, player A makes a proposal on how to split the dollar. If player B accepts, then the proposal is implemented and the game ends. If not, then player B gets to make a proposal in the subsequent period, that player A can accept or reject and so on.⁹ We will analyze the infinite horizon version of the game, which stipulates that the game goes on until players reach an agreement.

We assume that both players, A and B , are naive time inconsistent with the same $\beta - \delta$ preferences. That is, in each period they use the discount function $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \dots\}$ and they (naively) think that from the next period onwards they will use the discount function $\{1, \delta, \delta^2, \dots\}$. Moreover, in each period it is a common belief that players share (and will always share) the same preferences. This implies that each player is also naive about the other player's preferences.

A useful benchmark for our subsequent analysis will be the SPNE of the game when players have standard time consistent preferences with discount

⁹We assume that if a player is indifferent between accepting or rejecting, she will accept.

factor ρ . In each period t the SPNE of the time consistent game involves the following strategies for both players:¹⁰

$$\text{Propose } \left(\frac{1}{1+\rho}, \frac{\rho}{1+\rho}\right); \text{ accept } \frac{\rho}{1+\rho} \text{ or more.}$$

Notice that in the SPNE a player offers her opponent just enough to make her indifferent between accepting or rejecting. For example, if player B accepts the offer of player A in period 1, she gets $\frac{\rho}{1+\rho}$. If instead she rejects, she gets to make a proposal in period 2. The most she can get in period 2 is $\frac{1}{1+\rho}$, which after one period of discounting is worth $\frac{\rho}{1+\rho}$. Hence, player B will accept the proposal of player A in period 1 and the game will end immediately.

Equilibrium. Consider the following candidate equilibrium:

$$\text{In the current period } t: \text{ propose } \left(\frac{1}{1+\beta\delta}, \frac{\beta\delta}{1+\beta\delta}\right); \text{ accept } \frac{\beta\delta}{1+\beta\delta} \text{ or more.}$$

$$\text{From period } t + 1 \text{ onwards: propose } \left(\frac{1}{1+\beta\delta}, \frac{\beta\delta}{1+\beta\delta}\right); \text{ accept } \frac{\delta}{1+\beta\delta} \text{ or more.}$$

The first line refers to the strategies that players *actually* play in period t . The second line refers to the strategies that (in period t) players *plan* to play in the future. According to this candidate equilibrium, in each period t players behave as if they were playing the time consistent game with discount factor $\rho = \beta\delta$. One can verify that no player has a unilateral incentive to deviate from this strategy profile. Since each subgame looks like the game itself, the subgame perfection requirement is also satisfied and, hence, our candidate equilibrium is indeed an equilibrium.

We should note that despite the fact that these strategies constitute an equilibrium, neither player can rationalize why the other uses this strategy. For example, in period 1 player A cannot understand why player B would offer her only $\frac{\beta\delta}{1+\beta\delta}$ in period 2. In period 1 Player A may think as follows:

“If I accept the period-2 proposal, I will get $\frac{\beta\delta}{1+\beta\delta}$, which is worth today $\frac{(\beta\delta)^2}{1+\beta\delta}$. If instead I reject, I can get $\frac{1}{1+\beta\delta}$ in period 3, which is worth today $\frac{\beta\delta^2}{1+\beta\delta}$. Hence, in period 2 I will reject. Player B will end up with $\frac{\beta\delta}{1+\beta\delta}$ in period 3,

¹⁰A proposal of (x, y) means that the proposer concedes y to the other player and keeps x for herself.

which is worth to her today $\beta\delta^2 \frac{\beta\delta}{1+\beta\delta}$. Why then doesn't she offer me $\frac{\delta}{1+\beta\delta}$ in period 2, which I will accept and this will make her better off?¹¹ ”

What player A is missing, is that in period 2 she will not be as patient as she now thinks she will be. As a result, contrary to her current expectations, once period 2 rolls around she will be willing to accept $\frac{\beta\delta}{1+\beta\delta}$.

Naive Backwards Induction. The algorithm of backwards induction is applicable only in finite games of perfect information. Nevertheless, the basic logic behind NBI can be extended to accommodate the infinite horizon alternating-offers bargaining game (as well as the durable good monopoly that we consider next). To do this, we need to modify the second statement of the NBI definition as follows:

- For all $i \in \mathcal{B}$ and all periods t , the strategy $s_{i,t}$ is an *equilibrium* strategy in the subgame played by the period- t selves of all players.

As long as all the subgames played by the period- t selves of all players have a unique equilibrium, the definition does not leave any ambiguities. This is indeed the case in our bargaining game.

Suppose we are in period t and, without loss of generality, let it be player A 's turn to make a proposal. Both players agree that from period $t + 1$ onwards they will both be using the discount function $\{1, \delta, \delta^2, \dots\}$. Hence, they both agree that if they do not reach an agreement in the current period t , then in the subgame that begins in period $t + 1$ they will play the SPNE of the time consistent game and discount factor $\rho = \delta$. Therefore, if player B rejects the period- t proposal, she thinks that she can get $\frac{1}{1+\delta}$ in period $t + 1$. As a result, in period t she is willing to accept $\frac{\beta\delta}{1+\delta}$ or more. This leaves $1 - \frac{\beta\delta}{1+\delta}$ for player A . This is summarized by the following strategy profile:

In the current period t : propose $(1 - \frac{\beta\delta}{1+\delta}, \frac{\beta\delta}{1+\delta})$; accept $\frac{\beta\delta}{1+\delta}$ or more.

From period $t + 1$ onwards: propose $(\frac{1}{1+\delta}, \frac{\delta}{1+\delta})$; accept $\frac{\delta}{1+\delta}$ or more.

To confirm that these strategies survive NBI, we also need to note that:

$$1 - \frac{\beta\delta}{1+\delta} \geq \beta\delta \frac{\delta}{1+\delta} \tag{1}$$

¹¹To verify that player B will be better off, note that $\beta\delta(1 - \frac{\delta}{1+\beta\delta}) > \beta\delta^2 \frac{\beta\delta}{1+\beta\delta}$.

The LHS is the most that player A can get in period t , if she proposes a deal that player B will accept. The RHS is what player A can get in period $t + 1$, given the proposal that she naively expects player B will make. The direction of the inequality says that player A prefers to reach an agreement in period t , rather than stall the bargaining and accept the proposal of player B in period $t + 1$.

The possibility of delays in bargaining. As we pointed out earlier, under both solution concepts a naive player may not correctly predict how subsequent subgames will be played out. This raises the possibility that as a game progresses and naive players keep on being surprised (either by how they or others play), they may start questioning what they take for granted and may eventually learn that their future selves do not have the same preferences as their current selves. Notice, however, that in the context of our bargaining game both solution concepts are robust to this possibility. This is because both solution concepts predict that players will reach an agreement immediately. As a result, as long as players stay on the equilibrium/NBI path, they will never be surprised and, hence, they will not have a reason to question their preferences.

In what follows, we briefly discuss a version of the bargaining game with one sophisticated player, where players will fail to reach an agreement immediately. Consider now the case where the second player is sophisticated time inconsistent. To avoid confusion, let's call players A and R (rather than B). That is, in period t player R is aware that from period $t + 1$ onwards the future selves of both players will use the discount function $\{1, \beta\delta, \beta\delta^2, \dots\}$. The same argument we used above shows that, under NBI, in period 1 the naive player A will propose $(1 - \frac{\beta\delta}{1+\delta}, \frac{\beta\delta}{1+\delta})$. Now, player R knows that in period 2 player A will accept $\frac{\beta\delta}{1+\delta}$ or more, despite the fact that player A plans to accept $\frac{\delta}{1+\delta}$ or more. Hence, player R compares:

$$\frac{\beta\delta}{1+\delta} \leq \beta\delta(1 - \frac{\beta\delta}{1+\delta}) \quad (2)$$

The LHS is what player R will get in period 1 if he accepts. The RHS is what player R will get if he rejects and gets a chance to propose in period 2. The direction of the inequality means that player R will prefer to stall the bargaining. He rejects the proposal of player A in period 1 and then strikes a better deal in period 2.

4.2 Inter-temporal price discrimination in a durable good monopoly

A durable good monopolist faces a population of consumers. Each consumer is indexed by her willingness-to-pay (valuation), v , and can buy at most one unit of the good. The monopolist cannot observe the valuation v of any individual consumer, but he knows that consumer valuations are uniformly distributed in the interval $[0, 1]$. In each period t the monopolist sets a price p_t . Consumers decide whether to buy at the current-period price, p_t , or to postpone their decision. The monopolist's objective is to maximize his discounted stream of profits. Each consumer maximizes her discounted gains from trade, $v - p_t$. For simplicity, we will assume that marginal cost is zero and we will analyze the three period model.

Contrary to the standard model, our consumers are naive time inconsistent with $\beta - \delta$ preferences. That is, in period 1 consumers use the discount function $\{1, \beta\delta, \beta\delta^2\}$, whereas in period 2 they use the discount function $\{1, \beta\delta\}$. Also, in period 1 they erroneously believe that their period-2 selves will use the discount function $\{1, \delta\}$, i.e. they are naive time inconsistent. Moreover, in period 1 consumers erroneously believe that the monopolist shares their naive expectations about their future patience. That is, in period 1 consumers think that the monopolist thinks that period-2 consumers use the discount function $\{1, \delta\}$.

The monopolist knows all this. Also, he has standard geometric preferences with discount factor γ and this fact is common knowledge. Next, we will analyze how the game will be played out under both solution concepts.

The continuation game that starts in period 3. Suppose that in period 3 the monopolist faces these consumers whose valuations are uniformly distributed in the interval $[0, x]$, with $x < 1$. Then, the monopolist solves:

$$\max_{p_3} (x - p_3)p_3 \quad (3)$$

The solution is given by:

$$p_3(x) = \frac{x}{2} = v_3(x) \quad (4)$$

$$\pi_3(x) = \frac{x^2}{4} \quad (5)$$

where π_3 denotes the value of the objective function at the maximum and v_3 denotes the valuation of the *marginal* period-3 consumer. Equation (4) is the

monopolist's strategy as a function of the state x . The strategy of consumer v is to buy if and only if $v \geq v_3$.

The continuation game that starts in period 2. Suppose again that in period 2 the monopolist faces these consumers whose valuations are uniformly distributed in the interval $[0, x]$. Then, the monopolist solves:

$$\begin{aligned} & \max_{v_2, p_2} (x - v_2)p_2 + \gamma\pi_3(v_2) & (6) \\ \text{s.t.} \quad & v_2 - p_2 = \beta\delta(v_2 - p_3(v_2)) \end{aligned}$$

The objective function is the monopolist's profit from period 2 onwards. The constraint says that in period 2 there is a marginal consumer with valuation, v_2 , who is indifferent between buying in period 2 or in period 3. (Recall that once period 2 rolls around the consumer discount factor between periods 2 and 3 is $\beta\delta$ and the monopolist knows this.) Substituting expressions (4) and (5) into the monopolist's maximization problem and solving the constrained optimization yields:

$$p_2(x) = \frac{x}{2} \frac{(2 - \beta\delta)^2}{4 - 2\beta\delta - \gamma} \quad (7)$$

$$v_2(x) = x \frac{(2 - \beta\delta)}{4 - 2\beta\delta - \gamma} \quad (8)$$

$$\pi_2(x) = \frac{x^2}{4} \frac{(2 - \beta\delta)^2}{4 - 2\beta\delta - \gamma} \quad (9)$$

where π_2 denotes the value of the period-2 objective function at the optimum. Equation (7) is the monopolist's strategy as a function of the state x . A consumer's strategy is to buy if and only if her valuation is greater than that of the marginal consumer, v_2 . For future reference, we should note that expressions (7)-(9) are also functions of the product $\beta\delta$, the consumer discount factor between periods 2 and 3. We will, therefore, write $p_2(x, \beta\delta)$, $\pi_2(x, \beta\delta)$ and $v_2(x, \beta\delta)$.

So far we have not done anything different from what we would have done for the standard game with time consistent consumers. This is because with only two remaining periods time inconsistency does not matter. Hence, players' strategies are the same under both solution concepts and they coincide with the SPNE of the time consistent game when the consumer discount

factor equals the product $\beta\delta$.¹² As we shall now see, this will not be the case in period 1, where equilibrium and NBI lead to different predictions.

Equilibrium. In period 1 the monopolist faces the full population of consumers whose valuations are uniformly distributed in the interval $[0, 1]$. Under equilibrium the monopolist solves:

$$\begin{aligned} & \max_{v_1, p_1} (1 - v_1)p_1 + \gamma\pi_2(v_1, \beta\delta) & (10) \\ \text{s.t.} \quad & (v_1 - p_1) \geq \beta\delta(v_1 - p_2(v_1, \beta\delta)) \\ & (v_1 - p_1) \geq \beta\delta^2(v_1 - p_3(v_2(v_1, \beta\delta))) \end{aligned}$$

The objective function is the monopolist's profit from period 1 onwards. The two constraints say that in period 1 the marginal consumer, v_1 , (weakly) prefers to buy in period 1 than to buy in periods 2 or 3, *at the correct period-2 or period-3 prices which she takes for granted*. It can be shown that depending on the value of the parameters β , δ and γ either constraint could be the binding one.¹³ This is in contrast to what happens in the time consistent game where the period-1 marginal consumer is always determined by an indifference condition between periods 1 and 2 (i.e. the first constraint is always the one that binds). Furthermore, notice that when the first constraint happens to be the binding one, the equilibrium coincides with the SPNE of the standard time consistent game, where consumers use the discount function $\{1, \beta\delta, (\beta\delta)^2\}$.

To sum up, in equilibrium the game will be played out as follows. The solution to (10) determines the first period price and marginal consumer, p_1 and v_1 . We then enter period 2 with consumers whose valuations are uniformly distributed in the interval $[0, v_1]$. From period 2 onwards equilibrium strategies are given by $p_2(v_1, \beta\delta)$, $v_2(v_1, \beta\delta)$ and $p_3(v_2) = v_3(v_2)$ from equations (7), (8) and (4) respectively.

Naive Backward Induction. Under NBI consumers decide whether to buy in period 1, or to postpone their purchase, based on the period-1 price,

¹²Formally speaking, the solution to the game is not a *subgame* perfect Nash equilibrium. This is a game of imperfect information and the continuation game that starts in period 2 is not a proper subgame. It would be more precise to refer to the equilibrium as a strategy profile that also induces equilibrium behavior in all *continuation games* (as opposed to all *subgames*). Gul, Sonnenschein and Wilson (1986), as well as Besanko and Winston (1990), refer to the solution as a *SPNE* and we follow them.

¹³Sarafidis (2004) provides conditions that determine the structure of the equilibrium (i.e. which of the two constraints is binding) and provides economic intuition.

p_1 , and the prices that they (naively) predict for periods 2 and 3. Consumers may be naive about their period-2 preferences, but they are sophisticated game theorists. As such, by backwards induction in period 1 they realize that the period-2 price will be given by $p_2(\cdot, \cdot)$ from equation (7). However, because of their naivete they will evaluate this function at the wrong point.

To see this, recall that in period 1 consumers (naively) expect that in period 2 they use the discount function $\{1, \delta\}$. They also think that the monopolist shares their naive expectation. Therefore, in period 1 they believe that the period-2 price will be given by $p_2(v_1, \delta) = \frac{v_1}{2} \frac{(2-\delta)^2}{4-2\delta-\gamma}$ instead of the correct one $p(v_1, \beta\delta)$ from equation (7). Similarly, consumers think that the period-3 price will be given by $p_3(v_2(v_1, \delta))$ instead of the correct one $p_3(v_2(v_1, \beta\delta))$.

The monopolist knows all this. As a result, under NBI in period 1 he solves:

$$\begin{aligned} & \max_{v_1, p_1} (1 - v_1)p_1 + \gamma\pi_2(v_1, \beta\delta) & (11) \\ \text{s.t.} \quad & (v_1 - p_1) \geq \beta\delta(v_1 - p_2(v_1, \delta)) \\ & (v_1 - p_1) \geq \beta\delta^2(v_1 - p_3(v_2(v_1, \delta))) \end{aligned}$$

The objective function (the monopolist's profit from period 1 onwards) is, of course, the same as under equilibrium. The constraints say that the marginal consumer (weakly) prefers to buy in period 1 at the price p_1 than to buy in period 2 (first constraint) or in period 3 (second constraint) *at the prices that she naively predicts*.

The following argument shows that the first constraint will always be the binding one. In period 1 consumers think that the period-2 price, $p_2(v_1, \delta)$, will be set so as to make the consumer with valuation $v_2(v_1, \delta)$ indifferent between buying in periods 2 or 3. Moreover, in period 1 consumers expect the period-3 price to be $p_3(v_2(v_1, \delta))$ and they think that the discount factor between periods 2 and 3 is δ . Hence, we can write:

$$v_2 - p_2(v_1, \delta) = \delta(v_2 - p_3(v_2(v_1, \delta))) \quad (12)$$

Since $v_1 > v_2$, we can replace v_2 with v_1 in equation (12) and change the equality to a strict inequality (greater than, $>$). Finally, multiply both sides by $\beta\delta$ to get the desired result. (This is essentially the same argument as the one for the standard game with time consistent consumers.)

Hence, under NBI the game will develop as follows. The solution to the maximization problem in (11) determines the period-1 price, p_1 , and

marginal consumer, v_1 . (One can safely drop the second constraint from the maximization problem, because the first constraint will always be the one that binds.) From period 2 onwards, prices and marginal consumers are given by p_2 , v_2 and $p_3(v_2) = v_3(v_2)$ from equations (7), (8) and (4). As we pointed out earlier, once period 2 rolls around the consumer discount factor equals the product $\beta\delta$ and the expressions for p_2 , v_2 and $p_3(v_2)$ are evaluated at the point $(v_1, \beta\delta)$.

4.3 Equilibrium and naive backwards induction revisited

Having seen how our two solution concepts perform in the context of some well-known games, we now discuss their relative strengths and weaknesses.

As we pointed out earlier, under NBI players can *rationalize* their choices, in the sense of Bernheim (1984) and Pearce (1984), with a chain of justifications of the form: “I play x , because she plays y , because I play z ...” While this is an attractive feature of NBI in the context of our two-player bargaining game, it is a problematic feature in the context of the durable good monopoly. Do consumers really forecast future prices by putting themselves in the shoes of the monopolist (realizing that the monopolist will put himself in their shoes and so on)?

Maybe, a more realistic story is that consumers observe how the monopolist (or other monopolists in similar markets) have priced in the past and they make their purchasing decisions on the assumption that they will also have access to the same price path. This story is internally consistent as long as (given the consumers’ decisions) the monopolist finds it indeed optimal to follow the price path that consumers expect him to follow. This is precisely the story implied by equilibrium.

This equilibrium story, however, is problematic when applied to the bargaining game. In a market setting consumers have access to information pertaining to how the monopolist (or other monopolists in similar markets) have priced in the path. Consumers can then use this information to form beliefs about what prices they will face in the future. This type of information, i.e. how other players have played in the past, is harder to come by in the context of a bargaining game. Often the final outcomes of bargaining games are simply not revealed to outsiders. Even when the final outcome is revealed, it is hard, if not impossible, to know the exact path that led to this

outcome (i.e. what offers were made that were rejected and at which rounds). Moreover, if past bargainers really played according to the predictions of the theory, then they all reached agreements immediately. Since, there is no history off the equilibrium path, it is even harder to explain how a bargainer can form beliefs about how her opponent will behave in the future.

To conclude, we believe that it is pointless to debate about which of the two is a “better” solution concept in general. The economic and social context is crucial and this debate can be meaningful only within the context of a particular application.

5 Conclusion

This paper introduced agents (players) with time inconsistent preferences in non-cooperative games. When all players are sophisticated time inconsistent there is no need to tamper with the standard equilibrium concepts. However, when some agents are naive time inconsistent, it is not evident what the appropriate solution concept is. We proposed and formalized two solution concepts, which we termed *equilibrium* and *naive backwards induction*. The former is in the spirit of subgame perfection, whereas the latter relies on the algorithm of backwards induction. Unlike in standard games with time consistent preferences, these concepts (may, and most likely will) lead to different predictions.

We applied these solution concepts on two well-known economic applications, Rubinstein’s alternating-offers bargaining game and the durable good monopoly, and we discussed their pros and cons in the context of these applications. We argued that the social and economic context is important and each solution concept may be more or less attractive depending on the application one has in mind.

In closing, we would like to sketch two directions for future research. First, in the course of a dynamic game naive time inconsistent players will discover that their future selves violate plans that their previous selves have chosen. As a result, naive players may eventually become aware of their naivete (i.e. become sophisticated time inconsistent) and our solution concepts should be generalized to allow for this type of learning. Second, when a player is uncertain about whether her opponents are naive or sophisticated time inconsistent a game of incomplete information will arise. This set up may give rise to interesting signalling issues and reputational concerns.

References

- [1] Akin, Z. (2005). “Time Inconsistency and Learning in Bargaining Games”, *working paper*.
- [2] Bernheim, B. D. (1984). “Rationalizable Strategic Behavior”, *Econometrica*, 52, 1007-1028.
- [3] Besanko, D. and W. L. Winston. (1990). “Optimal Price Skimming by a Monopolist Facing Rational Consumers”, *Management Science*, 36 (5), 555-567.
- [4] Carrillo, J. and T. Mariotti. (2000). “Strategic Ignorance as a Self-Discipline Device”, *Review of Economic Studies*, 67, 529-544.
- [5] Chade, H., Prokopovych, P. and L. Smith. (2005). “Repeated Games with Present Biased Preferences”, *working paper*.
- [6] Chung, K. S. and J. Ely. (2005). “Foundations of Dominant Strategy Mechanisms”, *working paper*
- [7] Coase, R. H. (1972). “Durability and Monopoly”, *Journal of Law and Economics*, 15, 143-149.
- [8] Della Vigna, S. and U. Malmendier. (2004). “Contract Design and Self-Control: Theory and Evidence”, *Quarterly Journal of Economics*, 119 (2), pp. 353-402.
- [9] Gul, F., H. Sonnenschein and R. Wilson. (1986). “Foundations of Dynamic Monopoly and the Coase Conjecture”, *Journal of Economic Theory*, 39, 155-190.
- [10] Laibson, D. (1997). “Golden Eggs and Hyperbolic Discounting”, *Quarterly Journal of Economics*, 62 (2), 443-477.
- [11] Mas-Colell, A., M. Whinston and J. Green. (1995). *Microeconomic Theory*, New York, Oxford University Press.
- [12] O’Donoghue, T. and M. Rabin. (2001). “Choice and Procrastination”, *Quarterly Journal of Economics*, 116 (1), 121-160.

- [13] O'Donoghue, T. and M. Rabin. (1999). "Incentives for Procrastinators", *Quarterly Journal of Economics*, 114 (3), 769-816.
- [14] O'Donoghue, T. and M. Rabin. (1999). "Doing it Now or Later", *American Economic Review*, 89 (1), 103-125.
- [15] Pearce, D. (1984). "Rationalizable Strategic Behavior and the Problem of Perfection", *Econometrica*, 52, 1029-1050.
- [16] Phelps, E. S. and R.A. Pollak. (1968). "On Second-Best National Saving and Game-Equilibrium Growth", *Review of Economic Studies*, 35 (2), 185-199.
- [17] Pollak, R. A. (1968). "Consistent Planning", *Review of Economic Studies*, 35 (2), 201-208.
- [18] Rubinstein, A. (1982). "Perfect Equilibrium in a Bargaining Model", *Econometrica*, 50 (1), 97-110.
- [19] Sarafidis, Y. (2004). "Inter-temporal Price Discrimination with Time Inconsistent Consumers", *working paper*.
- [20] Strotz, R. H. (1956). "Myopia and Inconsistency in Utility Maximization", *Review of Economic Studies*, 23 (3), 165-180.
- [21] Wilson, R. (1987). "Game-Theoretic Approaches to Trading Processes", in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley, chap. 2, pp. 33-77, Cambridge University Press, Cambridge.