Noname manuscript No.
(will be inserted by the editor)

Predicting Human Cooperation in the Prisoner's Dilemma Using Case-based Decision Theory

Todd Guilfoos · Andreas Duus Pape

Received: date / Accepted: date

Abstract In this paper we show Case-based Decision Theory (Gilboa and Schmeidler, 1995) can explain the aggregate dynamics of cooperation in the repeated Prisoner's Dilemma, as observed in the experiments performed by Camera and Casari (2009). Moreover, we find CBDT provides a better fit to the dynamics of cooperation than does the existing Probit model, which is the first time such a result has been found. We also find that humans aspire to a payoff above the mutual defection outcome but below the mutual cooperation outcome, which suggests they hope, but are not confident, that cooperation can be achieved; and we find that cognitive capacity, and not memory storage space, is the more significant constraint on human problem-solving in this context. This is the first empirical test of Case-based Decision Theory in a strategic, dynamic, and multi-agent experiment. This provides a different frame of reference to view the results of cooperation in the iterated prisoner's dilemma that incorporates past experience as well as beliefs of future payoffs.

JEL codes: D01, D83, C63, C72, C88

Keywords Case-based Decision Theory · Prisoner's Dilemma · Learning · Agent-based Computational Economics · Experimental Economics

The order of the authors is not indicative of effort or contribution towards this article.

T. Guilfoos
Department of Environmental and Natural Resource Economics, University of Rhode Island, 219 Coastal Institute, 1 Greenhouse Road, Kingston, RI 02881
E-mail: guilfoos@mail.uri.edu

A. D. Pape
Economics Department, Binghamton University, PO Box 6000, Binghamton NY, 13902
Phone: 607-777-2660 Fax: 607-777-2681
E-mail: apape@binghamton.edu

1 Introduction

In this paper, we use Case-based Decision Theory (Gilboa and Schmeidler, 1995) to explain experimental data of human behavior in the repeated Prisoner's Dilemma game, and find that the aggregate dynamics are predicted by this theory. We are interested in explaining human behavior and we seek to match average cooperation rates over time. We establish that Case-based Decision Theory explains cooperation dynamics by comparing experimental data collected by Camera and Casari (2009) against simulated data generated by a computer program called the Case-based Software Agent (CBSA), which was introduced in Pape and Kurtz (2013). Pape and Kurtz show that CBSA (and therefore Case-based Decision Theory) explains individual human behavior a series of classification learning experiments from Psychology starting with Shepard, Hovland, and Jenkins (1961). Here we show that CBSA can explain human group behavior in a setting that is dynamic and strategic.

Case-based Decision Theory is a mathematical model of choice under uncertainty which has the following primitives: A set of *problems* or circumstances that the agent faces; a set of *actions* that the agent can choose in response to these problems; and a set of *results* which occur when an action is applied to a problem. Together, a problem, action, and result triplet is called a *case,* and can be thought of as one complete learning experience. The agent has a finite set of cases, called a *memory*, which it consults when making new decisions. The Case-based Software Agent is a *software agent*, i.e. "an encapsulated piece of software that includes data together with behavioral methods that act on these data Tesfatsion (2006)." CBSA computes choice data consistent with an instance of CBDT for an arbitrary choice problem or game, provided that the problem is well-defined and sufficiently bounded.

We analyze data from an experiment by Camera and Casari (2009), in which individuals are grouped into small 'economies' to play the repeated Prisoner's Dilemma. The purpose of their experiment is to vary the level of information available to players and measure the effect on cooperation. For example, in one treatment, the players are supplied unique identifiers for their opponents, so they know when they encounter the same opponent again. Because CBDT encodes the agent's information about the current choice directly (in the aforementioned problem variable), this is a particularly appropriate experiment to test with CBSA. We compare the simulated data to the real data and compute the mean squared difference in probability of cooperation over time. Like regression analysis, we then search the space of free parameter values for those that provide the best fit (minimizing mean squared error).[1]

We are able to establish four key facts about CBDT and its relationship with human choice behavior in Camera and Casari's Prisoner's Dilemma experiment.

We find:

(1) The choice behavior of this software agent (and therefore Case-based Decision Theory) correctly predicts the empirically observed trajectory of average cooperation rates over time across three different treatments. This shows that CBDT can predict human behavior in a strategic and dynamic setting.

(2) The choice behavior implied by CBSA is a closer fit to the empirical data than the best-fitting Probit model, and CBSA has only a fifth as many parameters to fit to data. This is a vote in favor of CBSA as a useful empirical description of human behavior in the repeated Prisoner's Dilemma and is a novel result in the literature.

(3) The best-fitting CBSA parameters suggest humans aspire to payoff value above the mutual defection payoff but below the mutual cooperation payoff, which suggests they hope, but are not confident, that cooperation can be achieved. In principle, the best-fitting aspiration values could have fallen into

[1] The free parameters of CBSA include two kinds of forgetfulness and an aspiration level. See Section 4.1 for details.

the 'unreasonable' range: namely greater than the best or lower than the worst possible payoff. The fact that this did not happen serves as an specification test of CBSA.[2]

(4) Consistent with Pape and Kurtz (2013), we find evidence that suggests that cognitive capacity, and not memory storage space, is the more significant constraint on human problem-solving. The evidence is that the best-fitting parameter values suggest that memory is imperfect in the following way: a relatively high fraction of events are encoded in memory ($78\% - 90\%$), but agents access a smaller percentage of stored, relevant event memories ($45\% - 70\%$) when they call upon their memory in a utility calculation.

These findings are useful in understanding the behavior of human subjects as well as developing a framework in which we can predict human behavior. For example, the infinitely iterated prisoner's dilemma can sustain cooperation when sufficiently patient agents employ a GRIM strategy, defecting forever if their partner defects, but this strategy does not seem to be played by human subjects. This paper is can be thought of as part of an effort to find alternative explanations of decision making which are more empirically valid.

Below, we review the relevant literature in Decision Theory, Game Theory, and the empirical study of the Prisoner's Dilemma (Section 2); we define CBSA precisely, show how it implements Case-based Decision Theory (CBDT) with imperfect memory (Section 3); we then explicitly describe the experiment of Camera and Casari (2009) and explain how we replicate it in CBSA (Section 4). Then we present and discuss our empirical results (Section 5) and conclude with some implications for future work (Section 6).

2 Related Literature

The central investigative tool of this paper is the case-based software agent (CBSA). It is a computational implementation of Case-Based Decision Theory (CBDT) introduced in Gilboa and Schmeidler (1995). Implementations produce agent choice behavior given a mathematical representation (von Neumann and Morgenstern, 1944; Savage, 1954). Designed correctly, the choice behavior produced by an implementation can be compared directly to empirical choice data of the same problem faced by humans. This can yield two classes of insights: First, the comparison can shed light on the question of whether and in what ways CBSA (and therefore CBDT) serves as a representation or 'explanation' of human behavior. Second, the comparison can shed light on the empirical phenomenon itself: for example, we learn what level of forgetfulness is consistent with the human behavior observed in the experiments found in Camera and Casari (2009).

Case-based Decision Theory—hereafter, CBDT—postulates that when an agent is confronted with a new problem, she asks herself: How similar is today's case to cases in memory? What acts were taken in those cases? What were results? She then forecasts payoffs of actions using her memory, and chooses the action with the highest forecasted payoff. The primitives are: a finite set of problems $\mathcal{P}$, a finite set of acts $\mathcal{A}$, and a set of results $\mathcal{R}$. A case is a triplet consisting of a problem, the act taken, and the outcome (result) of that act given the problem. A case can be thought of as a single, complete learning experience. The set of all cases is $\mathcal{C} = \mathcal{P} \times \mathcal{A} \times \mathcal{R}$.

CBDT representations are defined by four components.

The first component of a CBDT representation is the agents' memory. Memory is a set of cases which, in CBSA, can be thought of as the list of learning experiences the agent has had. An agent's memory is denoted $\mathcal{M}$.

[2] Like other specification tests, passing the test does not mean that the model is necessarily correctly specified; only that failing the test would have been evidence that it is misspecified.

The second component is the utility function $u : \mathcal{R} \to \mathbb{R}$. It is defined in the usual way.

The third component is the similarity function $s : \mathcal{P} \times \mathcal{P} \to [0, 1]$. The output value of the similarity function gives how much the input problems resemble each other in the opinion of the agent.

The fourth component is the aspiration level $H \in \mathbb{R}$. It is a reference level of utility, like expected value. However, while an expected value is the level of utility one believes on average is the most likely, the aspiration level should be thought of as the agent's *target* level of utility, which could, in principle, differ from the expected value. Mechanistically, it serves as a default value for forecasting utility of new alternatives. It also serves as a satificing level in the sense of Simon (1957): "Behaviorally, H defines a level of utility beyond which the [decision maker] does not appear to experiment with new alternatives (Gilboa and Schmeidler, 1996, page 2)."

Together, these four components define case-based utility:

$$CBU(a) = \sum_{(q,a,r)\in\mathcal{M}(a)} s(p, q)\,[u(r) - H]$$

Where $\mathcal{M}(a)$ is defined as the subset of the agents' memory $\mathcal{M}$ in which action a was taken, and $H \in \mathbb{R}$ is an aspiration level (see below). This utility represents the agent's preference in the sense that, for a fixed memory $\mathcal{M}$, a is strictly preferred to a' if and only if $CBU(a) > CBU(a')$.

Case-based Decision Theory was introduced for the main purpose of disposing of the state space: that is, the assumption that agents are able to list and reason about the set of all possible scenarios. Case-based Decision Theory limits the set the agent must reason about to the set of past experiences, and requires only that the agent be able to make similarity judgements between past experiences and new experiences. Therefore Case-based Decision Theory naturally incorporates cognitive constraints. Moreover, the implementation of CBSA here in this paper also includes forgetfulness explicitly.[3]

This paper contributes to a growing empirical literature testing the explanatory power of Case-based Decision Theory; these papers generally find support for CBDT.

In the paper most closely related to this one, Pape and Kurtz (2013) introduce CBSA, and find that imperfect memory, accumulative (not average) utility, a similarity function consistent with research from psychology, and a $80 - 85\%$ target success rate renders CBSA a good fit for human data in the classification learning experiment from the psychology literature.[4]

Ossadnik et al (2012) run a repeated choice experiment involving unknown proportions of colored and numbered balls in urns, which is the canonical ambiguous choice setting (i.e. Ellsberg, 1961). This is important because the authors of CBDT suspect that it is a better model of human behavior in settings of ambiguity versus risk. Ossadnik et al (2012) find that CBDT explains these data well compared to alternatives such as minimax (Luce and Raiffa, 1957) and reinforcement learning (Roth and Erev, 1995). Their method has some similarities with CBSA, in that they choose parameter values and functional forms of CBDT and calculate CBDT-governed agents' optimal choices, and compare those choices to aggregate human data. There are two important ways that the method differs from CBSA. First of all, the CBSA method sweeps parameter values, so provides many more candidate values for fitting the human data. Second, CBSA integrates forgetfulness and similarity functional forms from Psychology.

Bleichrodt et al (2012) provide a method to measure similarity weights which avoids parametric assumptions about the weights. Their method has a number of advantages, including testing CBDT in

[3] Incorporating cognitive constraints into economic models is a hallmark of work in the intersection of economics and psychology, such as Simon (1957), Simon et al (2008), Tyson (2008), Hanoch (2002), Ballinger et al (2011), and Cappelletti et al (2011).

[4] In particular, the 'SHJ' series of classification learning experiments, starting with Shepard, Hovland, and Jenkins (1961) and including Nosofsky (1986) and Nosofsky et al (1994).

more generality. An advantage of the CBSA approach is the ability to predict CBDT-governed behavior on arbitrary settings; the Bleichrodt et al (2012) approach implies a particular kind of experimental design. Therefore, we feel these methods are complementary: insights developed in their method can be applied to CBSA for application to other settings, for example.

Gayer et al (2007) investigate whether case-based reasoning appears to explain human decision-making using housing sales and rental data. They hypothesize and find that that sales data are better explained by rules-based measures because sales are an investment for eventual resale and rules are easier to communicate, while rental data are better explained by case-based measures because rentals are a pure consumption good where communication of measures are irrelevant.

Golosnoy and Okhrin (2008) investigate using CBDT to construct investment portfolios from real returns data and compare the success of these portfolios to investment portfolios constructed from expected-utility-based methods, and find some evidence that using CBDT aids portfolio success.

The Prisoner's Dilemma (PD) is perhaps the most famous game in game theory. It is a symmetric, simultaneous, two-player game with two actions, Cooperate and Defect, where (1) Defect strictly dominates Cooperate, but (2) the payoff for (Cooperate, Cooperate) Pareto dominates (Defect, Defect). Although there are benefits to cooperation, the individual incentive to defect means that mutual cooperation is not a Nash Equilibrium. Instead, (Defect, Defect) is the unique Nash Equilibrium. Because of this tension between what is best for the group versus the individual, the repeated Prisoner's Dilemma is used as a metaphor for cooperation in general and has been used to represent public goods problems, common pool resource depletion, and international politics. Much of the theoretical and empirical investigation into the repeated Prisoner's Dilemma has been about the question: when does cooperation occur and when is it sustainable?

There are many reasons why individuals might choose to cooperate, such as reputation building, altruism, or fear of reprisal.[5]

Experimentalists have been investigating the causes of cooperation through different treatments in experiments for some time now.[6] Most relevant to our investigation today, Camera and Casari (2009) show that punishment and information of past play history can lead to higher sustainable levels of cooperation. They show this by experimentally varying the level of information available to players and measuring the resulting levels of cooperation. We attempt to explain their data with CBSA.

The central investigative tool of this paper is a software agent, therefore this paper is part of Agent-based Computational Economics. There is a long history of using computational agents to explore the repeated Prisoner's Dilemma (PD): Axelrod (1980) ran a series of tournaments where academics and computer programmers submitted strategies that play against each other in a repeated PD in one of the earliest and most famous agent-based investigations. Similarly, Miller (1996) explores the evolution of

[5] The fear of reprisal was formalized in the famous Folk Theorem first suggested by Friedman (1971), where players cooperate in the infinitely repeated Prisoner's Dilemma in a sustainable equilibrium. The reprisal the players fear is that their opponent will defect in all future periods, therefore reverting to the suboptimal Nash outcome. This has been shown to be stable when players are sufficiently patient (Fudenberg and Maskin, 1986).

[6] The experimental literature on the PD and other repeated games is vast. A sample follows. Brosig (2002) show that signaling in face-to-face experiments may be effective at encouraging cooperation. In one-shot games there exists a low level of cooperation (Bereby-Meyer and Roth, 2006). Agents may learn to cooperate in repeated games, especially when monitoring of other players actions is allowed (Selten and Stoecker, 1986; Andreoni and Miller, 1993; Hauk and Nagel, 2001) but that cooperation breaks down during the the course of game. Evidence on the altruism motivation is mixed, some finding evidence for (e.g. Kreps et al, 1982) and some against (e.g. Cooper et al, 1996). Other papers of note include Ellison (1994); Bó (2005) and Bo and Fréchette (2011). Chong et al (2006) and Camerer and Hua Ho (1999) use an 'experience weighted attraction' model to study learning in a repeated trust game. This model postulates that players remember the history of previous play and form beliefs about what other players will do in the future and also are reinforced by how successful previous strategies have been. Monterosso (2002) uses false feedback to disrupt equilibria and measures the effects; this work seems also well-suited for analysis by CBSA and analyzing these data are a possible future extension.

strategies when computational agents pick from a predetermined set. They investigate which strategies survive over repeated play of the PD as information is varied. This strain of literature has typically involved agents following simple strategies that are tailor-made for this application, such as tit-for-tat or the grim trigger strategy. The contribution of CBSA to this literature, other than its striking empirical fit, is that CBSA is a general choice engine that can be used in other games and decision problems.

3 The Case-based Software Agent

The Case-based Software Agent was introduced in Pape and Kurtz (2013) and described here only briefly. For a more detailed presentation, we refer you that paper.

CBSA shares all primitives with CBDT: that is, a single choice problem experience is viewed as a triplet (p, a, r), where $p \in \mathcal{P}$ is a 'problem' or circumstance (formally: a vector of exogenous variables observable to the agent when the choice is made), $a \in Acts$ is an action taken in response to the circumstance, and $r \in \mathcal{R}$ is a the result or outcome of the action applied to the problem. CBSAs also have a utility function $u : \mathcal{R} \to \mathbb{R}$, which has the usual definition, and a similarity function $s : \mathcal{P} \times \mathcal{P} \to [0, 1]$ which describes how similar two circumstances are (in the mind of the CBSA); and an aspiration level $H \in \mathbb{R}$, which represents a target level of utility of the agent. In addition, and introduced here, CBSAs have an act randomization probability $\alpha \in [0, 1]$, which we define below.

Along with the decision primitives, CBSA defines the decision environment: i.e. those parts of the choice problem that are external to the agent.[7] In CBSA the decision environment is represented by function (algorithm) called the *problem-result map* or PRM. The PRM is the transition function of the environment. It takes as input the current problem $p \in \mathcal{P}$ the agent is facing, the action $a \in \mathcal{A}$ that the agent has chosen, and some vector $\theta \in \Theta$ of environmental characteristics. The PRM returns the outcome of these three inputs: namely, it returns a result $r \in \mathcal{R}$; the next problem $p' \in \mathcal{P}$ that the agent faces; and a potentially modified vector of environmental characteristics $\theta' \in \Theta$. I.e.:

$$PRM : \mathcal{P} \times \mathcal{A} \times \Theta \to \mathcal{R} \times \mathcal{P} \times \Theta$$

For example, in Savage's omelette problem (Savage, 1954, pp. 13-15), an agent chooses how to crack possibly rotten eggs for an omelette. The environment of the problem is the set of eggs that the agent is drawing from. The PRM is the function which, when called, serves up the next egg for cracking.

Figure 1 describes the choice algorithm which implements the core of CBSA, where the choice C represents Cooperation and D represents Defect. It is an algorithmic description of the choice process defined by CBDT, with two modifications. The modifications allow for imperfect memory. In Pape and Kurtz (2013), it was found that a match between CBSA and human data was only achieved by allowing for imperfect memory: otherwise CBSA solves the classification learning problem much faster than humans. We find that imperfect memory is also important for matching human data in this setting.

There are two kinds of imperfect memory. First, there is *imperfect recall*, governed by a probability $p_{\text{recall}} \in [0, 1]$. Imperfect recall corresponds to an inability to access all memory at any given time, and it is therefore associated with limited cognitive capacity. Second, there is *imperfect storage*, governed by a probability $p_{\text{store}} \in [0, 1]$. Imperfect storage corresponds to a failure to add some experiences to memory after they are experienced, and it is therefore associated with limited memory storage capacity.

[7] These need not be defined for a decision theory, so are not a formal part of CBDT, and need only be formally defined when one seeks to generate simulated choice behavior to compare to empirical data.

Input: problem p, memory $\mathcal{M}$.

1. For each $a \in \mathcal{A}$:
   (a) For each $(q, a, r) \in \mathcal{M}$, draw r.v. $b_{(q,a,r)} = \begin{cases} 1, & \text{with probability } p_{\text{recall}} \\ 0 & \text{otherwise.} \end{cases}$

$$\text{Construct } \mathcal{M}_a = \big\{(q, a, r)\big| b_{(q,a,r)} = 1, \text{ AND } \exists q \in \mathcal{P}, r \in \mathcal{R} \text{ s.t. } (q, a, r) \in \mathcal{M}\}$$

   (b) Let $U_a = \begin{cases} \sum_{(q,a,r)\in\mathcal{M}_a} s(p, q)\big[u(r) - H\big], & \text{if } \mathcal{M}_a \neq \emptyset \\ 0, & \text{otherwise} \end{cases}$
2. Construct set $BEST = \big\{a \in \mathcal{A}\big|U_a = \max_{b\in\mathcal{A}} \{U_b\}\big\}$
3. If $\#(BEST) = 1$ then let $a^\star$ be the sole entry in $BEST$.
   If $\#(BEST) = 2$, then C is chosen with probability α, D with probability $(1 - \alpha)$.

Output: Selected action $a^\star$

Fig. 1: The Choice Algorithm

In Figure 1, the agent faces a problem $p \in \mathcal{P}$ and has a memory $\mathcal{M} \subseteq \mathcal{C}$. In Step 1a, for each action a, she collects those cases in which she performed this act. Since her recall is imperfect, relevant cases are selected into the set $\mathcal{M}_a$ with probability p_{recall}, where relevant cases which are not recalled are simply ignored.[8] In Step 1b, she uses this subset of her memory $\mathcal{M}_a$ to construct a utility forecast of that act, called here U_a. The agent then chooses the action which corresponds to the maximum U. As seen in step 3, when there is a tie between C and D, C is chosen with an exogenous probability α. In the original formulation, it was assumed that $\alpha = .5$, but in this study we calibrate α to data (see Section 4.1 for calibration details).

Input: problem p, memory $\mathcal{M}$, characteristics θ.

1. Input p,$\mathcal{M}$ into choice algorithm (Figure 1). Receive output $a^\star$.
2. Let $(r, p', \theta') = PRM(p, a^\star, \theta)$.
3. With probability p_{store},
   Let $\mathcal{M}$' $= \mathcal{M} \cup \{(p, a^\star, r)\}$
   Else let $\mathcal{M}$' $= \mathcal{M}$

Output: problem p', memory $\mathcal{M}'$, characteristics θ'.

Fig. 2: A Single Choice Problem.

Figure 2 describes a single choice problem faced by the agent. It imbeds a reference to the choice algorithm described in Figure 1. Figure 2 embeds the agent in an environment and explicitly references that environment, in the call to PRM. In Step One, the agent selects an act, $a^\star$. In Step Two, the action is performed, in the sense that the environment of the agent reacts to the agent's choice: the PRM takes the current problem p, the action $a^\star$ selected by the agent, and the characteristics unobserved by the agent θ, and constructs a result r, a next problem p', and a next set of characteristics θ'. In Step Three, the agent's memory is augmented by the new case which was just encountered, so long as the agent does

[8] When $p_{\text{recall}} = 1$, the agent has perfect recall. It then corresponds to CBDT as it appears in Gilboa and Schmeidler (1995).

not have a 'write-to-memory error:' i.e., with probability p_{store}, the case that was just experienced is added to the set $\mathcal{M}$. With probability $(1 - p_{\text{store}})$, that case is discarded.

Since the choice problem depicted in Figure 2 maps a problem, a characteristic, and a memory vector to another vector in the same space, it can be applied iteratively. A series of such iterations, along with initial conditions and ending conditions, can then be used to produce a single time series of agent behavior, called a 'run.' Here, the initial conditions specify that agents have an empty memory, although it is simple to modify the algorithm such that the agent starts with some non-empty memory. The ending conditions can take on a variety of forms, and can be exogenous or endogenous. Often times the ending condition is simply a predetermined number of periods or a fixed probability that the cycle ends. That second condition is the one used in Camera and Casari (2009).

4 The Camera and Casari (2009) Experiment and CBSA

In this section we describe the details of the Camera and Casari experiment and how this experiment maps to the implementation of CBSA used here. Then, we describe how the simulated data are generated and how we compare them to the human data. We compare them along three dimensions: qualitative fit, quantitative fit, and model complexity.

		Player 2 C	Player 2 D
Player 1	C	25,25	5,30
Player 1	D	30,5	10,10

Fig. 3: The Prisoner's Dilemma used in Camera and Casari (2009)

4.1 Experiment details and implementation in CBSA

We analyze data from the experiment in Camera and Casari (2009). The specific parameterization of the Prisoner's Dilemma used in this experiment is shown in Table 3.[9] As usual, the first payoff listed is for player 1 and the second is for player 2. The treatments in the experiment vary the information information available to subjects. These treatments are ideal candidates for testing with CBSA. The treatments in the experiment imply different definitions of the set of problems $\mathcal{P}$. Therefore this experiment provides a set of *a priori* hypotheses that are empirically testable: that under the varying definitions of the set $\mathcal{P}$ defined by the treatments, CBSA's level of cooperation will move in tandem with humans.

In the experiment, the human subjects are put into 4-person groups. A group is called an 'economy.' Each economy plays one 'supergame.' A 'supergame' is a series of PD games among the four players, played for a random number of periods. Each period, the the 4 subjects are randomly paired and play PD.[10] After both pairs play and payoffs are given, with a fixed probability $(1 - \delta)$, the supergame immediately ends; and with continuation probability δ the game is played again. This repeats until continuation fails. Camera and Casari (2009) set $\delta = 0.95$ which implies that at all times, the conditional expectation is that there will be 20 more periods of play.

[9] See Camera and Casari (2009) for details about why these payoff values were selected.
[10] Randomly paired with a uniform probability.

We follow the payoff structure and the sorting of agents within a supergame as described above. In order to generate comparable data, we observe agents in games for 30 periods. We begin each supergame with agents that have a null memory. One simulation run is a supergame with four CBSAs that are randomly paired together each time period for 30 periods. We generate 1,000 runs for each treatment and evaluate the simulated data we have generated against the experimental data measured by Camera and Casari.

The definition of the treatments are *a priori* hypotheses. How to define the problem sets is a matter of interpretation; that is, the treatments do not *uniquely* identify corresponding problem sets. The following should be thought of as a set of testable hypotheses. Importantly, these hypotheses were chosen *a priori,* i.e. without regard for goodness-of-fit.

The treatments that the subjects were exposed to are designed to investigate how the level of anonymity influences their decision to cooperate. The level of information expands the set of possible cooperative equilibrium in the games, meaning we expect that having public information of players strategies allows a greater possibility for a cooperative equilibrium than when information is private. We investigate three of the four treatments tested by Camera and Casari.[11] In all treatments, we reason that players are aware of how far into the supergame they are; so all treatments' problem vectors include the period of play t.

Treatment 1 is *private monitoring*, which consists of anonymous subjects playing the supergame with no information about the player they are paired against or the players in the rest of the economy. Since no other information is available, $\mathcal{P}_1 = T$, with typical element $p_1 = (t)$, where $t \in T = \{1, 2, \ldots, 30\}$.

Treatment 2 is *anonymous public monitoring*, which gives information about the history of other players, including the highlighted history of ones current opponent. However, explicit identifiers are not available. In this treatment, we reason that the relevant information is the average cooperation rate of ones current opponent. So $\mathcal{P}_2 = T \times [0, 1]$, with typical element $p_2 = (t, \overline{a}(\theta'))$, where the dimension corresponding to the interval $[0, 1]$ is the average cooperation rate of the opponent, where θ' is the identity of the opponent, and where $\overline{a}(\theta')$ is the average cooperation rate of opponent θ'.

Treatment 3 is *non-anonymous public monitoring*, which consists of the information available in Treatment 2 as well as a unique ID of their opponent. We represent this as a vector of three binary variables (id_1, id_2, id_3), where at any time exactly one id variable is 1 and the others are 0. Therefore $\mathcal{P}_3 = T \times [0, 1] \times \{0, 1\} \times \{0, 1\} \times \{0, 1\}$, with typical element $p_3 = (t, \overline{a}(\theta'), id(\theta'))$, where θ' is the identity of the opponent, where $\overline{a}(\theta')$ is the average cooperation rate of opponent θ', and where $id(\theta')$ is a string of dummy variables that indicate the identity of opponent θ'.

In the definition of CBSA, the set of actions and the set of results must also be defined, but they are straight-forward. The set of actions is $\{C, D\}$ and the set of results is $\mathbb{R}_+$.[12] There is also an aspiration level H, which can be thought of as the target payoff from any play of the game.

In CBSA, the Problem-Result Map, or PRM (see Section 3), takes as input {an action $a \in \mathcal{A}$, a problem $p \in \mathcal{P}$, and a set of environmental characteristics $\theta \in \Theta$,} and delivers {a result in $r \in \mathcal{R}$, a 'next problem' p' the agent is to face, and a potentially modified vector of environmental characteristics $\theta' \in \Theta$}. In general in CBSA, θ describes the current state of the environment of each agent, i.e. exogenous, unknown forces that are acting on the agent.[13] In this case, the environment of the player is the identity of the opponent. Given this definition, the PRM can be defined as in Figure 4.

[11] The excluded treatment involves an additional stage game after the PD is played. We hope to study this treatment in a future extension.

[12] Arguably, the set of results is just the list $\{5, 10, 25, 30\}$. However, allowing for other values in the set of results changes nothing.

[13] *Known* exogenous forces acting on the agent are part of the problem vector.

Input: act a, problem p, characteristics θ.

1. Find action of opponent $a(\theta)$.
2. Set result r according to the game depicted in Figure 3, with actions $(a, a(\theta))$.
3. Next opponent is chosen θ'.
4. Set next problem p', by constructing the vector:

$$p' = \begin{cases} (t), & \text{if Treatment 1} \\ (t, \overline{a}(\theta')), & \text{if Treatment 2} \\ (t, \overline{a}(\theta'), id(\theta')), & \text{if Treatment 3} \end{cases}$$

Where:
- t is the time period,
- $\overline{a}(\theta')$ is the average cooperation rate of opponent θ', and
- $id(\theta')$ is a string of dummy variables that indicate the identity of opponent θ'

Output: result r, problem p', characteristics θ'.

Fig. 4: The Camera and Casari Prisoner's Dilemma Experiment PRM

There are five parameters of CBSA: two are set according to theory and three are estimated.

The first parameter set according to theory is the problem vector p, which was described above. The second is α, the probability of cooperating when indifferent. A brief theoretical analysis proves that α will be equal to the population average cooperation level in the first round of play. The logic is: at the beginning of the supergame, all agents are indifferent between C and D and therefore randomize with probability α. Therefore, *a priori* the first round of play will, on average, yield a fraction α of cooperators, which of course turns out to be true in the simulated data.[14]

The first and second estimated parameters are the imperfect memory parameters, namely p_{store}, which is the probability that an individual case is written to memory, and p_{recall}, which is the probability that an individual case is recalled when memory is accessed. The third estimated parameter is the aspiration level H, which is the target payoff level sought by the individuals.[15]

4.2 Choosing benchmark models

The purpose of a benchmark model is to present a version of CBSA which best explains the human data. There are three criteria we use to choose a model: qualitative fit, quantitative fit, and model complexity. Below, we describe these three criteria and then discuss the method by which we search alternative versions of CBSA for the 'best' model by these criteria.

Qualitative fit to human data is equivalent to matching "stylized facts" of human data. For example, we find that, empirically, Treatment 3 maintains a higher cooperation rate than Treatments 1 and 2 in all periods. A model which matches more of these regularities is said to have a greater qualitative fit.

[14] An alternative modeling choice is to posit a two types of agents: ones who always cooperate when indifferent and ones who always defect. Then one calibrates the relative size of the two populations to the known cooperation rate in the first round. This alternative modeling strategy does not make a large difference in the results so it is not presented here.

[15] There are also parameters we do not choose to vary. For example, we do not vary the functional form of similarity: we choose only the 'accumulative' form of similarity over 'average similarity,' as average similarity is found in Pape and Kurtz (2013) to cause the counterfactual behavior of individuals believing actions to be irrelevant. We also do not vary from the functional form of inverse exponential weighted Euclidean distance. There is a strong empirical case for this functional form in psychology, which is bolstered by the results found in Pape and Kurtz (2013). Please see that paper for more details.

Quantitative fit is a numeric evaluation of fit to human data: for any given set of simulated data, we construct Mean Squared Error (MSE) between the simulation average cooperation rate and human average cooperation rate. The lower the MSE, the better the quantitative fit. Even though perfect quantitative fit implies perfect qualitative fit, in practice quantitative fit can come at the cost of qualitative fit.

The model complexity criterion is also known as 'overfitting' in econometrics or 'model elegance' in theory. This third criterion rests on the observation that if a model is allowed to be arbitrarily complicated, then perfect qualitative and quantitive fit can be achieved, but such a model may be undesirable because it doesn't reveal insight into the phenomenon and is not generalizable out-of-sample. The 'model complexity' consideration leads us to select a simpler model over a more complicated one.

If a model has too many free parameters, one might be concerned that one is not truly 'testing' the model, and instead only running a 'calibration exercise.' So this criterion states that fewer free parameters, which must be fit to data, the better.[16]

Like linear regression, we seek a set of parameters of a mathematical model that best fit observed data by minimizing mean squared error. Unlike regression, there is no known closed-form function from the observed data to the parameters of CBDT.[17] Because there is no simple function, we repeatedly run CBSA with different parameter values, generate simulated data, and then measure those simulated data against the human data according to mean squared error (MSE), choose new parameter values, and repeat. Specifically, we explore the parameter space through an iterated grid-search: the parameter space is swept at a certain resolution, generating 1000 simulation runs for each parameter combination. Then, the part of the parameter space which contains the best fitting models is explored at a higher resolution, again with 1000 simulation runs per parameter combination. This is repeated until it appears we exhaust measurable improvements in MSE.

5 Experimental Results

In this section we present the results of two versions of CBSA and compare them to the Probit found in Camera and Casari (2009) and a constrained alternative Probit formulation. The purpose of this comparison is to validate CBSA: we seek to empirically 'benchmark' CBSA against some alternative explanatory model. We compare the four models by the model selection criteria described in the previous section: quantitative fit, qualitative fit, and model complexity. An overview of the results can be seen in Table 1. In this table, we provide parameter estimates for the two CBSA versions and degrees of freedom and goodness-of-fit (MSE) of all four available models. The predicted outcome variable, average cooperation rates over time, for all four models as well as the raw data are shown in Figure 5.

We are able to show (1) that the CBSA correctly predicts the empirically observed trajectory of average cooperation rates over time across different treatments, and (2) that the choice behavior implied by CBSA is a closer fit to the empirical data than either Probit model. We also find that the best-fitting parameters suggest (3) humans aspire to payoff value above the mutual defection outcome but below the mutual cooperation outcome, which suggests they hope but are not confident that cooperation can be

[16] For another point of view on this criterion, consider two competing theories that attempt to explain the same phenomenon: theory a and theory b. Suppose that there is some empirical data available about the phenomenon. Suppose theory a has many more free parameters than theory b does. Now suppose we calibrate theories a and b to the data, and we find, after calibration, that theory a and theory b both explain the same fraction of the variation in the data and explain the same qualitative phenomena. Then, under the model complexity criterion, theory b is preferred.

[17] I.e. the equivalent to $\beta = \frac{X'X}{X'Y}$

Fig. 5: Cooperation Rates over Time: Experimental Data Versus Alternative Models.

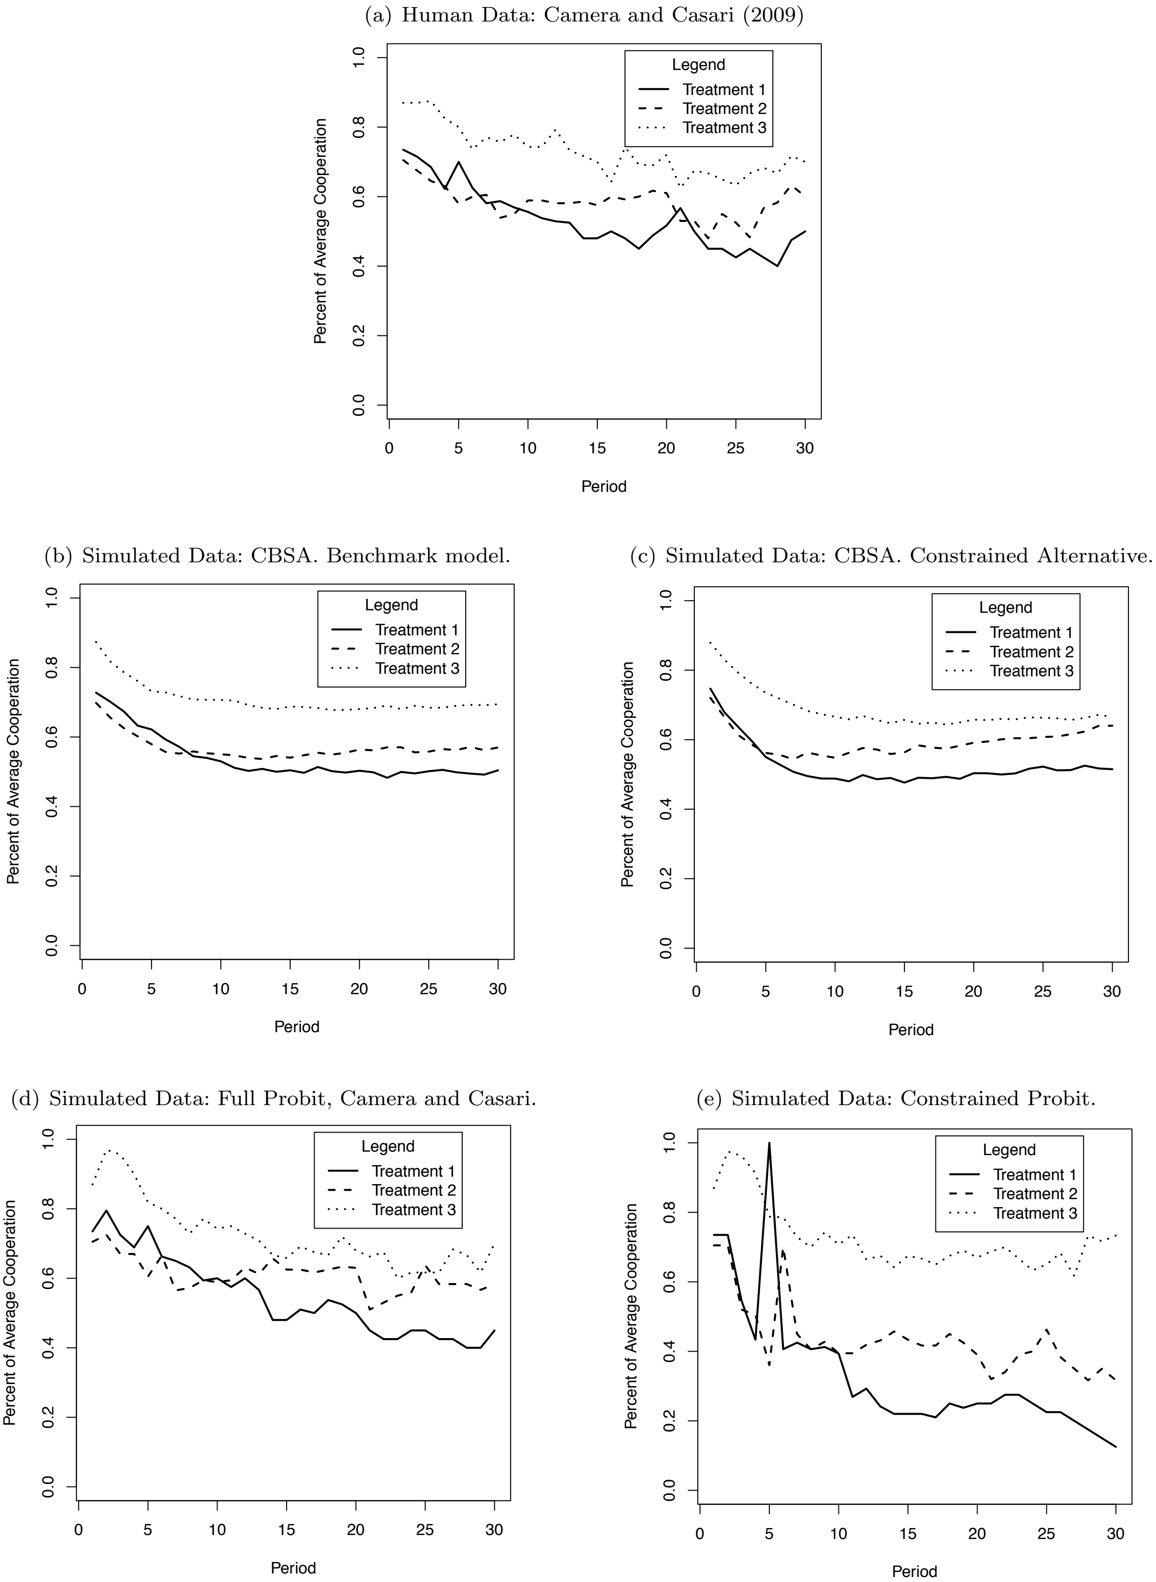

Table 1: CBSA and Probit Regression Predicting Average Cooperation Rates

	Benchmark CBSA	Constrained CBSA	Full Probit	Constrained Probit
Treatment 1				
H	10.8	10.5		
p_{recall}	0.39	0.6		
p_{store}	0.75	0.8		
α	0.735	0.735		
p	(t)			
Treatment 2				
H	14.8	14.8		
p_{recall}	0.45	0.6		
p_{store}	0.8	0.8		
α	0.705	0.705		
p	$(t, \overline{a}(\theta'))$			
Treatment 3				
H	12	13		
p_{recall}	0.7	0.6		
p_{store}	0.9	0.8		
α	0.87	0.87		
p	$(t, \overline{a}(\theta'), id(\theta'))$			
Other Estimated Parameters	None	None	See Camera & Casari (2009)	See Section 7
Number of Parameters	15	11	68	25
Mean Sq. Error (MSE)	0.0019	0.0033	0.0022	0.0289
Indiv. Fixed Effects?	No	No	Yes	No
Cycle Fixed Effects?	No	No	Yes	No
Predicted Cooperation	Fig. 5(b)	Fig. 5(c)	Fig. 5(d)	Fig. 5(e)

Notes:

α and p are set by theory and are not estimated. Only H, p_{recall}, and p_{store} are estimated.

In the definition of p,

t is the time period,

$\overline{a}(\theta')$ is the average cooperation rate of opponent θ',

$id(\theta')$ is a string of dummy variables that indicate the identity of opponent θ'.

p_{recall} and p_{store} are constrained to be identical across treatments in the Constrained CBSA.

achieved, and (4) that, consistent with Pape and Kurtz (2013), we find evidence that cognitive capacity, and not memory storage space, is the more significant constraint on human problem-solving.

Probit and CBSA are comparable in that they are somewhat similar methodologically: the coefficient estimates of a Probit emerge from an analytical maximization of a likelihood function given the data, and the CBSA parameter estimates emerge from a computational maximization of a quantitative measure of fit to data. In this sense, they are both predictive methods calibrated to data. Therefore, one could think that, if CBSA compares favorably to the Probit along some goodness-of-fit metric, then perhaps CBSA should be considered more seriously as an empirical method. On the other hand, if CBSA compares unfavorably to the Probit, it should be considered less seriously.

Note that this does not imply we seek to accept Probit over CBSA or the reverse. This is because CBSA and Probit could be, in a very real sense, both true. To see how this could work, suppose both the Probit and CBSA were good fits to data. One explanation could be that CBSA explains human learning behavior, and the memory, similarity, and utility of CBSA together encode strategies that the Probit measures, even though CBSA does not *explicitly* represent strategies.[18] Another explanation notes that Decision Theory is built on representation theorems: if behavior matches certain axioms, then a utility function, beliefs, etc., that represent that choice can be constructed. However, in decision theory, no representation theorem claims exclusivity: on the contrary, so long as the axiom sets of two representation theorems are not mutually exclusive, then the choice behavior can be represented by the structures in each theorem. So if the Probit and CBSA both seek to explain behavior in the 'representation theorem' sense, they could both be valid.[19,20]

This section proceeds in four parts: First, we describe the four models depicted here. Second, we compare the two CBSA models and the two Probit models along the dimensions of quantitative fit and model complexity using Table 1. Third, we compare the four models along the dimension of qualitative fit by using Figure 5. Steps two and three establish results (1) and (2) above, that CBSA fits well and compares favorably to the Probit. Fourth, we interpret the estimated parameters from the Benchmark and Constrained CBSA models and establish results (3) and (4) above.[21]

5.1 Model Descriptions

The Benchmark CBSA has 15 parameters, which corresponds to five per treatment. For each treatment, two parameters are selected according to theory and three are estimated. The parameters selected according to theory are the definition of the problem vector p and the likelihood of choosing cooperate when indifferent, α. The estimated parameters are: H, the aspiration level; p_{recall}, the probability that a given case is recalled from memory; and p_{store}, the probability that a given case is written to, or stored in, memory.

The Constrained CBSA is an alternative specification of CBSA. It comes from the following observation: given the fact that data were provided to the subjects of the experiment in all treatments in an

[18] This is possible if one notes that these computational structures of CBSA could 'encode strategies' in the way that a computer program encodes program behavior.

[19] Moreover, Matsui (2000) shows that Case-based Decision Theory and Expected Utility Theory can both represent the same choice behavior almost always. If the Probit can be likened to an expected utility perspective, then Matsui's result would suggest that both the Probit and CBSA could match.

[20] It is interesting to consider running the models on each others' outcomes: what strategies does a Probit suggest that the CBSA has encoded? And, in the other direction, what CBSA parameters emerge when CBSA seeks to predict the implied Probit strategies? We intend to consider these questions in a future extension.

[21] Camera and Casari provide their own interpretation of the Probit parameter estimates, so there is no need to reproduce it here.

identical way, perhaps memory is written to and accessed in an identical away across treatments. The Constrained CBSA formalizes this hypothesis by constraining that the probability of recall p_{recall} and the probability of storage p_{store} to be the same across all three treatments. (The aspiration level H, the initial rate of cooperation α, and the definitions of the problem vector vary across treatments.) As a consequence, the Constrained CBSA has 11 parameters total, only 5 of which are estimated.

The Full Probit refers to the Probit analysis which appears in in Camera and Casari (2009), Table 4, page 994. Camera and Casari propose three strategies players might use, and their Probit attempts to empirically identify the relative importance of these strategies in determining behavior. The three strategies are: reactive strategies, global strategies, and targeted strategies. Reactive strategies are: choosing to defect after ones opponent defects. Global strategies are: choosing to defect when any player in the economy defects. Targeted strategies are: choosing to defect against players who have defected against oneself, but ignoring defections against others. Each of these strategies is associated with a lag of one to five periods after a subject experiences a defection, so that the marginal response from a defection can change over those five periods. It also allows for players to abandon a defection after some time. Camera and Casari's Probit regression is designed to identify the marginal effects of these different strategies, where the binary outcome variable is the choice to cooperate and observations are people-periods. The probit also includes individual and cycle fixed effects, for a total of 68 estimated parameters.

We created the Constrained Probit as a variant to the Camera and Casari Probit. It is the same as the Full Probit, except that there are no cycle and individual fixed effects. The reason for the Constrained Probit is to make a fair comparison to CBSA along the following dimension: because CBSA does not fit parameters for different cycles and individual subjects, it can predict out-of-sample in its current form. However, the Full Probit allows independent fitting for individuals. These individual fixed effects would not be available for predicting out-of-sample. So the Constrained Probit is, in some sense, the best-fitting Probit which allows for the same freedom to predict out of sample as does CBSA. (The same reasoning could be applied to out-of-sample prediction for any econometric model. It guards against overfitting.) The Constrained Probit has 25 estimated parameters. The results of the Constrained Probit can be found in the Appendix Section 7.

For all four models–the two CBSAs and the two Probits–we construct the aggregate level of cooperation in each time period t:

$$CL_t = \sum_{i \in I_t} \frac{a_{i,t}}{N_t}$$

Where I_t is the set of subjects still playing at time t, N_t is the number of subjects still playing at time t, and where $a_{i,t}$ equals 1 if a player i chooses Cooperate and 0 if she chooses Defect. In the experiment, N_t varies with treatment and time period. In the CBSA results, N_t equals $4,000$ (we simulate $1,000$ economies, each with four players). To generate CL_t in the Full and Constrained Probit, we use the predicted value of the dependent variable, $\hat{Y}_i$, for each agent in the data. $\hat{Y}_i$ is the probability of cooperating. We record a binary value of 1 if $\hat{Y}_i \geq 0.50$ and 0 if $\hat{Y}_i < 0.50$. Then we construct the average cooperation level of all observations per period across economies as described above.

5.2 Model Comparison: Quantitative Fit and Model Complexity (Results 1 and 2)

Table 1 summarizes both the quantitative fit, MSE, and one measure of model complexity, the degrees of freedom, or number of estimated parameters.

With regard to quantitative fit, the order of Mean Squared Error (MSE) is:

$$\text{Benchmark CBSA} \leq \text{Full Probit} < \text{Constrained CBSA} << \text{Constrained Probit}$$

The Benchmark CBSA is a better fit than any other model available, although the Full Probit is close (only 20% larger MSE). The Constrained CBSA performs worse than the Benchmark CBSA (its MSE is 76% larger than the Benchmark CBSA and 48% larger than the Full Probit), which is not a surprise given that it has fewer free parameters. The Constrained Probit performs far worse than the others: its MSE is over 15 times that of the Benchmark CBSA and the other models.

These results are particularly striking when one considers the model complexity. The Full Probit is (slightly) worse at explaining the average cooperation rates despite having about five times the number of free parameters. The Full Probit, as opposed to the Constrained Probit, also takes advantage of individual fixed effects, which CBSA does not allow (CBSA assumes that all agents are ex-ante identical and differ only because of path dependence of experience over the course of the run.) Dropping individual fixed effects in the Constrained Probit reduces the model complexity to only twice that of the Benchmark CBSA, but at the cost of large amounts of predictive power.

5.3 Model Comparison: Qualitative Fit (Results 1 and 2)

Figure 5 visually depicts the actual average cooperation rates in the experiment versus the four models' predicted average cooperation rates. Figure 5(a) depicts the average cooperation rates over time in the human trials as found by Camera and Casari. Figures 5(b) and 5(c) depict the average predicted cooperation rates over time in the CBSA models; first the Benchmark, then the Constrained. Figures 5(d) and 5(e) depict the predicted average cooperation rates over time which arise from the Probit models; first the Full, then the Constrained.

Consider the following observations about the human cooperation rates as seen in Figure 5(a). First, Treatment 3 does not overlap with Treatments 1 and 2 and instead lies strictly above it for the entirety of the thirty periods. Second, Treatment 1 involves somewhat more cooperation than Treatment 2 in early periods, until some time between rounds 5 and 10, where Treatment 2 begins to involve more cooperation. Treatment 2 continues to involve more cooperation until the end, with the exception of a brief time around Period 20. Third, average cooperation rates appear to settle into their long-run averages around period fifteen or so. There is a possible upward drift in the three treatments, suggesting if the experiments were to have gone on longer, perhaps average cooperation rates would be u-shaped.[22]

The CBSA predicted cooperation rates, Figures 5(b) and 5(c), match the human data qualitatively quite well. First, Treatment 3 is strictly above the other treatments for the entirety of the run. Second, Treatment 1 has an early lead in cooperation over Treatment 2, but switches places around between periods 5 and 10, as in the human data. (It also ignores the brief reversal around Period 20, which looks like noise.) The crossover occurs a bit too early in the Constrained CBSA. Third, average cooperation rates settle into their long run averages fairly early, although apparently earlier than the human data (more likely by Period 10 to 15). There is also a possible upward drift in all treatments near the end, particularly in Treatment 2. The upward drift in Treatment 2 is perhaps too pronounced in the Constrained CBSA.

[22] In the future we hope to run the Camera and Casari experiment with a higher continuation rate to find whether a u-shape indeed emerges and compare it to CBSA.

The Full Probit, Figure 5(d), matches the human data somewhat well: Treatment 3 is largely above the other two treatments, Treatment 1 and 2 switch as they do in the human data, with overlap seen in the human data. Also, there appears to be some possible upward drift in all Treatments near the end. The most significant way that the Full Probit matches the human data better than the CBSA data is, Treatment 1's spike around period 5. The Full Probit has this spike (and the Constrained Probit has it in an even more exaggerated way). CBSA does not have this spike at all (and, in fact, CBSA seems significantly less noisy than the Probit). In the authors' opinion, this brief spike is likely just noise in the human data and not a meaningful feature to attempt to match. So the extent to which the Full Probit 'overfits' on such features, it would suggest that CBSA is a better qualitative fit. On the other hand, if one believes that the period 5 spike in the human data is not noise, that the Full Probit picked it up is to its favor.

This spike does cause the Constrained Probit to violate the ordering that Treatment 3 has higher cooperation rates than Treatments 1 and 2 over the whole trial. Finally, although the Full Probit matches the trajectory and final average cooperation rates fairly well, under the Constrained Probit, the average cooperation rates for the Probit under Treatments 1 and 2 falls much faster than either the human data or the benchmark CBSA, resulting in a final average level of cooperation much lower than the observed level. Finally, in Treatments 1 and 2, there is definitely no upward drift seen, although possibly in Treatment 3.

5.4 CBSA: Interpretation of Estimated Parameters (Results 3 and 4)

Here we interpret the values of the estimated parameters, p_{recall}, p_{store}, and H, in the best-fitting CBSA models and establish the results two and three: (2) that the aspiration levels are 'reasonable' and (3) the memory parameters are consistent with Pape and Kurtz (2013).[23] In the course of the presentation we investigate some robustness tests.

Result 3: The aspiration level H represents a target payoff level that the agent seeks in the PD Stage game (Table 3). For the Benchmark and Constrained CBSAs, the fitted values for all three treatments are above the (D, D) payoff (10) but below the (C, C) payoff (25). This suggests that agents hope to do better than the stage game Nash equilibrium but are not confident that they will. Aspiration levels can also be interpreted as the satisficing payment required for a subject to stop searching for a better outcome. This leads to an interpretation of the aspiration level as a weighted average of the two symmetric outcomes, (C, C) and (D, D).[24] In the benchmark model, this interpretation implies that the agent only hopes for the mutual co-operation outcome about 5% of the time in Treatment 1, 32% of the time in Treatment 2, and 13% of the time in Treatment 3. (The values are similar in the Constrained CBSA.) The aspiration level is higher in 2 than in 1, and higher in 3 than in 1. This is consistent with the interpretation that agents hope more monitoring will increase cooperation. However, that interpretation is somewhat undercut by the fact that the aspiration level is higher in Treatment 2 than in Treatment 3.

A robustness test of the aspiration level: In the Benchmark CBSA, when the aspiration level H is constrained set to be higher than achievable, 31, the mean squared error increases about sixteen-fold, to 0.030. When it is constrained to be below the worst payoff, 4, the MSE increases about twelve-fold,

[23] We do not interpret p, the problem set, and α, the probability of cooperating when indifferent, because they are set by theory; therefore, their values are not a "finding" and are therefore it is not appropriate to interpret them as one would estimated parameters. Please see Section 4.1 for how those parameters were chosen.

[24] One could also interpret the value as being induced by some 'hope distribution' over all four payoff values, $5, 10, 25,$ and 30.

to 0.023. The fact that reasonable aspiration values provide a better fit than 'unreasonable' aspiration values should be interpreted as a vote in favor of CBSA being the correct model specification for human choice here.

Result 4: The memory probabilities p_{recall} and p_{store} can be interpreted as follows: A low p_{recall} implies the relevant limitation is cognitive capacity, while a low p_{store} implies that the relevant limitation is memory space. These parameters vary across the three treatments in the Benchmark CBSA, and are constrained to be identical in the Constrained CBSA.

In all treatments of the benchmark and in the constrained CBSA, p_{recall} was found to be consistently smaller than p_{store}, which implies cognitive capacity is more of a constraint than memory storage. This pattern was also found when CBSA was mapped to human classification learning data in Pape and Kurtz (2013). This ordering consistency over problem domains should be taken as a vote of confidence in CBDT/CBSA as a model of human choice.[25]

Next, let us consider the ordering across treatments in the Benchmark CBSA. In recall probability p_{recall}, this is the ordering: $p_{\text{recall}}(Treat_1) < p_{\text{recall}}(Treat_2) < p_{\text{recall}}(Treat_3)$. This suggests *as the length of the problem vector (i.e. available information) increases, the probability of recall increases.* This seems counter-intuitive: the longer the problem vector, the more information must be recalled in the future. Presumably more information is more difficult to store and recall than less information, which would make it *harder*, not *easier*, to recall. On the other hand, adding information that is relevant to the problem could have the opposite effect. For example, if one has access to the opponent's ID, it may take less effort to store and recall cases from earlier periods because there is a 'marker' to attach those memories to: this ID variable. In any case, the apparent empirical fact is that the recall probability is inversely related to the length of the problem vector.

Under storage probability p_{store}, there is the ordering is inconsistent with any monotonic function of length of the problem vector: $p_{\text{store}}(Treat_2) < p_{\text{store}}(Treat_1) < p_{\text{store}}(Treat_3)$. It seems that little can be concluded from this. Perhaps it suggests that the Constrained CBSA should be preferred.

Robustness test of imperfect memory: When we only consider perfect memory, i.e. $p_{\text{recall}} = p_{\text{store}} = 1$, we find that the MSE worsens significantly. When memory is constrained to be perfect, the MSE increases about forty-fold, to 0.084. This suggests that imperfect memory may be more important than reasonable aspiration values.

6 Conclusion

In this paper, we use Case-based Decision Theory to explain the average cooperation level over time of the repeated Prisoner's Dilemma among random pairs of people in a small group; the work of Camera and Casari (2009). We are the first to use a method called the Case-based Software Agent and show that CBSA (and therefore CBDT) empirically explains human behavior in a strategic, dynamic, multi-agent setting. This provides a different frame of reference to view the results in the iterated prisoner's dilemma game that incorporates past experiences into explaining sustained cooperative outcomes. We establish four main results.

First, we find that CBSA predicts human behavior in this strategic and dynamic setting quite well. It has a good quantitative fit, and its predicted outcome variable has the main patterns of the human

[25] The best-fitting values of these parameters in Pape and Kurtz (2013) were approximately $p_{\text{recall}} \approx .7$, $p_{\text{store}} \approx 1$. That both memory probabilities are lower here suggests that cases in this experiment are more difficult to remember or process than those in Pape and Kurtz (2013).

data without overfitting. This is a vote in favor of CBSA and CBDT as a explanation of human behavior in this setting.

Second, we find that CBSA compares quite favorably with the Probit from Camera and Casari: it has an arguably stronger qualitative fit, a stronger quantitative fit, and fewer free parameters to achieve this fit. The Constrained CBSA, in which it is assumed that the memory parameters do not vary across treatments, also matches fairly well; albeit somewhat worse quantitative fit than the Full Probit, but an arguably stronger qualitative fit, and even fewer free parameters than the Full CBSA. The Constrained Probit, which is presented to consider a Probit which does not take advantage of individual-level calibration, fairs much worse than either CBSA. Combined with Result 1, a fairly strong empirical argument can be made that CBSA should be considered seriously as an empirical explanation of human behavior in the repeated Prisoner's Dilemma.

Third, we find best-fitting CBSA aspiration value H implies humans aspire to payoff value above the mutual defection outcome but below the mutual cooperation outcome, which suggests they hope, but are not confident, that cooperation can be achieved. In principle, the best-fitting aspiration values could have fallen into the 'unreasonable' range–greater than the best or lower than the worst possible outcome–which would have implied that CBSA is misspecified.

Fourth, and consistent with Pape and Kurtz (2013), we find evidence that suggests that cognitive capacity, and not memory storage space, is the more significant constraint on human problem-solving. The evidence is that the best-fitting parameter values suggest that memory is imperfect in the following way: a relatively high fraction of events are encoded in memory ($78\% - 90\%$), but agents have a lower probability of accessing stored, relevant event memories ($45\% - 70\%$) when they call upon their memory in a utility calculation.

This paper was predicted in the concluding paragraph of Pape and Kurtz (2013): "This computational implementation of Case-based Decision Theory"–that is, CBSA–"can be calibrated to and tested against human data in any existing experiment which can be represented in a game-theoretic form[.] This suggests a model for future studies. As these studies accumulate, we will learn whether and when Case-based Decision Theory provides an adequate explanation of human behavior in other decision settings and may also learn which parameters appear to vary by setting and which, if any, remain constant across settings." With more studies like this one, their stated goal may be achieved: "This could lead to a version of CBDT which can be used to simulate human behavior in a variety of economic models."

References

Andreoni J, Miller JH (1993) Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence. The Economic Journal 103(418):570–585

Axelrod R (1980) Effective choice in the prisoner's dilemma. Journal of Conflict Resolution 24(1):3–25

Ballinger TP, Hudson E, Karkoviata L, Wilcox NT (2011) Saving behavior and cognitive abilities. Experimental Economics 14(3):349–374

Bereby-Meyer Y, Roth AE (2006) The speed of learning in noisy games: partial reinforcement and the sustainability of cooperation. The American Economic Review 96(4):1029–1042

Bleichrodt H, Filko M, Kothiyal A, Wakker PP (2012) Making case-based decision theory directly observable

Bó PD (2005) Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. The American Economic Review 95(5):1591–1604

Bo PD, Fréchette GR (2011) The evolution of cooperation in infinitely repeated games: Experimental evidence. The American Economic Review 101(1):411–429

Brosig J (2002) Identifying cooperative behavior: some experimental results in a prisoner's dilemma game. Journal of Economic Behavior & Organization 47(3):275–290

Camera G, Casari M (2009) Cooperation among strangers under the shadow of the future. The American Economic Review pp 979–1005

Camerer C, Hua Ho T (1999) Experience-weighted attraction learning in normal form games. Econometrica 67(4):827–874

Cappelletti D, Guth W, Ploner M (2011) Being of two minds: Ultimatum offers under cognitive constraints. Journal of Economic Psychology 32(6):940 – 950

Chong JK, Camerer CF, Ho TH (2006) A learning-based model of repeated games with incomplete information. Games and Economic Behavior 55(2):340–371

Cooper R, DeJong DV, Forsythe R, Ross TW (1996) Cooperation without reputation: experimental evidence from prisoner's dilemma games. Games and Economic Behavior 12(2):187–218

Ellison G (1994) Cooperation in the prisoner's dilemma with anonymous random matching. The Review of Economic Studies 61(3):567–588

Ellsberg D (1961) Risk, ambiguity, and the savage axioms. Quarterly Journal of Economics 75(4):643–669

Friedman JW (1971) A non-cooperative equilibrium for supergames. The Review of Economic Studies pp 1–12

Fudenberg D, Maskin E (1986) The folk theorem in repeated games with discounting or with incomplete information. Econometrica: Journal of the Econometric Society pp 533–554

Gayer G, Gilboa I, Lieberman O (2007) Rule-based and case-based reasoning in housing prices. The BE Journal of Theoretical Economics (Advances) 7(1)

Gilboa I, Schmeidler D (1995) Case-Based Decision Theory. The Quarterly Journal of Economics 110(3):605–39

Gilboa I, Schmeidler D (1996) Case-based Optimization. Games and Economic Behavior 15:1–26

Golosnoy V, Okhrin Y (2008) General uncertainty in portfolio selection: A case-based decision approach. Journal of Economic Behavior & Organization 67(3):718–734

Hanoch Y (2002) "neither an angel nor an ant": Emotion as an aid to bounded rationality. Journal of Economic Psychology 23(1):1–25

Hauk E, Nagel R (2001) Choice of partners in multiple two-person prisoner's dilemma games an experimental study. Journal of conflict resolution 45(6):770–793

Kreps DM, Milgrom P, Roberts J, Wilson R (1982) Rational cooperation in the finitely repeated prisoners' dilemma. Journal of Economic theory 27(2):245–252

Luce RD, Raiffa H (1957) Games and Decisions. WIley, New York

Matsui A (2000) Expected utility and case-based reasoning. Mathematical Social Sciences 39(1):1–12

Miller JH (1996) The coevolution of automata in the repeated prisoner's dilemma. Journal of Economic Behavior & Organization 29(1):87–112

Monterosso J (2002) The fragility of cooperation: A false feedback study of a sequential iterated prisoner's dilemma. Journal of Economic Psychology 23(4):437 – 448

von Neumann J, Morgenstern O (1944) Theory of Games and Economic Behavior. Princeton University Press, Princeton, NJ

Nosofsky R (1986) Attention, Similarity, and the Identification–Categorization Relationship. Journal of Experimental Psychology: General 115(1):39–57

Nosofsky R, Gluck M, Palmeri T, McKinley S, Glauthier P (1994) Comparing Models of Rule-Based Classification Learning: a Replication and Extension of Shepard, Hovland, and Jenkins (1961). Memory and Cognition 22:352–352

Ossadnik W, Wilmsmann D, Niemann B (2012) Experimental evidence on case-based decision theory. Theory and Decision pp 1–22, DOI 10.1007/s11238-012-9333-4, URL http://dx.doi.org/10.1007/s11238-012-9333-4

Pape AD, Kurtz KJ (2013) Evaluating case-based decision theory: Predicting empirical patterns of human classification learning. Games and Economic Behavior 82:52–65

Roth AE, Erev I (1995) Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. Games and Economic Behavior 8(1):164–212

Savage LJ (1954) The Foundations of Statistics. Wiley

Selten R, Stoecker R (1986) End behavior in sequences of finite prisoner's dilemma supergames a learning theory approach. Journal of Economic Behavior & Organization 7(1):47–70

Shepard R, Hovland C, Jenkins H (1961) Learning and Memorization of Classifications. Psychological Monographs 75:1–41

Simon HA (1957) Models of man; social and rational.

Simon HA, Egidi M, Marris RL (2008) Economics, bounded rationality and the cognitive revolution. Edward Elgar Publishing

Tesfatsion L (2006) Handbook of Computational Economics, vol 2, Elsevier B.V., chap 16

Tyson CJ (2008) Cognitive constraints, contraction consistency, and the satisficing criterion. Journal of Economic Theory 138(1):51–70

7 Appendix: Details about the Constrained Probit

Table 2: Probit Regression on Individual Choice to Cooperate: Marginal Effects

Dependent variable: 1= cooperation, 0=defection

	Treatment 1 Private monitoring	Treatment 2 Anonymous public monitoring	Treatment 3 Public monitoring (non-anonymous)
Reactive strategies			
Grim trigger	-0.485***	-0.309***	0.046
	0.007	0.093	0.035
lag 1	0.079*	-0.052***	-0.084***
	0.0434	0.003	0.037
lag 2	0.088**	-0.089***	-0.175***
	0.036	0.028	0.044
lag 3	0.066	-0.071***	-0.089***
	0.047	0.019	0.018
lag 4	0.038*	-0.062*	-0.075
	0.023	0.034	0.075
lag 5	0.005	-0.069***	-0.031
	0.003	0.008	0.052
Global strategies			
Grim trigger	-	-0.248	-0.121***
	-	0.171	0.019
lag 1	-	0.194***	0.021
	-	0.065	0.067
lag 2	-	0.184	0.026
	-	0.109	0.046
lag 3	-	0.191***	0.051***
	-	0.059	0.019
lag 4	-	0.153***	0.003
	-	0.032	0.034
lag 5	-	0.131***	-0.027
	-	0.034	0.067
Targeted strategies			
Grim trigger	-	-	-0.408***
	-	-	0.039
lag 1	-	-	-0.062***
	-	-	0.017
lag 2	-	-	-0.079***
	-	-	0.033
lag 3	-	-	-0.034
	-	-	0.024
lag 4	-	-	-0.062***
	-	-	0.007
lag 5	-	-	-0.079***
	-	-	0.009
Sample Size	3,520	5,080	4,280

Similar to the Camera and Casari (2009) paper, this is a Probit regression to identify the marginal effect these different strategies. 2 shows the results of the different strategies from the 'constrained Probit,' which differs from the one in Camera and Casari (2009) only in that individual and cycle fixed effects are excluded. To make clear the construction of the variables included: the grim trigger is coded as 1 for all periods following a defection. The lag variables are to control for the five following periods after a defection. The lag 1 contains a 1 for the first period after a defection by an opponent but zero for all other periods, and a lag 2 contains a 1 for the second period after a defection by an opponent. If the player chooses to defect after observing a defection by an opponent then the expectation would be a negative coefficient on at least one of the grim trigger or lags.