



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Games and Economic Behavior 52 (2005) 373–385

GAMES and
Economic
Behavior

www.elsevier.com/locate/geb

Identifying individual differences: An algorithm with application to Phineas Gage

Daniel Houser^{a,*}, Antoine Bechara^b, Michael Keane^c,
Kevin McCabe^a, Vernon Smith^a

^a *Interdisciplinary Center for Economic Science & Department of Economics, George Mason University,
4400 University Drive, MSN 1B2, Fairfax, VA 22030, USA*

^b *Department of Neurology, University of Iowa, IA, USA*

^c *Department of Economics, Yale University, CT, USA*

Received 18 June 2004

Available online 26 April 2005

Abstract

In many research contexts it is useful to group experimental subjects into behavioral “types.” Usually, this is done by pre-specifying a set of candidate decision-making heuristics and assigning each subject to a heuristic in that set. Such approaches might perform poorly when applied to subjects with prefrontal cortex damage, because it can be hard to know what cognitive heuristics such subjects might use. We suggest that the Houser, Keane and McCabe (HKM) robust classification algorithm can be a useful tool in these cases. An important advantage of this classification approach is that it does not require one to specify either the nature or number of subjects’ heuristics in advance. Rather, both the number and nature of the heuristics are discerned directly from the data. To illustrate the HKM approach, we draw inferences about heuristics used by subjects in the well-known gambling task [Bechara, A., Damasio, A.R., Damasio, H., Anderson, S., 1994. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition* 50, 7–12].

© 2005 Elsevier Inc. All rights reserved.

JEL classification: C11; C15; C91

* Corresponding author.

E-mail address: dhouser@gmu.edu (D. Houser).

1. Introduction

The unhappy circumstances of Phineas Gage are by now well known. Briefly, as related by Antonio Damasio (1994) in his *Descartes Error*, Gage was working as the foreman for a railroad construction team in Vermont in 1848, when an explosion blew an iron bar through his left cheek, skull and the front of his brain. The bar exited the top of his head at high speed, and Gage managed to survive the blast. Although his post-accident IQ remained high according to standard measures, he nevertheless underwent radical personality changes and, perhaps more interestingly, seemed to lose the ability to make good decisions. In particular, he systematically made decisions that were, by any objective measure, not in his long-run best interest. He eventually lost his job and family, and spent much of the rest of his life working as a sideshow attraction for a circus.

We now know that Gage suffered damage to the ventromedial (VM) area of his prefrontal cortex (see, e.g., Damasio et al., 1994). People with damage in this area typically maintain good memory and score well across a wide range of personality and intelligence tests. However, they tend to have difficulty in making “good” decisions. That is, they often make decisions that seem clearly contrary to their best interest, even when they claim that they know this is the case.

Investigating the natures of the differences between VM and normal decision making has proved challenging, because VM patients perform as well as normal patients on many standard diagnostic tests. However, Bechara et al. (1994) describe one laboratory experiment in which VM patients perform remarkably differently than control subjects. This experiment has been dubbed the “gambling task,” because it involves turning over cards sequentially and earning and losing money, according to the markings on each card. Bechara et al. (1994) report that VM patients choose cards from “bad” decks systematically more often than people without such brain damage. In their experiment, a bad deck is one that yields high immediate rewards but higher future losses, so that on average a person playing a bad deck will lose money. A good deck, on the other hand, provides lower immediate rewards but even lower future costs, so that on average a person drawing from the good deck will earn money. The main result reported by Bechara et al. (1994) is that about 60% of VM patients draws are from bad decks, while this is true for only about one-third of their control subjects.

Bechara et al. (2000) investigate three reasons, not mutually exclusive, for differences in behavior in the gambling task. These are that VM patients might be relatively (i) hypersensitive to reward; (ii) insensitive to punishment; or (iii) insensitive to future consequences. To discriminate these hypotheses they designed a new experiment, a variant of the gambling task, such that the bad decks yield low immediate punishment and even lower future earnings, while the good decks yield high immediate punishment and even higher future reward. Analysis of this experiment’s data allows them to conclude that neither (i) nor (ii) is supported by the experimental data, and that (iii) is a simple hypothesis consistent with the evidence.

In this paper we discuss an alternative procedure for drawing inferences about the heuristics used by VM and control patients when playing the original gambling task. Our

approach is to analyze data from the original environment using the statistical classification algorithm suggested by Houser et al. (2004). The goal of our analysis is not to provide new results about the behavior of people with VM damage. Indeed, experimentation over the last decade by Bechara and others has expanded the knowledge of VM behavior far beyond what one can expect to gain by a statistical analysis of a relatively old data set. Rather, in this paper we demonstrate that the Houser, Keane and McCabe (HKM) classification procedure can be used to discern behavioral patterns that were not originally teased out of this data set, and that those patterns line-up well with what subsequent experimentation has already discovered. In particular, we show that hypothesis (ii) above, that VM patients might be relatively insensitive to losses, can be informed through an HKM analysis of the original gambling-task data.

There are several reasons that behavioral researchers in all fields, including economics, psychology and neuroscience, might be interested in the HKM statistical approach. One is that HKM does not require the researcher to pre-specify the nature or number of the heuristics used by subjects. This is in marked contrast to many approaches to type-classification that require the investigator to pre-specify the universe of possible decision rules (e.g., the popular strategy suggested by El-Gamal and Grether (1995)). Especially when analyzing the behavior of people with brain damage, it seems likely that the usual introspective process that generates this universe may fare quite badly.¹ In addition, HKM does not require that all subjects with a particular brain condition (in the present case, VM and control subjects) use the same heuristic. As discussed below, the idea behind the procedure is to group subjects according to similarities in their decision-making behavior, regardless of any known physical abnormalities they might possess.²

The data set analyzed in this paper is small and unbalanced. It consists of seventeen VM patients, and eight lesion control subjects who have brain damage in an area outside of the ventromedial prefrontal cortex (in particular, to the left-somatosensory cortex). Nevertheless, we demonstrate that a simple analysis can be conducted that groups subjects according to similarities in their decision-making strategies (or heuristics), and that allows inference with respect to whether these heuristics differ in terms of their sensitivity to losses.

We allow for two types of heuristics in our population. Our results indicate that fifteen VM patients and two controls use one type of heuristic, while two VM patients and six controls use the other. The two heuristics do not differ with regard to the way they respond to losses, which lines up well with the results of subsequent experimentation reported by Bechara et al. (2000).

¹ This is not to say that introspection necessarily works well when trying to explain the behavior of “normal” subjects: even in very simple environments it can be extremely difficult to write down a model that explains or predicts individual decision making well. Also, while quite common, introspection is not the only procedure available to determine a universe of possible heuristics. Objective evidence from neuroeconomic studies of decision making might provide useful insights into the cognitive strategies used by both brain damaged and normal subjects (see, e.g., McCabe et al., 2001).

² It is possible, of course, to incorporate this information into the HKM classification procedure.

2. Statistical methodology

The statistical procedure used in this paper is developed in detail in Houser et al. (2004), and will be only briefly described here. Papers that discuss closely related procedures for inference in multinomial choice frameworks include Geweke and Keane (1999a), Geweke et al. (2001) and Houser (2003). The HKM approach is useful whenever an investigator is interested in drawing inferences about the nature of behavioral heterogeneity in a population, but does not feel comfortable taking a strong stand with respect to the nature of that heterogeneity. In particular, under relatively weak assumptions, the HKM algorithm draws inferences about both the nature and number of heuristics (or, equivalently, decision rules) used by subjects in a given population.

A decision rule is a map from information to action. For example, if people sitting in a theatre are given the information that the theatre is burning, many will likely decide to act by leaving the building. Behavioral heterogeneity might exist even here: a few might decide to stay. Intuitively, the HKM approach allows one to draw inferences about both the nature and number of relationships that exist between the information people have and the actions they take, at least within a given context.

While many interesting types of decisions are easily observed, it is usually the case that the information that resulted in a particular action is not. This is less the case in laboratory experiments. There, much (even most) of the information that is relevant to subjects' laboratory decisions is under the control of, and therefore known to, the researcher. We exploit this control to specify the form of the heuristics that we investigate below.

2.1. The HKM classification procedure³

In order to provide additional detail about the HKM algorithm, we consider an experiment where subjects solve a T period dynamic decision problem. The “gambling task” analyzed below is an instance of this environment, although the discussion in this section is more general. Suppose that each period subjects choose either alternative A or B , each of which results in a finite monetary reward. Payoffs can be stochastic, but the realizations of the random variables in period t occur before the decision at t is made, while the realizations of period $(t + 1)$'s random variables occur after the decision at t . Each subject's total payoff is the sum of the rewards earned over the T periods. Subjects have complete information regarding the stochastic link between their current choices and future payoffs, but the link is complicated and it is difficult to determine the decision rule that maximises expected total payoffs.

The goal is to learn about the dynamic decision rules that subjects actually use when solving this difficult problem. To do this, Houser et al. (2004) (henceforth, HKM) begin by assuming that subjects are rational in a weak sense. In particular, a subject will choose alternative A in period t if and only if, in period t , the value that they place on choosing A is greater than the value they place on choosing B . Because the problem is dynamic, the value that subjects place on A and B depend both on the immediate reward

³ This section follows Houser (2003) closely.

to each choice and on the way subjects believe that choice will impact their future pay-offs.

The idea is that choices, in general, have both immediate and future costs and benefits, and that even when people have the same information they might use it differently and draw different inferences about the immediate and future consequences of their decisions. This is the fundamental intuition that guides the HKM algorithm. To implement the HKM type classification procedure one posits that individual alternative valuations are additively separable into a contemporaneous component that captures the immediate net benefit of a choice (e.g., the alternative's immediate monetary payoff) and a "future component" that depends on a subject's information and accounts for the subjective way information is used to value alternatives. We use the term "future component" because it is natural for economists to think that this function might map a subject's information to his or her expected net future benefits from taking a particular action. However, this need not be the case. For example, if a subject is "myopic" in the sense that they are focused only on immediate rewards, then the future component would be identically zero. A key advantage of the HKM algorithm is that it does not require one to take a strong stand on the nature of the future component.

In the laboratory it is often reasonable to assume that the contemporaneous payoff structure is known (and is equal to the cash value of a choice), so that differences in choice behaviour between subjects who face otherwise identical contemporaneous returns can be traced to differences between the subjects' future components. The idea put forth by HKM (Houser et al., 2004) is to cluster subjects into groups that seem to have similar future components, while simultaneously drawing inferences about the future components' forms. In this way, HKM avoid taking a strong a-priori stand on either the nature or number of decision rules at use in the population subjects.

HKM (Houser et al., 2004) model the unobserved future component of each alternative's value as a parametric stochastic function of the subject's information set I_{nt} . The information set can include anything the researcher believes is relevant to subjects when making their decisions, such as choice and payoff histories. Then, the value that subject n assigns to alternative $j \in \{A, B\}$ in period t , $V_{njt}(I_{nt})$, given that they use decision rule k , can be written

$$V_{njt}(I_{nt} | k) = w_{njt} + F(I_{n,t+1} | I_{nt}, j, \pi_k, \zeta_{njtk}),$$

$$I_{n,t+1} = H(I_{nt}, j).$$

Here, w_{njt} is the known immediate reward associated with alternative j . $F(\cdot)$ represents the future component. It depends on the alternative j and information set I_{nt} , and is characterised by a finite vector of parameters π_k , whose values determine the nature of decision rule k , and a random variable ζ_{njtk} that accounts for idiosyncratic errors subjects make when attempting to implement decision rule k . (The researcher must specify the distribution of the idiosyncratic errors.) The function $H(\cdot)$ is the information set's stochastic law of motion. It provides the dynamic link between current information and actions and future information, and it is exogenous with respect to the decision rule.

We denote the choice in period t of subject n following decision rule k with information I_{nt} by:

$$d_k(I_{nt}) = \begin{cases} A & \text{if } Z_{nt}(I_{nt} | k) > 0, \\ B & \text{otherwise,} \end{cases} \quad \forall k \in K,$$

where $Z_{nt}(I_{nt} | k) = V_{nAt}(I_{nt} | k) - V_{nBt}(I_{nt} | k)$.

The goal is to draw inferences about the parameters π_k ($\forall k \in K$), and about the probability with which each subject uses each decision rule. To do this HKM construct the likelihood function associated with this framework. This requires knowing the probability, conditional on a subject's information set, that they will choose A or B .

The probability that subject n using decision rule k chooses alternative A at period t , given that they have information I_{nt} is given by

$$P(d_k(I_{nt}) = A) = P(V_{nAt}(I_{nt}) > V_{nBt}(I_{nt})) = P(w_{nAt} - w_{nBt} + f(I_{nt} | \pi_k) > 0),$$

where $f(\cdot)$ is a stochastic function that represents the differenced future components $F(I_{n,t+1} | I_{nt}, A, \pi_k, \zeta_{nAtk}) - F(I_{n,t+1} | I_{nt}, B, \pi_k, \zeta_{nBtk})$. The conditional probability that B is chosen is one minus the conditional probability that A is chosen.

With conditional choice probabilities in hand it is straightforward to construct the likelihood function needed to draw inferences about the different decision rules at use in the population, and the probability with which each subject uses each rule. Under the distributional assumptions made by HKM, the likelihood function corresponds to a mixture of normals probit model. This likelihood can be computationally burdensome to maximise, and numerical procedures such as Gibbs sampling are typically required. Interested researchers should consult Houser et al. (2004) for discussion on this point.

3. The gambling task

Bechara's gambling task (Bechara et al., 1994) is a sequence of static decision problems under ambiguity. The experimenter begins by giving a subject \$2,000 in play money. The experimenter places four decks of cards in front of the subject, and tells him/her that they can earn more play money by turning over cards, and that his/her goal is to earn as much play money as possible. The subject is told that every card they choose will result in them earning some amount of money, and that there will be occasional cards that impose costs on them. The subject is told nothing else. The subject then begins turning over cards, one-by-one, until they are told to stop by the experimenter. The stopping point is after 100 cards have been selected, although the subject does not know this in advance.

The subject is told nothing about the payoff or cost distributions within any of the decks of cards. In fact, the decks have been constructed in a very particular way. The first two decks, call them A and B , provide a positive payment of \$100 for each card. However, they also have occasional very high costs. On average, turning over 10 cards in the A or B decks will have a net cost of \$250. The C and D decks have lower rewards per card, \$50, but also have lower occasional costs. On average, turning over 10 cards in the C or D decks yields a positive return of \$250. For this reason, we will refer to decks A and B as the "bad" decks, and C and D as the "good" decks.

The main result reported by Bechara et al. (1994) (see also Bechara et al., 1997) is that VM patients choose from the bad decks statistically significantly more often than normal subjects. On average, around 60% of all VM patients' draws are from the bad decks, while

this is true of only about one-third of the normal patients' draws. This led to much speculation about the source of the behavioral difference. One question was whether VM patients were relatively insensitive to losses, and if this insensitivity could explain the difference. Subsequent research by Bechara et al. (2000), which used a new experiment designed to address this question, suggested that differences in loss aversion-behavior were not likely the source of the different choices. The results we report below provide convergent evidence for this conclusion.

4. A simple model

The Houser et al. (2004) approach to type classification requires that subjects' relevant information sets, and the link between information and action, be specified. We assume that each subject has a subjective value associated with draws from each deck of cards, and that they draw a card from the deck on which they place the highest value. The way that values are formed can be modeled in any way that the researcher chooses, subject only to identification issues. In this paper, because our intent is only illustrative, we use a simple model that nevertheless allows us to address whether VM subjects respond to losses differently than lesion controls. We noted in the introduction that this was one of the primary hypotheses advanced to explain the behavioral patterns observed in the original gambling task data. We also noted that this hypothesis was not supported by results from subsequent experiments.

Denote the deck by j (with total number of decks J), the subject by n and the current draw by t . Assume that subjects assign values to draws in H different ways (that is, there are H valuation heuristics used in the population). With this notation, we model the subjective value that subject n assigns to drawing a card from deck j at round t , assuming they use heuristic h , as:

$$V_n(j, t; h) = b_{1jh} + b_{2jh}I(\text{last draw was from deck } j \text{ \& } t > 50) \cdot \text{Loss}(t - 1) \\ + b_{3jh}I(\text{last draw was from deck } j \text{ \& } t > 50) \cdot \text{Reward}(t - 1) \\ + e_n(j, t; h),$$

where e is an identically and independently distributed Gaussian random variable that represents idiosyncratic noise, which arises due to failures to implement the heuristic perfectly. Because this is a situation of ambiguity, the model assumes that the subject uses the first 50 draws to gain experience in each deck. Inferences with respect to loss and reward effects are based on the final 50 draws experienced by each subject. Finally, the function $I(\cdot)$ represents an indicator function that takes value one if the condition inside the brackets is true, and is otherwise zero. This model simply posits that the value a subject places on drawing from deck j depends on a constant, noise, and his/her most immediate previous experience with that deck.

It is possible to use the HKM algorithm to draw inferences about the number H of heuristics in the population, the nature of each heuristic h in H (that is, the coefficient values), and to determine the probability with which each subject uses each heuristic. A specific way to do this is detailed in Houser et al. (2004), and involves a Bayesian

analysis of a mixture of probits model (for more on mixtures of probits see, e.g., Geweke and Keane, 1999b).

For this paper's purposes, however, we assume that there are exactly two heuristics at use in the population. There are two reasons for this decision. First, the results of substantial previous research with this population suggest that there are in fact two types of behavioral heuristics in this population, and using these previous results to inform our current model is reasonable. At the same time, note that there is no necessary reason to expect that all VM patients will follow the same heuristic, or that all normal controls will follow the same heuristic. For example, some VM patients might follow a strategy that looks very similar to the control subjects. The HKM procedure allows for this and other possibilities.

A second reason to assume that there are two types of decision rules in this population is that, as a practical matter, it would be difficult to interpret the finding that there are three or more heuristics in the population. The reason is that our sample size is rather small (8 controls and 17 VM patients), and evidence of more than two heuristics might not be robust to a larger sample, or the nature of the heuristics that we estimate might be a quite biased reflection of the true heuristics at use in the population, given the relatively small number of subjects that would be assigned to each.

This highlights an important feature of the HKM approach to type classification. Because it is a robust approach, in the sense that both the nature and number of heuristics are determined endogenously, it can be less efficient than procedures that take a stand on the heuristics subjects use. Of course, if such a stand is wrong, and the model consequently misspecified, then the efficiency gain might come at the cost of specification error bias.

4.1. Implementation and identification

Although there are two good decks, and two bad decks in the actual experiment, in this section we report results based on a model that treats each pair as one. Equivalently, we model the individual as making a choice between choosing a deck with \$100 payoffs or \$50 payoffs, and then randomizing across the two decks within that choice. Hence, we set $J = 2$, which turns out to mean that there are three identified coefficients, along with one variance term with a pegged value, that characterize each heuristic.

Note that the value function described above requires both location and scale normalization for identification. Location normalization is achieved by differencing:

$$\begin{aligned}
 &V_n(1, t; h) - V_n(2, t; h) \\
 &= b_{11h} - b_{12h} \\
 &\quad + b_{21h} I(\text{last draw was from deck 1} \ \& \ t > 50) \cdot \text{Loss}(t - 1) \\
 &\quad - b_{22h} I(\text{last draw was from deck 2} \ \& \ t > 50) \cdot \text{Loss}(t - 1) \\
 &\quad + b_{31h} I(\text{last draw was from deck 1} \ \& \ t > 50) \cdot \text{Reward}(t - 1) \\
 &\quad - b_{32h} I(\text{last draw was from deck 2} \ \& \ t > 50) \cdot \text{Reward}(t - 1) \\
 &\quad + e_n(1, t; h) - e_n(2, t; h).
 \end{aligned}$$

The differenced constants are not separately identified, but are estimated as a single constant. Similarly, the differenced error component is treated as a single noise term. Also, because the nature of the experiment induces little variation in rewards, the coefficients on lagged rewards are only weakly identified. Consequently, we choose to drop them for the remainder of our analysis. Finally, scale normalization is achieved by pegging the variance of the error at a fixed value.

To implement the Bayesian version of HKM as described in Houser et al. (2004) one must specify priors on the coefficients b that appear in the value expression above, along with the fraction of each type that exists in the population. We follow Houser et al. (2004) and use Gaussian priors with means of zero and standard deviations of one for the intercepts, and 0.1 for the coefficients on losses. We use a diffuse Dirichlet prior centered at $1/2$ for the fraction of each type in the population.

5. Data and results

Our data set consists of 25 subjects who played the gambling task one time. 17 of our subjects are VM patients, and 8 are lesion controls with damage to the left somatosensory cortex. The data were collected by Antoine Bechara and colleagues at the University of Iowa, and represent a subset of data that has been previously published in various books and journals. Figure 1 compares the frequency with which the two types of patients draw from the bad decks (the \$100 decks). As has been previously reported, VM patients draw from the bad decks statistically significantly more often than the normal patients. Moreover, as seen in Fig. 2, the rate at which VM patients draw from the bad deck seems roughly constant over the entire experiment. The rate at which LC's draw from the bad deck is similar to the VM rate over the first 20 or so draws, but then declines substantially, but stays roughly constant over the last 70 or so draws.

5.1. Results

Our results are derived through the use of a Gibbs sampling algorithm. Details about the Gibbs sampler, and the way in which it can be implemented to draw inferences in the




Fig. 1. Mean number of bad deck choices by frontals and controls.

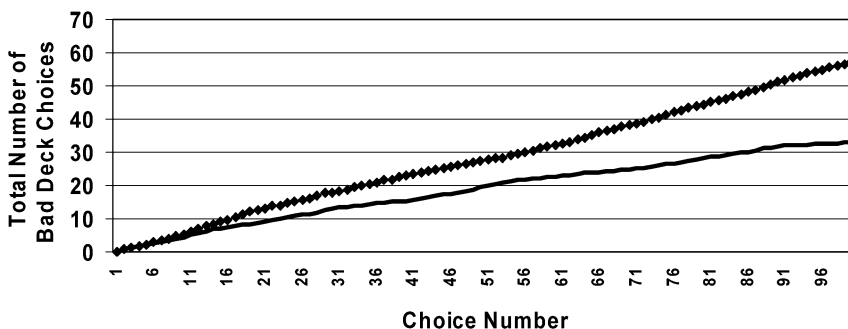


Fig. 2. Cumulative number of bad deck choices by frontals (+) and controls (—).

present environment, are presented in Houser et al. (2004) and will not be repeated here. Briefly, the Gibbs sampler is a recently developed numerical procedure for drawing inferences within the context of statistical models like ours. The Gibbs sampler is attractive because it provides accurate inferences under weak regularity conditions, and those conditions are satisfied by our model. Our results are based on a Gibbs sampling algorithm that we coded in FORTRAN 77 and that makes extensive use of IMSL subroutines. We ran the sampler for a total of 500 cycles. The results reported below are based on the last 250 cycles.⁴

Consider first the way in which we type-classify subjects. We based our subject classification on the posterior mean probability that they were each type. If the posterior probability of being the VM type is greater than or equal to 0.50,⁵ then they are classified as that type. Otherwise, they are classified as the lesion control type. The posterior type probabilities favored one type over another by only a few percentage points for most of our subjects. The highest posterior mean probability across all subjects of being the VM type was about 66%, and the smallest was about 36%.

Table 1 provides the results of our typing procedure. We call one of the two estimated heuristics the “VM” heuristic, simply because most of the subjects assigned to it are VM patients. We denote the other heuristic as the “LC” heuristic for the same reason. It turns out that 17 subjects are classified as VM types, which is identical to their frequency in the data. However, two subjects classified as VM are in fact lesion controls. Thus, four subjects are “misclassified,” in the sense that their actual brain condition is not reflected by the label of the heuristic that they use.

Table 2 describes the marginal posterior distributions for the coefficients of each heuristic. Notice first that the marginal posterior distributions of the coefficients for the amount lost in the previous period have the majority of their mass to one side of zero for both

⁴ Visual inspection of the draw sequences suggested that convergence had been achieved by cycle 250. Complete draw sequences are available from the authors on request.

⁵ Three subjects had a posterior probability of 0.496 of being the ventromedial type, and in each case we assigned them to the ventromedial type. Two of these three are actually VM, so reversing the classification for these three reduces the number of total VM types by 3 (from 17 to 14), and induces one additional type classification error (from 4 to 5).

Table 1
Number of subjects of each classified type by actual brain condition

Classification	Actual brain condition		
	VM	LC	Total
VM	15	2	17
LC	2	6	8
Total	17	8	

Table 2
Marginal posterior distributions

	LC heuristic		VM heuristic	
	Mean	SD	Mean	SD
Constant	-0.05362	0.09747	0.04513	0.08600
Loss bad deck	0.00062	0.00028	0.00066	0.00033
Loss good deck	-0.00381	0.00138	-0.00400	0.00125

Note. These coefficients correspond to the differenced value function described in Section 4.1, where the bad deck is deck 1 and the good deck is deck 2. Hence, the constant is $b_{11} - b_{12}$, loss bad deck is b_{21} , and loss good deck is $-b_{22}$.

the LC and VM heuristics. Moreover, the values of these coefficients are very similar. This suggests immediately that, as reported by Bechara et al. (2000) based on a different experimental design, both VM and lesion control patients respond to losses incurred in the previous period, and that these responses are similar. On the other hand, the posterior means of the constant terms for the two heuristics differs by about 0.1, and the posterior means lie on different sides of zero. Given the small sample size, this provides some evidence that the baseline rate at which VM's and LC's choose from the bad deck differs.⁶ In particular, the coefficient estimates imply that VM's choose from the bad deck at a baseline rate of about 55%, while lesion controls choose from the bad deck at about a 45% baseline rate.

In addition to the baseline rates, our estimates imply that both types of subjects are more likely to choose from the same type of deck after experiencing a loss in that deck, than they would be otherwise.⁷ For example, a subject who turns over a card in the bad deck and receives a cost of \$1,250 will choose from one of the bad decks again with probability 0.82 if they are using the VM heuristic, and with probability 0.77 under the LC heuristic. Experiencing a cost of \$250 from the good decks generates probabilities of bad deck choices of 16 and 15% for the VM and LC heuristics, respectively. Overall then, these findings suggest that VM patients choose cards from the bad decks at a higher baseline rate

⁶ Note that the standard deviations of the intercepts are large relative to their means: zero lies in a high marginal posterior density region for each.

⁷ This might be counterintuitive. One possible explanation for this behavior is that subjects come to expect that it is unlikely to experience two losses in a row in a given deck. Alternatively, this result might reflect an aggregation effect embedded in our statistical model. In particular, it is possible that subjects are in fact switching decks after a loss, but not switching reward amounts.

than the LC subjects. Consequently, they experience large losses more frequently, and these losses lead to yet more frequent choices from the bad decks.

Our analysis leaves unanswered the question of why VM patients would tend to choose from the \$100 decks at a higher baseline rate than the lesion controls. Further experimentation by Bechara et al. (2000) suggests that the reason may be that VM damage leaves one unable to assess the future negative consequences of one's actions accurately. This is consistent with the patterns observed in Fig. 2, which suggests LC patients reduce their baseline choice rates in the bad deck after the first 20 draws or so, while this reduction is not apparent in the VM subjects' data.

6. Concluding comments

Uncovering the nature and number of behavioral heuristics that people use, even in very narrow contexts, presents one of the most important current challenges to the behavioral sciences (see Houser (2003) for an elaboration of this point). A standard approach to this involves somehow determining a universe of possible ways that people might act, and then determining which one among this universe fits each person's behavior best (see, e.g., El-Gamal and Grether, 1995). While this approach has been shown to work well in some circumstances (see, e.g., Houser and Winter, 2004), there are some environments in which its success is less likely. The study of brain damaged people is one such environment, because it does not seem likely that introspection by a person with a normally functioning brain could provide accurate guidance on the heuristics that might be used by someone with a brain abnormality. The HKM classification procedure is a robust alternative. We have demonstrated in this paper that the results obtained by application of the HKM statistical procedure to data from the original gambling-task design line-up well with results from subsequent new experiments with VM patients.

Although this paper focused on a behavioral study, it is important to point out that the HKM algorithm has broad applicability that extends beyond the analysis of behavioral data. In particular, it would be straightforward to apply the algorithm to data that include behavioral decisions and neuronal firing information, say as might be collected during an fMRI imaging experiment. Analyzing such a data set holds the promise of identifying jointly both the behaviors that neuroeconomists should seek to explain, along with the neural structures that support those behaviors.

Acknowledgments

We thank an anonymous referee for valuable comments. We also received useful feedback from participants in seminars at the University of Arizona, George Mason University, University of Mannheim, University of Minnesota (at the 2002 Conference on Economics and Neuroscience) and the 2001 AIER Symposium on Behavioral Economics.

References

- Bechara, A., Tranel, D., Damasio, H., 2000. Characterization of the decision-making deficit of patients with ventromedial prefrontal cortex lesions. *Brain* 123, 2189–2202.
- Bechara, A., Damasio, A.R., Damasio, H., Anderson, S., 1994. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition* 50, 7–12.
- Bechara, A., Damasio, H., Tranel, D., Damasio, A.R., 1997. Deciding advantageously before knowing the advantageous strategy. *Science* 275 (5304), 1293–1295.
- Damasio, A.R., 1994. *Descartes Error*. G.P. Putnam's Sons, New York.
- Damasio, H., Grabowski, T., Frank, R., Galaburda, A.M., Damasio, A.R., 1994. The return of Phineas Gage: The skull of a famous patient yields clues about the brain. *Science* 264, 1102–1105.
- El-Gamal, M., Grether, D., 1995. Are people Bayesian? Uncovering behavioral strategies. *J. Amer. Statist. Assoc.* 90, 1137–1145.
- Geweke, J., Keane, M., 1999a. Bayesian inference for dynamic discrete choice models without the need for dynamic programming. In: Mariano, R., Schuermann, T., Weeks, M. (Eds.), *Simulation Based Inference and Econometrics: Methods and Applications*. Cambridge University Press. Also available as Federal Reserve Bank of Minneapolis working paper #564, January, 1996.
- Geweke, J., Keane, M., 1999b. Mixture of normals probit models. In: Hsiao, Lahiri, Lee, Pesaran (Eds.), *Analysis of Panels and Limited Dependent Variable Models*. Cambridge University Press, pp. 49–78. Also available as Federal Reserve Bank of Minneapolis staff report #237.
- Geweke, J., Houser, D., Keane, M., 2001. Simulation based inference for dynamic multinomial choice models. In: Baltagi, B. (Ed.), *Companion to Theoretical Econometrics*. Blackwell, pp. 466–493.
- Houser, D., 2003. Classification of types of dynamic decision makers. In: Nadel, L. (Ed.), *Encyclopedia of Cognitive Science*, vol. 1. Nature Publishing Group, London, pp. 1020–1026.
- Houser, D., Winter, J., 2004. How do behavioral assumptions affect structural inference? Evidence from a laboratory experiment. *J. Bus. Econ. Statist.* 22 (1), 64–79.
- Houser, D., Keane, M., McCabe, K., 2004. Behavior in a dynamic decision problem: An analysis of experimental evidence using a Bayesian type classification algorithm. *Econometrica* 72 (3), 781–822.
- McCabe, K., Houser, D., Ryan, L., Smith, V., Trouard, T., 2001. A functional imaging study of cooperation in two-person reciprocal exchange. *Proc. Natl. Acad. Sci.* 98 (20), 11832–11835.