Dynamic Decision Makers, Classification of Types of

Intermediate article

Daniel Houser, University of Arizona, Tucson, Arizona, USA

CONTENTS

A dynamic decision problem is one that requires a sequence of decisions in a setting where pay-offs and alternatives available for later decisions depend on earlier choices. Analysis of data from dynamic behavioral experiments can shed light on the nature of the different types of dynamic decision makers in the population.

INTRODUCTION

Types

Suppose a general orders his troops into battle against an approaching enemy and valiantly leads the charge himself. One of his soldiers takes advantage of the momentary confusion surrounding the order by slipping away from the battlefield to an area of relative safety. One reasonable explanation for the different decisions made by the general and soldier is straightforward: they are different 'types' just in that they face different incentives. Victory in battle means lasting glory and honor for the general, while the soldier might receive little but the chance to fight again another day. It is less straightforward, however, to explain why one soldier flees while his comrades, who are in the same situation, rush forward with the general into combat. As a casual description, we might also say that 'observationally' identical soldiers who make different decisions are different 'types'.

In order to decide whether to flee or fight, soldiers must solve a dynamic decision problem. It is dynamic because their decision affects in a nontrivial way the alternatives available to them at later times. For example, the soldier who fights might be in a position to save the life of a wounded comrade on the battlefield. The soldier who flees might save his own life, but at the risk of being punished for desertion. Their eventual decision rests on idiosyncratic characteristics such as preferences and subjective assessments of battlefield risks.

Game theory defines people who have different preferences as different 'types'. Unfortunately, preferences are not observable. In practice, it is more useful to define people as different types of dynamic decision makers if, like the soldiers, they make different decisions in observationally identical dynamic situations.

Interest in classifying and characterizing dynamic decision makers has grown as the importance of accounting for type heterogeneity in dynamic economic models has become apparent. This importance stems from the fact that most of dynamic economics has as its final goal policy analysis. That is, the goal is to predict how different sorts of incentive structures (e.g. the tax system) affect dynamic decisions (e.g. work and educational choices). As the example of the soldiers makes clear, not everybody responds to the same incentives in the same way. Economists increasingly recognize that analyses, which assume that firms and societies can be described as collectives and modeled as though they were single agents, can often lead to very misleading conclusions and policy recommendations (Furubotn and Richter, 2000). Models that take account of type heterogeneity have the potential to improve policy analysis substantially.

Decision Rules

Economists use so-called 'decision rules' to describe the way a person's actions depend on personal information. 'Information' here should be thought of broadly as everything a person knows (including demographic variables) that could be relevant to a decision. The soldiers above, for example, can be thought of as having the choice either to fight or flee. When presented with the information that the enemy is approaching, the decision rules of some soldiers generate the decision to flee, while those of others generated the decision to fight. In general, people in observationally identical situations who make different decisions are viewed as using different decision rules, and a person's 'type' is defined by the decision rule they use. Increasingly, research is directed towards identifying, at least within narrow contexts, the number and nature of different decision rules, or types, that exist in the population.

Differences in decision rules can arise from differences in preferences. This is a good explanation in many situations, particularly when trying to explain idiosyncratic differences in tastes for, say, coffee and tea. On the other hand, differences along dimensions such as propensities to cooperate, which are well documented and involve higher-order cognitive processing, seem less naturally attributable to preferences. There is some evidence that differences in decision rules associated with higher-order functions are due to different cognitive algorithms employed to determine actions (e.g. McCabe *et al*., 2001). Such differences are analogous to the difference between human and computer decision making. When a human plays chess against a computer, both parties have the same objective and information, yet their decision rules differ because they use different algorithms to determine their moves.

Expectations

A higher-order task of particular importance to dynamic decisions is expectation formation, because all dynamic decision problems require some sort of forward-looking behavior. Economists have been using the 'rational expectations' assumption to model forward-looking behavior for decades. However, numerous studies in economics and psychology suggest that expectations are not formed rationally. Moreover, it is straightforward to show that different expectation formation mechanisms lead to different dynamic decision rules.

An important, and often-replicated, finding in the literature on static decision making is that, except in very simple cases, people do not assess objective probabilities correctly (e.g. Camerer, 1995). Since probability assessment is cognitively difficult, it is presumably accomplished with idiosyncratic heuristics. Moreover, because probability assessment is fundamental to expectation formation, it seems likely that expectations are formed with idiosyncratic and imperfect heuristics. Although research in this area is still in its early stages, it seems plausible that heterogeneity in expectation formation underlies much of the idiosyncratic variation in dynamic decision rules.

TYPE ELICITATION

Stopping Experiments

Dynamic decision problems (DDPs) faced by individuals provide perhaps the simplest interesting environment in which to study dynamic decision rule heterogeneity. In these environments an individual makes several decisions sequentially, and the decisions made early in the problem affect the nature of the decision task later in the problem. For example, a person might first decide whether to bicycle or walk to work, and then decide on the route to follow. This is a DDP, because the set of candidate routes depends on the outcome of the first decision. (Economists contrast DDPs with sequential 'static' decision problems, where one makes a series of unrelated decisions.) It is important to understand how people actually solve DDPs, particularly when the dynamic nature of the problem involves deciding between different pay-offs at different times (so-called 'intertemporal' decision problems), because many actual consumption, savings, and labor supply decisions must be made within this context.

'Stopping problems', a widely studied class of DDPs, have proved useful tools in classifying and characterizing dynamic decision rules. In a simple stopping experiment, subjects receive payment offers sequentially from the experimenter until they accept one, at which time the experiment ends. Many variants of this basic design have been studied. For example, subjects might have to pay for offers; they might not know the distribution from which offers are generated; they might be able to accept previously rejected offers; and they might not know the exact amount of the offer, but only whether it is higher or lower than other offers. An advantage of this framework is that theoretical predictions about behavior under various decision rules are straightforward to derive. Different decision rules often imply different stopping points. Hence, observing stopping times allows one to draw simple and compelling inferences about the sorts of decision rules that are used in the population.

Analysis of stopping experiment data, using techniques such as that of El-Gamal and Grether (1995) discussed below, show that there is great heterogeneity in the ways subjects solve experimental stopping problems. There is little evidence to suggest that people decide in ways that are consistent with rational expectations. Instead, subjects seem to use sophisticated, nonstationary reservation pay-off heuristics. This means that subjects stop as soon as their pay-off is sufficiently high, where 'sufficiently' depends, for example, on the number of times they have had to search and on whether they are paying search costs.

There is a small set of reservation pay-off decision rules into which most subjects' behavior seems to fall. Moreover, there are two features that most of these rules share. Firstly, subjects who use them tend to stop searching somewhat earlier than a rational expectations searcher would. Secondly, these heuristics work well in the sense that subjects who use them earn only slightly less on average (often about one percent) than they would have if they had followed the rational expectations rule. Since the rational expectations rule is cognitively very complex to implement, there may be a sense in which using reservation pay-off heuristics is in fact optimal. For a survey of results from the experimental stopping experiment literature, see Cox and Oaxaca (1996).

The Voluntary Contribution Mechanism

Types can also be discerned in game environments where multiple subjects interact and make strategic decisions that affect each other's pay-off. The voluntary contribution mechanism (VCM) is an important example of such a game. There are N players, and player n has endowment w_n. Player n contributes g_n to the public good and leaves the remainder in a private account. The total contribution to the public good is $G = \sum_n g_n$. The interesting feature of the VCM is that the return on investment in the public account differs from that on investment in the private account. Without loss of generality we can suppose that the return to each player on the total investment in the public account is given by r while the return on the private account is set to unity. This means that the pay-off function for player n is

$$\Pi(g_n, \hat{g}_n) = (w_n - g_n) + rG \tag{1}$$

where $\hat{g}_n$ represents the vector of contributions of everyone except player n. Provided that $r < 1$ it is easy to see that, given any arrangement of contributions by the other subjects, each player maximizes his or her individual pay-off by 'free-riding', or contributing zero to the public good. Hence, free-riding is a dominant Nash equilibrium strategy. But if $rN > 1$ then it is Pareto optimal for each player to contribute everything to the public good, and this strategy Pareto-dominates free-riding. The parameter r is the marginal per-capital return (MPCR). When designing VCM experiments, the MPCR and the number of subjects are usually chosen to exploit the tension between free-riding and Pareto optimality.

Experimental research with the VCM, has generated many widely-replicated results, including clear evidence of decision rule heterogeneity (for a survey see Ledyard, 1995). In particular, there is usually a subset of subjects, 'free-riders', whose decision is to contribute very little to the public good in every round, and another subset, 'cooperators', who systematically contribute a large fraction of their endowment to the public good. A third subset uses 'reciprocal' rules, trying to match others' contributions.

The presence of reciprocal decision rules suggests that group dynamics might be influenced by the type composition of groups. For example, the presence of players who contribute little or nothing to the public good could lead to decreasing aggregate contributions over time if reciprocators attempt to match free-riders' small contributions. Hence, a feedback system that is sensitive to the proportion of each type within the group could be created, and could affect the extent to which a group is able to sustain cooperation.

Recent research has found that type composition seems to have important effects on group dynamics (e.g. Gunthorsdottir *et al.*, 2001). In particular, the number of free-riders in a group influences that group's path over the course of a game. Without free-riders, groups are capable of sustaining high levels of contribution to the public good; while the presence of free-riders often pushes groups towards successively lower levels of contributions.

TYPES AND PERSONALITY SURVEYS

The ability to learn about a person's behavioral type from a personality survey would be useful, since the dependence of group outcomes on type composition implies that knowledge of types could be used to design groups (such as school classes) efficiently. However, whether behavioral types broadly and systematically correlate with personality surveys is an open question, and experiments have generated widely conflicting results. Nevertheless, some personality variables seem to correlate with propensities to cooperate in experiments. Personality dimensions displaying this correlation include Machiavellianism, self-monitoring, and three of the 'big five' personality traits.

Machiavellianism

Inspired by the writings of Niccolo Machiavelli (1469–1527), and first developed by Christie and Geis (1970), the Machiavellianism (or Mach) scale measures the extent to which a person agrees that the end sanctifies the means. People who score highly on the Mach scale tend to be manipulative, opportunistic, and rational. Low Machs tend to be more emotional and less likely to depart from social norms in order to pursue their own self-interest. While high Machs tend to be competitive and exploitative, low Machs are usually more willing to cooperate (Gunnthorsdottir, McCabe and Smith, 2002).

Self-monitoring

The 'self-monitoring' scale is an measure of the dependence of an individual's behavior on the social context. High self-monitors work to create the impression needed to obtain their social goals, while low self-monitors are less concerned about the impression they make. High self-monitors have been found to be more likely to cooperate, particularly in experiments where repeated interactions with the same counterpart are possible.

The Big Five

The 'big five' personality traits are extraversion, agreeableness, conscientiousness, neuroticism, and openness. Among these, extraversion and agreeableness seem to be positively correlated with cooperativeness, while neuroticism seems to be negatively correlated. The relation of the other two traits with cooperativeness is not clear.

Many other personality variables, including self-esteem and locus of control, have been studied in relation to cooperation, but without clear results. For further discussion on the connection between types and personality surveys, see Kurzban and Houser (2002).

STATISTICAL METHODS FOR TYPE CLASSIFICATION

Sophisticated statistical procedures are not usually required to determine whether subjects in an experiment behave according to a particular decision rule. Intuitively, all that is required is to compare actual decisions with those that would arise under a hypothesized behavior. Although the details depend on the experimental design, formal procedures to accomplish this sort of comparison are typically straightforward. A more difficult task, and one that typically requires sophisticated statistics, is to determine how people actually make decisions in a given dynamic environment.

Attempting to characterize actual decision making requires, at least, allowing for multiple decision rules to be used in the population. The task is then to determine the number of decision rules, and to assign each subject to a decision rule. Broadly, there are two ways in which this can be done. We will briefly summarize the two approaches, and then discuss in greater detail an instance of each of them.

One approach, exemplified by a procedure suggested by El-Gamal and Grether (1995), requires one to specify in advance a set of candidate decision rules. A statistical procedure is then used to choose a 'best' subset of these rules. Finally, each subject is assigned to one of the rules in the subset. An advantage of this sort of procedure is that it is relatively straightforward to implement. However, unless it is feasible to include all of the rules that subjects might possibly use in the prespecified superset, a potential drawback is that the right rules might not be included. Misspecification could mask underlying commonalities in subjects' play.

An alternative approach, exemplified by a method suggested by Houser *et al*. (2001), requires no assumptions about the number of decision rules used in the population, the nature of each decision rule, or the assignment of subjects to decision rules. This approach requires cluster analysis: subjects are clustered according to commonalities in their behavior.

The goal of both these approaches is to put each subject into a behavioral category. El-Gamal and Grether require one to specify the categories in advance, while Houser *et al*. allow the categories to be determined by the data. Of course, it may not be easy to assign behavioral labels to groups that follow statistically similar decision rules.

The Classification Procedure of El-Gamal and Grether

Suppose one has data from a behavioral laboratory experiment where each of N subjects makes T decisions. Let C^K denote the prespecified set of K heuristics (i.e., decision rules) that subjects might use to make these decisions, and let $c \in C^K$ denote a particular heuristic. The idea is to determine, for each subject, the number of decisions consistent with each possible heuristic, and then assign that subject to the heuristic that best fits his or her behavior.

To implement the procedure one assumes that each subject follows exactly one of the heuristics in C^K. In practise, of course, a subject's behavior may not be perfectly consistent with any of the heuristics in C^K. El-Gamal and Grether (1995) circumvent this problem by assuming that subjects follow their heuristics with error.

Heuristics are chosen that specify a subject's decision uniquely from his or her state. A subject's 'state', which generally changes after each decision, is a vector that summarizes all of the person's decision-relevant information. Let x_t^c be an indicator variable that takes the value one if the subject's tth decision agrees with heuristic c and takes the value zero otherwise. Assume that the decisions are made independently with common error rate ε.

Let $(x_{n1}, \ldots, x_{nT})$ be a vector denoting subject n's actions, and let $(x_{n1}^c, \ldots, x_{nT}^c)$ be a vector of zeros and ones that summarizes the consistency of the subject's choices with c. That is, assume that $x_{nj}^c = 1$ if decision rule c predicts decision x_{nj}, and $x_{nj}^c = 0$ otherwise. Then set $X_n^c = \sum_t x_{nt}^c$. The likelihood function for the subjects' actions is then:

$$f^c(x_{n1}, \ldots, x_{nT}) = (1 - \varepsilon/2)^{X_n^c} (\varepsilon/2)^{T - X_n^c} \tag{2}$$

It is natural to assign each subject to the heuristic from the candidate set that maximizes his or her likelihood function.

This model can be 'overfit': including a large number of heuristics in C^K would allow the statistical model to fit a sample arbitrarily well. Overfitting usually leads to results with little external validity. To avoid overfitting, El-Gamal and Grether suggest penalizing the likelihood for each additional heuristic that is included in the set of candidate heuristics. Let C^k denote a subset of $k \leq K$ decision rules. El-Gamal and Grether argue that a reasonable penalized log-likelihood is obtained by forming the Bayesian posterior that arises under the following priors: (1) the probability that the population includes exactly k heuristics is $1/2^k$; (2) all possible k-tuples of heuristics in any C^k are equally likely (each with probability $1/K^k$); (3) all allocations of heuristics to subjects are equally likely (each with probability $1/k^N$); (4) all error rates (between zero and one) are equally likely; and do not depend on the number of rules used in the population or on the way those rules are assigned. This generates the following penalized log-likelihood function:

$$\log\left(\prod_n \max_{c_n \in C^k} f^{c_n}(x_{n1}, \ldots, x_{nT})\right) - k \log 2 - N \log k - k \log K \tag{3}$$

Determination of the population of heuristics as well as the assignment of subjects to heuristics is accomplished by simply maximizing the above expression over the set of all possible k-tuples that can be formed from the set of K decision rules.

The Classification Procedure of Houser, Keane, and McCabe

Suppose that subjects solve a 15-period DDP. At each period, subjects choose either A or B, each of which results in a nonnegative monetary reward. Pay-offs are stochastic. The realizations of the random variables for period t occur before the decision at t is made, and the realizations of the random variables for period $t+1$ occur after the decision at t. Each subject's total pay-off is the sum of the rewards earned over the 15 periods. Subjects have complete information regarding the stochastic link between their current choices and future pay-offs, but the link is complicated and it is difficult to determine the decision rule that maximizes expected total pay-offs.

The goal is to learn about the dynamic decision rules that subjects actually use when solving this difficult problem. Houser *et al*. (2001) 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 the person places on choosing A is greater than the value he or she places on choosing B. Because the problem is dynamic, the values that subjects place on A and B depend both on the immediate reward to each choice and on the way subjects believe that choice would influence their future pay-offs. Houser *et al*. assume that alternative valuations are additively separable into a 'present' component, which captures immediate rewards (in this case the immediate monetary pay-off), and a 'future' component, which captures any benefits expected to accrue in subsequent periods as a result of that choice (in this case future monetary pay-offs).

Since the present pay-off structure is known for each agent, differences in behavior result only from differences in the future component. Hence, all differences in decision rules between subjects can be captured by differences in the future component. Houser *et al*. propose clustering subjects into groups that seem to have similar future components, while simultaneously drawing inferences about the future components' forms. In this way, they avoid the need to prespecify the nature of the decision rules used by the subjects.

Drawing on Geweke and Keane (1999), Houser *et al*. model the unobserved future component of each alternative's value as a parametric function of the subject's information set I_{nt}. The information set can include anything the researcher believes is relevant to the subject when making his or her decision, such as choice and pay-off histories. Then, the value that subject n assigns to alternative $j \in \{A, B\}$ in period t, V_{njt} (I_{nt}), assuming that the person uses decision rule k, can be written

$$V_{njt}(I_{nt}|k) = w_{njt} + F(I_{n,t+1}|I_{nt}, j, \pi_k, \varsigma_{njtk}) \quad (4)$$

$$I_{n,t+1} = H(I_{nt}, j) \quad (5)$$

Here, w_{njt} is the known immediate pay-off associated with alternative j. $F(\cdot)$ represents the future component. It depends on the alternative j and information set I_{nt} and is characterized by a finite vector of parameters π_k, whose values determine the nature of decision rule k, and a random variable ς_{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 (possibly stochastic) law of motion. It provides the dynamic link between current information, actions and future information. Note that it does not vary with the decision rule.

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

$$d_k(I_{nt}) = \begin{cases} A & \text{if } Z_{nt}(I_{nt}|k) > 0 \\ B & \text{otherwise} \end{cases} \quad \text{for all } k \in K, \quad (6)$$

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 $\pi_k (k \in K)$, and about the probability with which each subject uses each decision rule. To this end, Houser *et al*. construct the likelihood function associated with this framework. This requires knowing the probability, conditional on a subject's information set, that he or she will choose A or B.

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

$$\begin{aligned} P(d_k(I_{nt}) = A) &= P(V_{nAt}(I_{nt}) > V_{nBt}(I_{nt})) \\ &= P(w_{nAt} - w_{nBt} + f(I_{nt}|\pi_k) > 0) \end{aligned} \quad (7)$$

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

Knowing the conditional choice probabilities, it is straightforward to construct the likelihood function needed to draw inferences about the different decision rules used in the population, and the probability with which each subject uses each rule. Under the distributional assumptions made by Houser *et al*. the likelihood function corresponds to a mixture of normals probit model. Unfortunately, this likelihood can be computationally difficult to maximize. Further discussion of this point (and estimation strategies) can be found in Houser *et al*. (2001) and Geweke and Keane (1999).

CONCLUSION

Economists say that people who make different decisions in observationally identical situations are different 'types'. Decision rules form the link between a person's situation and decisions, and it is natural to define a person's type by the decision rule he or she uses. Investigating the nature of the various decision rules at use in the population is important, because the effects of incentives on behavior depend on the decision rules that incentives act upon.

Many economists are particularly interested in the decision rules people use in dynamic environments. Experimental studies have found that a small number of decision rules seem to explain most observed behavior in very narrow contexts, and that these rules do not usually include the rational expectations rule. Further research is needed to determine the nature and number of decision rules in the population, the relationship between decision rules used in different contexts, and consequent implications for individual and group outcomes and incentive structures. Such research may employ sophisticated statistical procedures that group people according to common behavioral patterns. These patterns may be either specified in advance or discerned directly from experimental data.

References

Camerer C (1995) Individual decision making. In: Kagel J and Roth A (eds) *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.

Christie R and Geis F (1970) *Studies in Machiavellianism*. New York, NY: Academic Press.

Cox JC and Oaxaca R (1996) Testing job search models: the laboratory approach. *Research in Labour Economics* **15**: 171–207. Greenwich, CT and London: JAI Press.

El-Gamal M and Grether D (1995) Are people Bayesian? Uncovering behavioral strategies. *Journal of the American Statistical Association* **90**: 1137–1145.

Furubotn EG and Richter R (2000) *Institutions and Economic Theory: The Contribution of the New Institutional Economics*. Ann Arbor, MI: University of Michigan Press.

Geweke J and Keane M (1999) Bayesian inference for dynamic discrete choice models without the need for dynamic programming. In: Mariano, Schuermann and Weeks (eds) *Simulation Based Inference and Econometrics: Methods and Applications*. Cambridge, UK: Cambridge University Press.

Gunthorsdottir A, Houser D, McCabe K and Ameden H (2001) Disposition, history and contributions in public goods experiments. [Working paper, University of Arizona.]

Gunnthorsdottir A, McCabe K and Smith V (2002) Using the Machiavellianism instrument to predict trustworthiness in a bargaining game. *Journal of Economic Psychology* **23**: 49–66.

Houser D, Keane M and McCabe K (2001) How do people actually solve dynamic decision problems? [Working paper, University of Arizona.]

Kurzban R and Houser D (2002) Individual differences in cooperation in a circular public goods game. *European Journal of Personality* **15**: 37–52.

Ledyard J (1995) Public goods: a survey of experimental research. In: Kagel J and Roth A (eds) *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.

McCabe K, Houser D, Ryan L, Smith V and Trouard T (2001) A functional imaging study of cooperation in two-person reciprocal exchange. *Proceedings of the National Academy of Science* **98**: 11832–11835.

Further Reading

Andreoni J (1995) Cooperation in public goods experiments: kindness or confusion? *American Economic Review* **85**: 891–904.

Geweke J, Houser D and Keane M (2001) Simulation based inference for dynamic multinomial choice models. In: Baltagi B (ed.) *Companion to Theoretical Econometrics*. Oxford, UK: Blackwell.

Geweke J and Keane M (1997) Mixture of normals probit models. [Federal Reserve Bank of Minneapolis staff report 237.]

Geweke J and Keane M (2001) Computationally intensive methods for integration in econometrics. In: Heckman J and Leamer E (eds) *Handbook of Econometrics*, vol. V. North Holland.

Gilks WR, Richardson S and Spiegelhalter DJ (1996) *Markov Chain Monte Carlo in Practice*. London, UK: Chapman & Hall.

Krusell P and Smith AA (1995) Rules of thumb in macroeconomic equilibrium: a quantitative analysis. *Journal of Economic Dynamics and Control* **20**: 527–558.

Lettau M and Uhlig H (1999) Rules of thumb versus dynamic programming. *American Economic Review* **89**: 148–174.

McLachlan GJ and Basford KE (1988) *Mixture Models: Inference and Applications to Clustering*. New York, NY: Marcel Dekker.

Sargent TJ (1987) *Dynamic Macroeconomic Theory*. Cambridge, MA: Harvard University Press.

Stokey NL and Lucas RE (1989) *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press.