

Faculty Research



University
of Michigan
Business
School

WORKING PAPER SERIES

Supporting Learning in Evolving Dynamic Environments

Faison P. Gibson
University of Michigan Business School

Working Paper 03-011

Supporting learning in evolving dynamic environments

Faison P. Gibson
University of Michigan Business School
701 Tappan Street
Ann Arbor, MI 48109-1234
fpgibson@umich.edu

August 22, 2002

Abstract

In dynamic decision environments such as direct sales, customer support, and electronically mediated bargaining, decision makers execute sequences of interdependent decisions under time pressure. Past decision support systems have focused on substituting for decision makers' cognitive deficits by relieving them of the need to explicitly account for sequential dependencies. However, these systems themselves are fragile to change and, further, do not enhance decision makers' own adaptive capacities. This study presents an alternative strategy that defines information systems requirements in terms of enhancing decision makers' adaptation. In so doing, the study introduces a simulation model of how decision makers learn patterns of sequential dependency. When a system was used to manage workflows in a way predicted by the model to enhance learning, decision makers in a bargaining experiment learned underlying patterns of sequential dependency that helped them adapt to new situations. This result is rare if not unique in the study of dynamic decision environments. It indicates that a shift, *away from substituting for short-term deficits and toward enhancing pattern learning*, can substantially improve the effectiveness of decision support in dynamic environments. Based on the specific findings in this study, this shift has important implications for designing information system workflows and potential future applications in interface design.

1 Introduction

The growth of electronic and telephone-based communications has increased the prevalence of interactive dynamic decision tasks such as customer support, direct sales, and electronically mediated bargaining. A distinguishing feature of these tasks is that decision makers must execute sequences of interdependent decisions as the environment evolves. For instance, over the course of a day, a telephone-based credit collector may face many classes of delinquent debtors with whom she must bargain. The pitches required to convince these debtors to resolve their delinquencies differ. Getting debtors to settle difficult delinquencies with many interdependencies typically requires lengthy offer, counter-offer sequences while simpler delinquencies may require only one offer.

The data indicate that inexperienced decision makers in dynamic environments ignore these sequential dependencies, leading to poor performance (for reviews, see Sterman, 1994, 2000). One solution is to replace decision makers with software agents optimized for specific tasks (e.g., supply chain management, Kimbrough, Wu, & Zhong, 2002). However, when responses to particular sequences of decisions or the relevant sequences themselves change frequently, these agents become hard if not impossible to maintain, suggesting a continued, important role for human decision makers (Beam, Segev, Bichler, & Krishnan, 1999).

We might be able to design systems to help inexperienced decision makers in dynamic environments with sequential dependencies if we better understood the root of their difficulties. One hypothesis (the deficits hypothesis) is that decision makers in dynamic environments have trouble with sequential dependencies because explicitly accounting for them exceeds their cognitive capacity (Sterman, 1989a, 2000; Sweeney & Sterman, 2000). An effective decision support strategy consistent with the deficits hypothesis is to have the information system handle the computations needed to account for the dependency (e.g., Davis & Kottemann, 1995; Sengupta & Abdel-Hamid, 1993). However, systems based on this strategy suffer from the same brittleness in evolving environments as narrowly defined software agents.

Perhaps more importantly, the deficits hypothesis appears not to reflect skilled performance. Skilled decision makers in dynamic environments with sequential dependencies make decisions by

recognizing features of past decisions in the current one and then recalling a decision that worked well; an activity requiring comparatively little cognitive effort and well-suited to the time constraints imposed by functioning dynamic environments (Klein, Orasanu, Calderwood, & Zsombok, 1993).

However this observation alone provides no guidance for how to design systems to support and reinforce inexperienced decision makers' recognition and recall abilities. As a first step toward enabling such support, this paper develops a novel hypothesis about how decision makers learn to recognize patterns of sequential dependency in the form of a sequence-learning (S-L) simulation model. The S-L model significantly extends recent work in sequence recognition (Cleeremans & McClelland, 1991; Rohde & Plaut, 1999) and cognitive game theory (Fudenberg & Levine, 1998). It focuses on how decision makers learn to recognize sequences of cause and effect in highly interactive environments such as direct sales, customer support, and bargaining. As such, it represents a fundamental departure away from understanding decision makers' support requirements purely in terms of deficits and toward reinforcing capabilities that they bring to the task. As examined in this study, this shift has important implications for designing information system workflows and possible future application in interface design. The resulting impact on decision makers' performance in an *evolving* dynamic bargaining environment examined here is extremely rare if not unprecedented.

The next section reviews prior data on decision maker performance in dynamic decision environments and develops the S-L model. After that, an interpersonal bargaining experiment is conducted to test the S-L model's recommendations for how to design system workflows. Finally, limitations of this study are addressed, and the study's implications for system design are further elaborated.

2 Learning in evolving dynamic environments

By definition, dynamic environments evolve (Brehmer, 1995). Within a given decision episode, decisions early in a sequence may impact the effectiveness of later decisions (Brehmer, 1995; Sterman, 1989a, 1989b). Further, between episodes, the length and interrelationships of dependencies between decisions may themselves change (Diehl & Sterman, 1995).

Generally, decision makers in dynamic environments experience two *short term* problems that

appear traceable to their limited cognitive capacity. First, decision makers have difficulty integrating information into an appropriate cognitive framework that accounts for the sequential dependency (Serman, 1989a, 2000; Sweeney & Serman, 2000). Second, even if decision makers integrate the sequential dependency into an appropriate cognitive framework, they have difficulty explicitly predicting its effects (Serman, 2000).

Supporting these observations, in one recent experiment only 44% of subjects were able to indicate the existence of a production lag in their formulation of an inventory management problem, even though the lag was clearly identified in the problem description (Sweeney & Serman, 2000, Manufacturing Case). In this same group of subjects, only 10% were subsequently able to properly account for the lag when plotting inventory's future trajectory relative to orders. Similar patterns of findings have been reported for more complex inventory management tasks (Davis & Kottemann, 1995; Diehl & Serman, 1995), supply chain management (Serman, 1989b), capital stock management (Serman, 1989a), software project management (Sengupta & Abdel-Hamid, 1993), and firefighting (Brehmer, 1995).

Decision support strategies that reduce demands on decision makers' cognitive capacity have been effective in improving decision makers' short-term performance in these tasks. When subjects acting as software project managers were supplied with a system that clearly forecast the impacts of their staffing decisions, they appropriately adjusted staffing earlier in the sequence (Sengupta & Abdel-Hamid, 1993, pp. 421–422). A similar intervention has been shown effective in inventory management (Davis & Kottemann, 1995). However, designing such systems requires extensive knowledge of the task environment that may not always be available or changes so frequently that it is hard to capture in a maintainable system.

Further, *the performance of experienced decision makers in functioning dynamic environments does not correlate with traditional measures of their cognitive capacity, leading to some doubt that this is the problem to address* (Kanfer & Ackerman, 1989). Experienced decision makers engage in a cognitively less demanding cycle of observing the current situation, recalling a similar past situation, and choosing from a small set of actions that worked before in that situation; repeating

this process as necessary (Klein et al., 1993). Early learners appear much less skilled in recalling appropriate past situations (Sweeney & Sterman, 2000). This difference between early learners and skilled performers has been observed in air traffic control (Kanfer & Ackerman, 1989), emergency dispatch (Joslyn & Hunt, 1998) and firefighting (Klein et al., 1993).

Although there is evidence that experienced decision makers in functioning dynamic environments rely primarily on recall in making decisions, how novices can develop this ability has not been well specified. Few, if any, attempts to develop information systems that support this process have been documented. The next section develops the S-L model of how decision makers in dynamic environments learn to recognize sequential dependencies. This model appears unique: (1) in addressing how decision makers learn to recognize sequential dependencies with experience in dynamic environments; and (2) in proposing ways to improve this learning that can be incorporated into information system design.

3 The S-L Model

Figure 1 shows the S-L model. It treats learning sequential dependencies as a process of learning to associate patterns of events over time. In so doing, it makes use of a simple analogy with the brain, common in cognitive neuroscience, in which its inputs, outputs, and internal representations are conceived as patterns of activation across neurons (McClelland, 2001). The decision maker experiences and reacts to the patterns of activation but is not assumed to explicitly perform the calculations that simulate them in the model.

A high-level interpretation of the decision processes captured by the S-L model is as follows. First, the decision maker observes cues in the environment ($stimulus_k$) and simultaneously combines this perception with previously noted higher level features ($feat(j, t - 1)_k$), leading to an internal representation that may be composed of a number of higher level features ($feat_j$). In the case of the credit collector, examples of the higher level features formed through this process might be that the customer has said no twice already to offers to fully resolve the debt; or that the customer has agreed to a partial resolution but has said no to an offer to fully resolve the delinquency and now

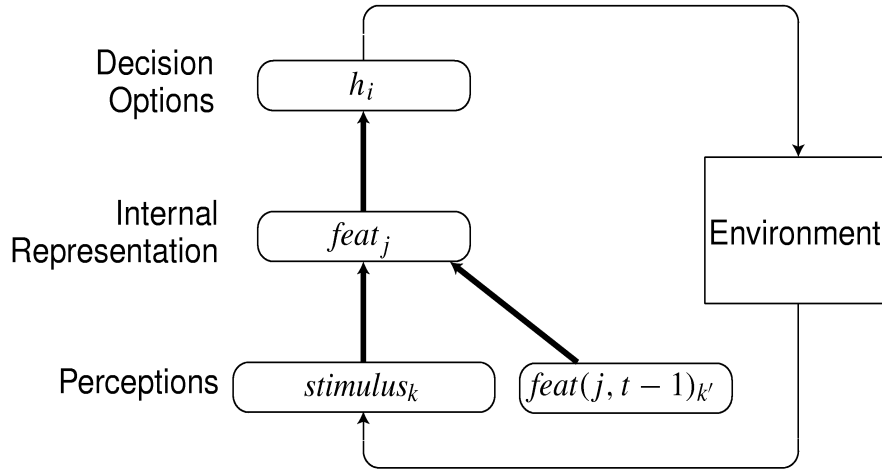


Figure 1: The S-L model. Details of the model's functioning are contained in the text.

appears to be backing off the partial resolution. The decision maker's internal representation then feeds into his or her consideration of which decision options (h_i) to favor, with different patterns of higher level features tending to lend more and less support to different decision options (e.g., ask for more money or offer easier payment terms). An option is then chosen in competition with other options based on its relative support. As is common in cognitive game theory, the model assumes that decision makers learn based on the success and failure of their decisions, with positive reinforcement for successful decisions raising the conditional likelihood that they will be chosen in the future (Fudenberg & Levine, 1998).

3.1 Mechanics of the S-L Model

Making Decisions. Turning to the mechanics of the S-L model, as is common in studies of repeated decision making, it assumes that decision makers' task is to choose one of a competing set of options at each decision point (Dienes & Fahey, 1995; Logan, 1988; Roth & Erev, 1995). The likelihood, act_{h_i} , that a given decision option, h_i , will be chosen from among m such options is based on the decision maker's relative assessment of how likely it is to lead to a desired outcome using the following equation:

$$act_{h_i} = \frac{e^{weighted_sup_i}}{\sum_{l=1}^m e^{weighted_sup_l}} \quad (1)$$

where $weighted_sup_i$ represents the weighted support an option is currently receiving from the

decision maker’s internal representation. Choice over options is then competitive, with the winner determined by random selection over act_{hi} , a frequent assumption in modeling decision maker’s performance in sequential prediction tasks (Cleeremans & McClelland, 1991). Under this mechanism, when the decision maker has more than one option that has worked in the past, it is as if he or she were wavering. When one option has consistently succeeded and others not, this mechanism leads to more certain decisions.

$weighted_sup_i$ in Equation (1) is calculated as the weighted, $w_{i,j}$, sum of the activation of higher level features, act_{feat_j} , that form the decision maker’s internal representation of environmental and other stimuli:

$$weighted_sup_i = \sum_{j=1}^n input_j w_{i,j} \quad \text{where} \quad input_j = act_{feat_j} \quad (2)$$

The activation of a given higher level feature indicates whether a set of stimuli in the decision context that it represents have successfully predicted outcome of given decision options in the past (Rumelhart, Durbin, Golden, & Chauvin, 1995). For instance, the activation of a higher level feature representing that a customer has twice refused offers depends on how well it has performed in predicting the success of different decision options. In the work reported here, activation of features is a “stretched” binomial version of Equation (1) frequently used to model human performance in classification tasks (Rumelhart et al., 1995):

$$act_{feat_j} = \frac{2}{1 + e^{weighted_sup_j}} - 1 \quad (3)$$

where $weighted_sup_j$ is calculated as in Equation (2) but using stimuli, $stimulus_k$, in place of act_{feat_j} , and weights, $w_{j,\{k,k'\}}$, connecting stimuli with higher level features in place of $w_{i,j}$.

Learning. In conformance with the well-known “law of effect”, decision makers learn by adjusting their evaluations of decision options based on success or failure (Dienes & Fahey, 1995; Fudenberg & Levine, 1998; Roth & Erev, 1995). Since the evaluation of decision options is fully determined by the weights connecting stimuli to higher level features and higher level features to decision options,

learning is accomplished by modifying these weights as in a number of studies that model human learning (Rumelhart et al., 1995). First, the weights between the decision options and the higher level features are modified using the following rule:

$$w_{i,j} = \eta (\delta_i) (input_j) + w_{i,j} \quad (4)$$

where η is a learning parameter that can be set to determine the size of the weight change. In this rule, δ_i is the difference between the outcome for a decision option, t_i , and the model's estimated probability that that decision option would succeed:

$$\delta_i = t_i - act(h_i) \quad (5)$$

where the outcome, t_i , is as follows:

$$t_i = \begin{cases} 1 & , \text{ if } h_i \text{ chosen and success} \\ 0 & , \text{ if } h_i \text{ chosen and no success} \\ 0 & , \text{ if } h_{j \neq i} \text{ chosen and success} \\ \frac{e^{weighted_sup_i}}{\sum_{\ell \neq j} e^{weighted_sup_\ell}} & , \text{ if } h_{j \neq i} \text{ chosen and no success} \end{cases} \quad (6)$$

This assumption for t_i is consistent with hill climbing because it indicates that decision makers gravitate toward what has worked and away from what has not based purely on current feedback (Sterman, 2000). The δ_j for adjusting the weights between the feature units and the stimuli is calculated as the weighted sum of the δ_i for the m decision options, h_i , multiplied by the derivative of act_{feat_j} (Rumelhart et al., 1995):

$$\delta_j = \sum_{i=1}^m \delta_i 2 act_{feat_j} (1 - act_{feat_j}) w_{i,j} \quad (7)$$

This δ_j is then used to change the weights between the higher level features and the stimuli, including

the past values of act_{feat_j} , by substituting the relevant values into Equation (4).

3.2 Implications for Decision Maker Learning

As just described, the central assumption in the S-L model is that decision makers in dynamic environments represent situations by combining their current perceptions with their representations of immediately prior situations. This representation process is essentially one of recall. The weighted support internal representations of the current situation receive is based on their past performance in predicting the success of decision options. The decision makers' representations then recall decision options based on these options' past utility in situations receiving that representation. Choice among decision options is then competitive based on each option's level of support. As such, the S-L model conforms to and refines the observation that experienced decision makers in functioning environments make decisions essentially through recall. By construction, the model does not account for any effort decision makers may expend in lengthy reasoning since they are assumed not to have the time in evolving environments.

The internal representation step in the S-L model is an important theoretical refinement. It implies that managing the flow of sequential dependencies the decision maker experiences can facilitate learning. Since internal representations are partially based on previous representations, starting with short sequential dependencies may be advantageous (Elman, 1990). Representations for these shorter dependencies can be learned more quickly and thereby provide more coherent inputs to representations for longer dependencies. However, implementing this progression by necessity reduces time spent learning the longer dependency. It may therefore be more beneficial to start immediately with the longer dependency so that learning time with it is increased. Which of these implications dominates for any given dynamic environment must be derived empirically through simulation. Simulation results then serve as behavioral predictions for human decision makers.

As detailed more fully in the experimental results section below, simulation with the S-L model in a dynamic decision environment with sequential dependencies suggests that time on task with the longest dependencies is critical to performance. Models that spent the longest time in tasks

with longer sequential dependencies were able to learn higher level features of these dependencies. In later interactions, models were able to recognize these higher level features in novel situations and make better decisions than models that had progressed from shorter to longer dependencies.

4 Experiment

This section first describes the method used for subjects, a multi-day Internet bargaining task. It then presents the S-L model's predictions and compares them with actual decision makers' learning in two conditions that test the effect of ordering the progression of sequential dependencies on learning: (1) *Evolving* — Decision makers bargained with opponents whose behavior evolved to include longer and longer sequential dependencies; (2) *Consistent* — Decision makers bargained with opponents who consistently behaved according to the longest dependencies observed in the Evolving condition.

4.1 Subjects Method

Thirty-six subjects participated in a four session bargaining experiment over consecutive days. During the experiment, subjects' bargaining opponents either accepted or rejected the offers they made. The two experimental conditions, Evolving and Consistent, differed in the response rules opponents used during the first three sessions of the experiment. Table 1 summarizes the progression of response rules by condition. In the Evolving condition, opponents' response rules progressed from having no dependency on the previous pattern of offers and responses in the first session (0-offer rule), to depending on the last offer-response pair in the second session (1-offer rule), and finally to depending on the sequence of the last two offer-response pairs in the third session (2-offer rule). As described in more detail below, the 0-offer and 1-offer rule were derived from this 2-offer rule.

In the first three sessions of the Consistent condition, subjects dealt only with opponents using the 2-offer rule with no progression. In the fourth session, subjects in both the Evolving and Consistent conditions dealt with opponents who responded differently (2-offer-prime rule) according to the same sequential dependencies that conditioned the 2-offer rule.

	Evolving ($n = 18$)	Consistent ($n = 18$)	Control ($n = 24$)
Session 1	0-offer	2-offer	NA
Session 2	1-offer	2-offer	NA
Session 3	2-offer	2-offer	NA
Session 4	2-offer-prime	2-offer-prime	2-offer-prime

Table 1: Response rule progression by experimental condition. Control subjects provided a comparison for Evolving and Consistent subjects in Session 4.

For the first three sessions, this design led to three within-subjects measures of performance crossed by the two between-subjects response rule conditions. In the fourth session, subjects performance was measured in four consecutive intervals leading to a four-within by two-between design.

To help determine whether response rule condition and experience influenced subject performance in the fourth session, 24 control subjects were selected from the same subject pool as Evolving and Consistent subjects and run in a one-day experiment where they bargained solely with opponents who responded according to the 2-offer-prime rule.

4.1.1 Detailed Procedure

Subjects were recruited from among students at the University of Michigan and paid \$15 for each session they attended. There were no financial incentives for performance, a fact that was not expected to affect the patterns of mean performance between conditions (Camerer & Hogarth, 1999).

All sessions of the experiment took place using an Internet-based bargaining environment. Subjects came to a computer lab and used a web-browser to connect to a server that supposedly allowed them to interact with other subjects. Similar interactive web interfaces are becoming increasingly common in tasks requiring interpersonal interactions such as bargaining, customer service, and sales.

At the start of the experiment subjects were informed that they would be playing the role of debt collectors against other “debtor” subjects who were behind on their credit card payments. These other subjects were in fact computer algorithms that generated the debtor’s response according

to the response rules outlined above and further described below. During each session, subjects made contact with 20 *different* debtors whom they bargained with for twelve speaker turns each. Subjects' goal on each speaker turn was to get the debtor to agree to pay as much money in as short a time as the debtor would likely agree to. As observed in the functioning environment on which this task is based, even if debtors agreed to terms, subjects had to continue bargaining (Gibson & Fichman, 2002; Sutton, 1991). Debtors could grow noncommittal and reject terms already accepted (the reason they were delinquent in the first place). Alternatively, they could evolve toward accepting more demanding terms that would more quickly resolve the delinquency. Similar patterns of evolving sequential dependency, widely observed in interpersonal negotiation, are an important challenge for designing systems to support decision makers in these environments (Beam et al., 1999).

4.1.2 Bargaining Task Interface and Debtor Behaviors

Figure 2 shows the bargaining task interface. On each speaker turn, subjects had four seconds to make their offer by clicking one of three options: (L)ow (\$100 in 8 days); (M)edium (\$300 in 5 days); or (H)igh (\$900 in 2 days), and then clicking on a talk button. The short time frame and the categorical decision making are typical of the functioning task environment as well as a broader range of functioning environments such as police dispatch (Joslyn & Hunt, 1998), air traffic control (Kanfer & Ackerman, 1989), and firefighting (Klein et al., 1993). After that, subjects both saw on screen and heard through headphones the debtor's response of accept or reject to the offer.

General debtor behaviors. The debtor state transition diagrams (STDs) displayed in Figure 3 implemented the 0-offer, 1-offer, 2-offer, and 2-offer-prime response rules. In this task, nodes in the STDs represented the debtor's state resulting from the subject's most recent offer (arcs). The response heard by the subject was generated according to the node's label of *accept* or *reject* 90% of the time and the opposite response the other 10% of the time.

For each node in each debtor STD in Figure 3, there is one offer that leads to the highest payment with the highest certainty. This move can be referred to as the "best" move for maintaining agreement from the subject's perspective and corresponds to the goal supplied to subjects in their

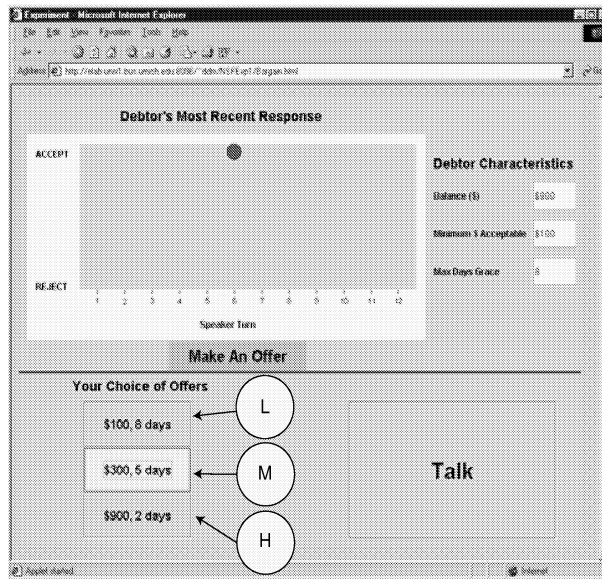


Figure 2: A single debtor contact. The figure shows a hypothetical subject's information display as she is selecting her seventh offer. The labels L(ow), M(edium), and H(igh) are superimposed on the figure for ease of exposition and were not displayed to subjects.

instructions.¹ The absolute number of best moves indicated subjects' overall mastery of the system, perfect mastery indicated by all moves being best moves. Positive increase in best moves was an indication that subjects were becoming more proficient.

2-offer and 2-offer-prime Debtors. We focus first on the 2-offer and 2-offer-prime debtors because they are related and because all the other debtor STDs were derived from the 2-offer debtor's STD. For debtors using the 2-offer and 2-offer-prime STDs, the best offer could be determined with a high degree of certainty from the last *two* offer-response pairs.

The dependencies linking these offer-response pairs were based on foot-in-the-door (FID), door-in-the-face (DIF), and good-cop-bad-cop sequential behavior patterns observed in many bargaining environments (Cialdini, 1984; Gibson & Fichman, 2002; Rafaeli & Sutton, 1991; Sutton, 1991). In FID, bargainers get their opponents to accede to a relatively low request (e.g., L in Figure 3) and then move them up to higher levels (e.g., M or H in Figure 3) that they could not have obtained on one request alone. In DIF, bargainers make a burdensome request (e.g., H) that will almost certainly be rejected and then get their opponents to accept a request (e.g., M) that they would not

¹At the very start of a contact, this "best" move involved making an offer that would lead to the possibility of getting an offer accepted on the next speaker turn.

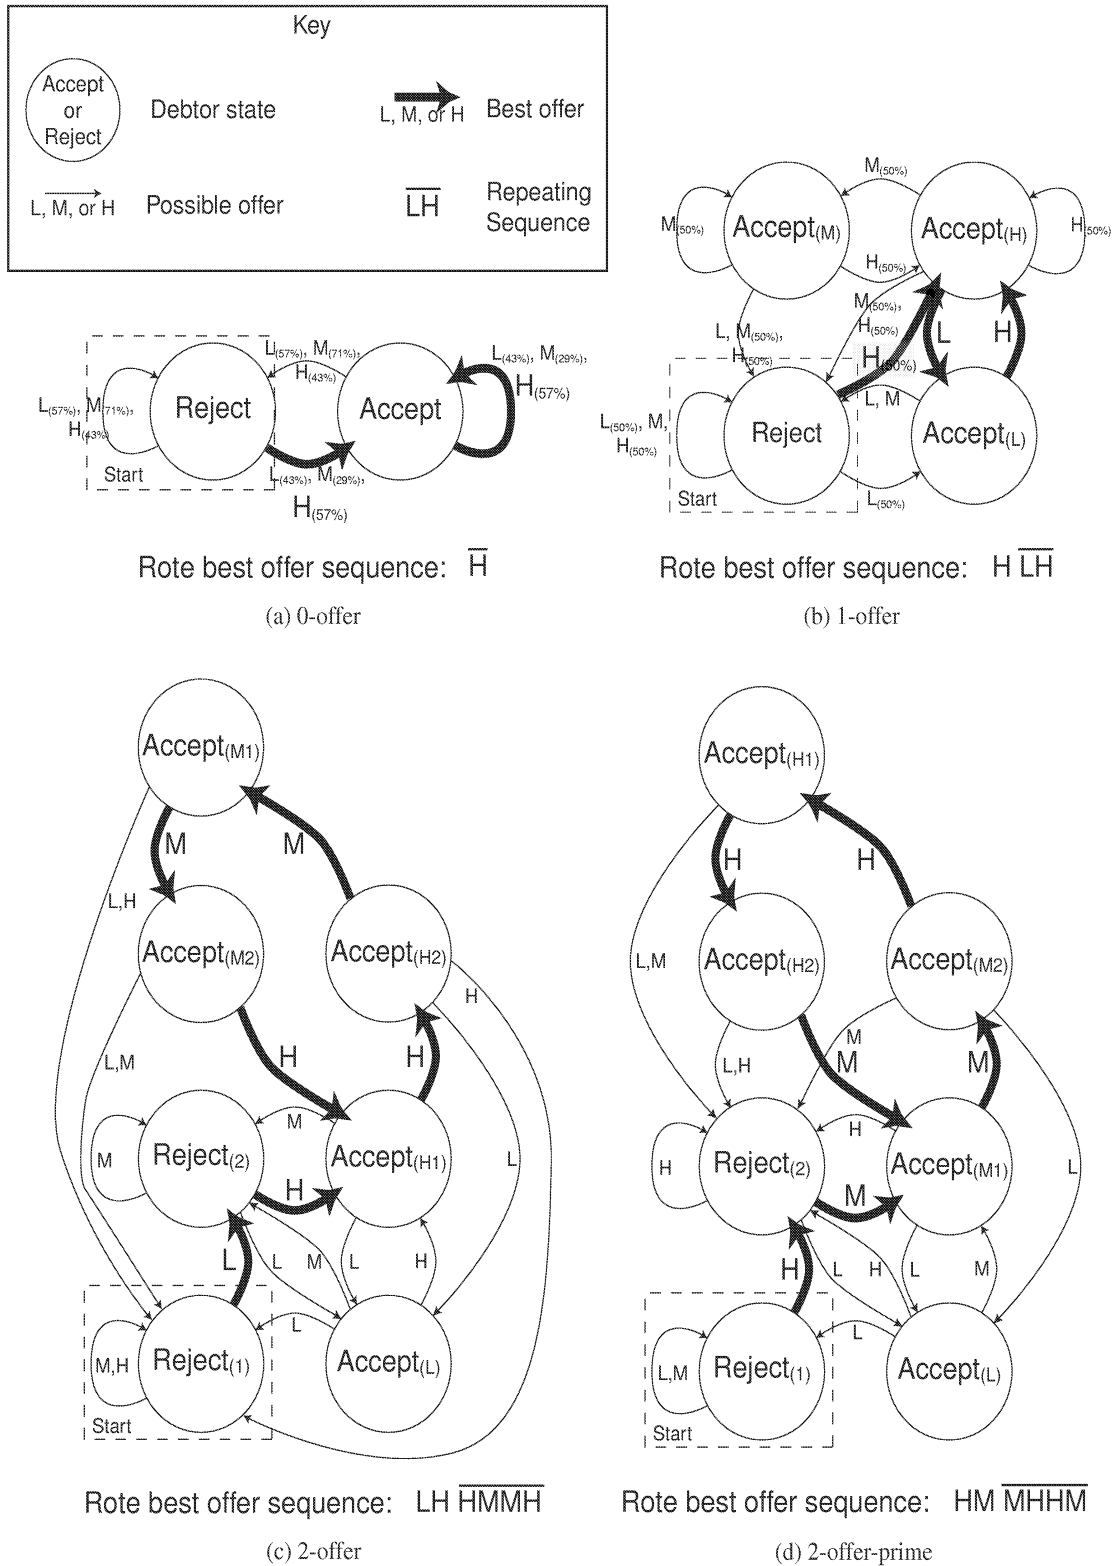


Figure 3: Debtor state transition diagrams (STD) used in the experiment. Subject offers (L, M, or H) caused debtors to transition to new states, with percentages indicating probabilistic transitions. Debtors always started a given contact in the corner of the transition diagram marked by the dashed box. The “Rote best offer sequence” for each STD indicates the sequence of highest offers with the highest probability of payoff.

have accepted if made by itself without the first request. In good-cop-bad-cop, bargainers respond to altering displays of sternness and leniency that corresponds to making high demands (e.g., H) and then backing off (e.g., to M) or vice versa (Rafaeli & Sutton, 1991).

A detailed examination of the 2-offer STD shows how FID and DIF combined with good-cop-bad-cop can be used to produce plausible explanations for the sequential patterns of debtor behavior that subjects had to learn to anticipate in order to perform effectively. In this STD, debtors are at first unwilling to accept any offer until the subject has made a show of compassion by offering L (good-cop). Even though 2-offer debtors reject this offer, it primes them to accept a high offer (H) that will move them much closer to resolving their delinquency (FID with initial rejection). 2-offer debtors will then accept one more H offer before getting the jitters and rejecting a third H offer (delayed DIF). Even though an H offer has been twice accepted, if subjects learn to pre-emptively back off to M (and not all the way to L) they can maintain agreement at a reasonably high level and avoid having to start all over building the debtor to agreement. In the functioning environment from which the task was drawn, this twist corresponds to collectors' observation that they had to maintain urgency while intermittently backing down (bad-cop then good-cop, Rafaeli & Sutton, 1991). After the subject backs off for one offer, the debtor is both willing to entertain a higher offer or completely abandon the negotiation (FID with hesitation). Similar reasoning applied to the 2-offer-prime debtor who more closely followed DIF.

As indicated in Figure 3, subjects could stay on the path of best offers with both the 2-offer and 2-offer-prime debtors by following a rote sequence that was different for each debtor. However, the repeating (overlined) portion of each debtor's sequence was generated by the same underlying sequential dependencies indicated in Table 2. For example, subjects in the repeating portion of the sequence who had just had a high offer accepted after a medium offer could infer with a high degree of certainty that their next best offer was high (H) with *both* the 2-offer and 2-offer-prime debtors. *Subjects who learned to recognize the underlying pattern of sequential dependencies, as opposed to just the rote sequence of best offers specific to the individual debtors, could transfer a high level of performance between the 2-offer and 2-offer-prime debtors.*

Evidence	Likely Debtor State	Best Offer
Accept _(H) after Accept _(M)	Accept _(H1)	H
Accept _(H) after Accept _(H)	Accept _(H2)	M*
Accept _(M) after Accept _(H)	Accept _(M1)	M
Accept _(M) after Accept _(M)	Accept _(M2)	H

* Subjects could pre-emptively lower their demands from previously accepted levels after having learned to anticipate debtors' tendency to back out of high commitments and force subjects to start over from the beginning with them (see text for details).

Table 2: Shared sequential dependencies between the 2-offer and 2-offer-prime debtors.

0-offer and 1-offer Debtors. The 0-offer and 1-offer debtors' STDs represent full and partial collapses respectively of the 2-offer debtor's STD. As such, they were meant to provide tasks in which feedback was more immediate and coherent internal representations could be formed for later learning of sequential dependencies. In the context of the decision environment on which the task is based, they represented debtors requiring simpler tactics such as might be found in the earlier stages of delinquency (Sutton, 1991).

In the 0-offer STD, debtor responses were based solely on the subjects' most recent offer. This STD was derived by first tabulating how many times a given offer led to accept across all seven nodes in the 2-offer STD and then dividing by the number of nodes. This computation gave the unconditional probability that a given offer would lead to an accept node, if all nodes were infinitely randomly sampled, an offer then randomly selected, and its outcome node observed. The probability of a given offer leading to a reject node was then computed by subtracting this probability from one, and the STD in Figure 3(a) was constructed.

For the 1-offer debtor, responses were based on the subject's offer in conjunction with the last offer-response pair. The derivation of the 1-offer STD was slightly more complex because it involved tabulating according to four cases based on the immediately preceding offer and the response it received: (1) any preceding rejected offer; (2) preceding offer low and accepted (Accept_(L)); (3) preceding offer medium and accepted (Accept_(M)); and (4) preceding offer high and accepted (Accept_(H)).² To compute the conditional likelihood of a new offer leading to an accept node for

²Other approaches for collapsing the sequential dependency were possible. This approach treated the reject nodes from the 2-offer STD equivalently to the Accept_(H) and Accept_(M) nodes.

each case, the number of times a given new offer led to an accept node was tabulated and divided by two if the offer could lead to both an accept and a reject and one otherwise. The conditional probability of the offer leading to a reject node was then one less this amount. Finally, the STD in Figure 3(b) was constructed using these probabilities.

4.1.3 Procedure and Task Summary

Evolving subjects went from debtors whose responses displayed no sequential dependencies to debtors whose responses depended on the last two offer-response pairs. This progression was tantamount to using an information system to manage subjects' workflow so that they started with easier cases that then gained in difficulty. By contrast, Consistent subjects started immediately with debtors who responded to their offer in conjunction with the last two offer-response pairs, tantamount to using the information system to manage workflows so that these subjects always bargained with the same difficulty debtor.

From a theoretical perspective, Evolving subjects were given a chance to develop internal representations of shorter dependencies before progressing to longer ones. Since the S-L model's formulation suggests that representations for shorter dependencies are the building blocks for longer ones, this approach could aid learning. However, Consistent subjects were not given the opportunity to focus first on these building blocks. Instead, they spent all of their time trying to develop internal representations suitable to the longer dependencies.

4.2 Simulation Method

To produce predictions for human decision makers, 18 instances of the S-L model were run in each of the experimental conditions and 24 in the control condition. Instantiating the S-L model required: (1) determining the models' representation of the offer options (H, M, or L); (2) hypothesizing the representation of the information considered by decision makers in making each offer; (3) and determining the number of higher level features the model could recognize. To maintain the flow of the exposition, these points are addressed in Appendix A.

4.3 Simulation and Human Subject Results

The S-L model provided predictions for decision makers' best move performance. These predictions can be evaluated: (1) by whether the S-L model correctly predicted the effects exhibited by decision makers; and (2) how closely the model fit decision maker performance. Predictions for the main effects of debtor type, learning trend, and the interaction of debtor type and learning trend were calculated using 1 *df* contrasts (Judd & McClelland, 1989).


4.3.1 Pattern of Effects

Figure 4(a) shows the number of best moves models and subjects made in the first three sessions of the experiment. In each session, performance was bounded at 240 best moves with 80 expected based on random offer selection. Overall, model and subject performance fell between these two bounds.


Models predicted that decision makers dealing with Consistent debtors would improve between sessions one and three while those dealing with Evolving debtors would decrease ($t_{34} = -4.80$, $p < 0.001$), and subjects confirmed this prediction ($t_{34} = -5.93$, $p < 0.001$). Further, the model predicted that decision makers dealing with Evolving debtors would show a significant non-linear trend in learning as characterized by the spike in performance in Session 2 ($t_{34} = 21.39$, $p < 0.001$), again confirmed by subjects ($t_{34} = 1.88$, $p < 0.05$, one-tailed). Finally, models predicted that decision makers dealing with Evolving debtors would underperform those dealing with Consistent debtors in the third session ($t_{34} = -17.4$, $p < 0.001$) as was also confirmed by subjects ($t_{34} = -4.63$, $p < 0.001$). However, a small number of Consistent subjects attained almost perfect performance by the last session while no models did. As discussed below, attaining perfect performance required deterministic application of decision rules that generally only appears late in decision makers' learning, while by design the S-L model's choice rule remains probabilistic.

Figure 4(b) shows models' and subjects' transfer and control performance. For each set of 5 contacts shown in the figure, performance was bounded at 60 best moves with 20 expected based on random offer selection.

We first examine control performance and then contrast it with transfer performance. The trans-



(a) Best moves by session for the first three sessions.



(b) Best moves by group of five contacts for transfer and control tasks.

Figure 4: Model and subject performance during training and transfer.

fer task provided an opportunity to assess the extent to which models and subjects just memorized rote sequences *or* learned the underlying dependency displayed in Table 2. For Controls, model performance was beneath the prediction for random guessing with no significant linear trend in performance but a significant quadratic trend ($t_{34} = 3.16, p < 0.01$). This prediction was largely confirmed by subjects whose overall performance was not significantly different from random guessing but with no significant trends.

In addition to differences between the two experimental conditions, of interest in transfer is the degree to which decision makers with experience differed on average from relatively naïve decision makers in the control condition. The S-L model predicted that decision makers from both experimental conditions would significantly outperform controls during the transfer task ($t_{58} = -7.1, p < 0.001$), and subject performance confirmed this prediction ($t_{58} = -4.07, p < 0.001$). The model also predicted that decision makers who had bargained with Consistent debtors during the first three sessions would significantly outperform those who had bargained with Evolving debtors ($t_{34} = -3.68, p < 0.001$), again confirmed by subjects ($t_{34} = -2.09, p < 0.05$). Finally, the S-L model predicted that decision makers in both conditions would show positive linear rates of learning ($t_{34} = 6.51, p < 0.001$) that were greater than for controls ($t_{58} = -3.64, p < 0.001$) with no difference between conditions. All three of these predictions were supported, with human subjects showing positive learning rates ($t_{34} = 4.50, p < 0.001$) that were greater than for control subjects ($t_{58} = -3.46, p < 0.001$) and no significant difference in rate between Evolving and Consistent conditions.

4.3.2 Model Fit

The second indication of how well the S-L model predicted decision maker performance was how closely it fit subjects and in which direction it deviated. For the first three sessions, the model generally underpredicted subject performance ($\overline{diff} = 37.07, t_{35} = 5.76, p < 0.001$) with the obvious exception of Evolving subjects in Session 2 (see Figure 4(a)). Taking the difference of the squared residuals within subjects and then computing the mean of this difference across subjects indicates that the model provided a better fit for decision makers bargaining with Evolving debtors

than Consistent debtors ($t_{34} = 4.56, p < 0.001$).

Similarly, the model under-predicted transfer best-move performance ($\overline{diff} = 14.47, t_{35} = 7.09, p < 0.001$). In transfer, the model provided slightly closer fits for Evolving than Consistent subjects ($t_{35} = 1.89, p < 0.05$, one-tailed). As indicated in Figure 4(b), a significant difference between both the S-L model and human subjects is that some human subjects again displayed performance indicative of a perfect execution of the underlying STD while no model instances did. The closest model instances only displayed this behavior for 1/2 of their offers.

4.4 Discussion

The S-L model predicted all significant patterns in decision makers' performance. As predicted by the model, decision makers dealing with Evolving debtors outperformed those dealing with Consistent debtors in the first two sessions of the experiment, when their debtors did not display sequential dependencies. This result partially replicates previous findings where decision makers in tasks with shorter sequential dependencies outperformed those in tasks with longer dependencies (Diehl & Serman, 1995; Gibson, 2000). In Session 3, when Consistent and Evolving subjects both bargained with 2-offer debtors whose response rule depended on longer sequential dependencies, Consistent decision makers outperformed Evolving decision makers.

This last result was not just due to the fact that Consistent decision makers had memorized idiosyncrasies in debtors' response patterns over the three sessions. As predicted by the S-L model, Consistent decision makers also outperformed Evolving decision makers when placed against a new debtor in Session 4 who displayed a different sequence of responses. *Consistent decision makers had learned the pattern of sequential dependencies shared by the two debtors better than Evolving decision makers.* Finally, even Evolving decision makers who only had short experience with sequential dependencies outperformed Controls who had no experience with sequential dependencies.

For all of this success, the model underpredicted the overall level of performance decision makers achieved, suggesting the S-L model does not provide a complete account of decision maker behavior in this task. As further discussed under limitations below, by construction, the S-L model

is limited to an account of the role of sequential pattern discovery in decision maker performance. Even with this limitation, the model was able to differentiate the impact of two learning interventions designed to improve performance and suggest the most effective intervention.

5 General Discussion and Conclusion

To describe how decision makers learn to account for sequential dependencies in evolving dynamic environments, this paper developed the S-L model. As predicted by the S-L model, decision makers' ability to bargain against opponents who displayed significant sequential dependencies was proportionate to the amount of time they spent bargaining with those opponents. Starting with "easier" opponents who displayed shorter sequential dependencies and building to "harder" opponents who displayed longer dependencies did not lead to more effective learning. As further predicted by the S-L model, decision makers did not just learn the rote sequences of behavior these opponents displayed. Rather, they learned to recognize higher level features of their opponents' behavior (signature patterns of interaction) that enabled them to identify sequential dependencies and respond appropriately to them when they encountered them in new opponents. This last result is rare if not unique in the study of dynamic decision environments.

This work has several limitations that we examine before going on to a discussion of its implications for information systems design.

5.1 Limitations

Limitations in this study relate principally to the task used and the modeling perspective applied. While the task has many important elements of dynamic tasks, it has clearly been simplified from the functioning environment on which it is based and is also simpler than many other dynamic tasks. In particular, the task abstracts the interaction between debtors and collectors to one of making offers and hearing responses. In the functioning setting of credit collections, collectors do much more, for instance: engaging in informal chit-chat to break the ice, specifically probing for information, making threats and giving encouragements (Gibson & Fichman, 2002). The limiting focus applied here is justified partly by the decomposition collectors themselves apply to their profession; they

consider bargaining a separable activity. Further, an analysis of interactions between debtors and collectors in the functioning environment indicates that one of the most significant factors in getting a debtor to discuss resolving his or her debt is proposing and re-proposing offers (Gibson & Fichman, 2002).

Do these task limitations allow the work to address findings from generally more complex environments studied elsewhere (e.g., supply chain management, Sterman, 1989b)? As noted earlier, the almost uniform finding from these environments has been that decision makers have trouble learning sequential dependencies. Decision makers in the task used here presented the same characteristics early in learning. Thus, the task, although simpler, partially replicates these previous findings.

Further, the results reported here suggest time on task as a possible reason for decision makers' difficulty learning sequential dependencies in earlier studies. Diehl and Sterman (1995) studied decision maker performance in repeated sessions with an inventory management task where the length of the dependency changed between successive sessions, similar to the Evolving condition in this experiment. Although Diehl and Sterman's (1995) subjects performance improved between sessions, they showed almost no improvement in their ability to account for sequential dependencies. The sequence learning process embedded in the S-L model suggests that, like Evolving subjects in this study, Diehl and Sterman's (1995) subjects did not have sufficient experience dealing with any given length dependency to effectively learn the higher level features that identified it.

Turning to the S-L model, its major limitation is that it was *explicitly* restricted to an account of how decision makers learned to recognize patterns of dependency between decision instances. Accepting this limitation as an attempt at a first order approximation of decision maker cognition appears warranted. Previous studies of performance in dynamic environments suggest that decision makers quickly toss aside theory-driven, explicit approaches in attempting to achieve goals (e.g., Diehl & Sterman, 1995). Further, even decision makers who are able to develop effective decision rules do so only after sustained periods of high performance during which they are not able to state coherent rules (Stanley, Mathews, Buss, & Kotler-Cope, 1989).

One possibility, consistent with the S-L model, for this last result is that until particular decision options prove to be consistently effective, decision makers entertain more than one possibility at each decision point, making it hard to state a rule. The S-L model's learning mechanism predicts that in environments where specific decision options prove to be consistently effective, decision makers will tend more and more to these options, leading to a systematic pattern of behavior. When decision makers are asked to reflect over this consistent pattern of behavior shortly after periods of sustained high performance, stating a rule, essentially just providing a description of their own consistent pattern of behavior, should be easier. Just as the S-L model predicts that decision makers' performance in dynamic tasks is a function of their learning higher level features that consistently identify sequential dependencies, so may be decision makers' ability to state rules about how to achieve high task performance.

5.2 Implications

To effectively support decision makers in dynamic environments, our conceptualization of their capabilities is critical. To date, most decision support efforts have focused on the deficits hypothesis: decision makers' poor performance in dynamic tasks is due to their limited capacity to explicitly represent sequential dependencies and forecast their effects (Sterman, 2000; Sweeney & Sterman, 2000). The work reported here indicates that we may profit by shifting our focus toward how decision makers learn to recognize and respond to patterns of sequential dependency in dynamic environments. Such a shift has important implications for designing information systems to support these decision makers, relating primarily to workflow design with possible future applications to interfaces.

First, the results reported here indicate that designing system workflows to enhance decision makers' capacity to detect patterns of sequential dependency improves performance in evolving environments. This capacity exists even at the novice level but appears to require significant experience to refine. For instance, in a recently studied inventory management task, subjects who had prior experience with a more complex inventory task displayed significantly better ability in predicting the overall future shape of inventory's trajectory than subjects who did not have this

experience (Sweeney & Sterman, 2000, Manufacturing Case). However, these subjects still did not display a level of understanding that would have led to re-equilibration of inventory. In the work reported here, when novices' workflow focused them on debtors displaying a consistent pattern of sequential dependencies over a longer period of time, they were able to recognize this signature pattern in a new type of bargaining opponent and quickly achieve a high level of performance.

From a practical standpoint, the question is how to implement this workflow strategy in functioning environments. The work reported here suggests that, in defining system requirements, task experts should be able to identify the higher level features they use to identify dependencies and thereby the data that helps them determine whether those features are present. In the functioning environment on which the task used here was based, workflows were derived using factors that experts felt affected the likelihood of the debtor responding to different negotiating strategies such as time delinquent and amount of delinquency (Gibson & Fichman, 2002).

How does this workflow strategy fare against the level of change inherent in dynamic environments? Three distinct cases that may occur in combination present themselves. In the first, the same features remain relevant for identifying the dependency, but the effectiveness of different decision options changes. For instance, customers who once readily responded to efforts to get them to raise their initial commitments may become willing to raise these commitments only half as much. In this case, information systems implementing the workflow strategy will not require change because the same factors remain relevant for identifying sequential dependencies. However, decision makers will have to learn new responses to the dependency. The S-L model suggests that this learning will be proportionately less than suggested by the performance of the Consistent subjects in this study because decision makers will not have to relearn higher level features that identify the dependencies.

In the second case, a more serious challenge presents itself when previously identified higher level features are no longer relevant in recognizing sequential dependencies, but the data required to recognize the dependencies does not change. This case is similar to that faced by Diehl and Sterman's (1995) subjects and Evolving subjects in this study who dealt with sequential dependencies of different lengths in each session. As in the first case, the system does not have to be

altered. However, lengthier time on task, proportionate to that for Consistent subjects in this study, is required for decision makers to relearn higher level features that identify relevant sequential dependencies.

In the third case, the most serious challenge of the three cases presents itself when different data elements become important for identifying relevant higher level features. For example in credit collections, similarly to other sales and customer service tasks, the natural or economic conditions in a given geographic region may change causing customers to behave differently. In such cases, knowing the customer's region may become more important for identifying the state of the interaction than it was previously. Any solution to this case involves both changing the system and relearning higher level features that identify relevant dependencies. In such cases, an information system that cannot be easily reconfigured to display more relevant information will be a hindrance. Of course such flexibility is not without costs, and the trade off is likely to be domain specific.

Finally, the work reported here suggests future research that can inform the design of interfaces for dynamic decision environments. The task interface used in this study represented state variables and offer options as discrete entities. Most, if not all, previous experiments with dynamic tasks have used continuous indicators for current state and decision options (e.g., expected software project duration and number of software developers, Sengupta & Abdel-Hamid, 1993). The use of continuous state indicators and decision variables is justified by the observation that, in the aggregate, many business phenomena can be viewed in terms of continuous flows (Sweeney & Sterman, 2000). The problem with using continuous indicators is that they may make recognizing states and their associated actions more difficult by making them seem less similar. For instance, it is less clear that an individual who in a sequence of offers agreed to pay \$150 and then \$225 is similar to one who agreed to pay \$175 and then \$250 than it is if both individuals are indicated as having agreed to pay the lowest possible amount (equivalent to an L offer in the work reported here) and then an amount equivalent to their usual monthly payment (equivalent to an M offer here). In tacit recognition that decision makers perform better when similarities between states are more salient, many customer

service organizations gear their information systems toward providing categorical data that facilitates drawing similarities and distinctions over patterns of behavior (e.g., the offer category scheme just mentioned). Formally testing these conjectures concerning the differential effectiveness of interfaces that use discrete and continuous variables to represent dynamic environments is a topic for future research.

5.3 Conclusion

In increasingly prevalent dynamic decision environments such as direct sales, customer support, and electronically mediated bargaining, decision makers must make sequences of interdependent decisions. Past decision support efforts have focused on relieving the decision maker of the need to explicitly account for dependencies between decisions. However, such systems themselves are fragile to change and further do not focus on enhancing decision makers' capacities to deal with change. This study has presented an alternative strategy that defines information systems requirements in terms of enhancing decision makers' ability to learn patterns of sequential dependency between decisions. In so doing, the study has introduced a simulation model of how decision makers learn sequential dependencies in order to predict the relative impact of different strategies for improving decision makers' performance. When a system was used to manage workflows in a way predicted by the model to enhance learning, decision makers were able to learn underlying patterns of sequential dependencies, not just their superficial manifestation. This result suggests that a fundamental shift, *away from substituting for cognitive deficits and toward enhancing sequential pattern learning*, is appropriate in how we support decision makers in dynamic environments. As detailed in this study, this shift has important implications for designing system workflows and potential future applications in interface design.

Appendix A: Detailed Simulation Settings and Assumptions

Instantiating the learning model required: (1) determining the models' representation of the offer options (H, M, or L); (2) hypothesizing the representation of the information considered by decision makers in making each offer; (3) and determining the number of higher level features the model

could recognize. We address each of these in turn and then briefly examine the mechanics of how models made offers and learned in the simulation study.

The three offer possibilities were represented as separate action options using Equation (1). The principal information subjects used as they made offers appeared to be their last offer paired with the debtor's response. Subjects were assumed to perceive offer-debtor-response pairs as a single discrete unit, a representation commonly assumed in cognitive game theory (Fudenberg & Levine, 1998, Chapter 4), repeated decision making (Dienes & Fahey, 1995), and skill automatization (Logan, 1988). For instance, if the subject made a high offer and the debtor rejected, the subject would encode that offer-response pair as high-offer-reject. In the simulation, a 1-of-n vector representation was used to capture this perceptual encoding. Since there were three possible offers and two possible responses, the vector had a length of six with 1 being placed in the position of the offer-response pair that had occurred and 0 in the other positions.³

Given the construction of the task, higher level features were patterns of offers and responses. As for determining the number of higher level features the model should be able to recognize, theory is relatively silent, even under idealized circumstances (Mitchell, 1997, pp. 218–220). Therefore, the number of higher level features that the model was able to recognize was varied between five and fifteen in increments of five. No effect was found on the overall pattern of results, but models with fifteen higher level features appeared to learn more reliably. Therefore, models capable of recognizing 15 higher-level features were used.

At the start of each simulation experiment, the model weights were set to small random values close to 0. This assumption caused models to display individual learning characteristics since the different weights for each model represented different priors concerning which offer to select at the start of learning (Bishop, 1996). Therefore, eighteen models were run in each condition to estimate average performance for one and two-stage models.

Since the model's behavior was fully determined by the weights, the weights' initial random

³Alternative representations are possible, in particular representing the last offer and debtor response separately. This and other alternative representations were tried without affecting the general pattern of simulation results reported presently.

values caused the model to initially display random decision behavior. This behavior does not follow Raiffa's (1982) observations of convergence during bargaining which suggest that bargainers may follow a strategy of asking for more if the previous offer is accepted and asking for less if the previous offer is rejected. To instantiate this initial behavior in models, they were trained to make offers that were one level lower than offers that had immediately just been rejected and to make offers that were one level higher than offers that had been immediately accepted for five contacts before commencing the simulation experiment of human subjects.

The mechanics of both this initial pretraining and the learning during the simulated experiment were identical. For each offer, a vector representing the decision context was presented as the input to the higher level features. It contained the 1-of-n representation of the last offer-response pair and the activations of the higher level features at the moment of making the last offer.⁴ The current activation of each higher level feature was then calculated by substituting the input vector and appropriate weights into Equation (2), and then substituting that result into Equation (3). The activation of the higher level features then served as the inputs into the choice of offer to make. The choice of offer was determined by calculating each offer's cumulative evidence using Equation (2), estimating the probability of the offer being chosen using Equation (1), and finally choosing randomly from the offers based on the likelihoods just estimated.

After the offer was made the debtor algorithm responded using one of the STD's from Figure 3 based on session and experimental condition for each offer in a contact. This response was used by the backpropagation algorithm in learning as specified by Equations (4)–(7). In the results reported, the learning rate η from Equation (4) was set to 0.08 with similar patterns of results produced when η was varied between 0.05 and 0.12.

References

Beam, C., Segev, A., Bichler, M., & Krishnan, R. (1999). On negotiations and deal making in electronic markets. *Information Systems Frontiers*, 1(3), 241–258.

⁴At the start of each contact, no offers had been made, so the value for the offer-response pair and all prior higher level feature activations were set to 0.

- Bishop, C. M. (1996). *Neural networks for pattern recognition*. New York: Oxford University Press.
- Brehmer, B. (1995). Feedback delays in complex dynamic decision tasks. In P. Frensch, & J. Funke (Eds.), *Complex problem solving: The European perspective*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty*, *19*(1–3), 7–42.
- Cialdini, R. B. (1984). *Influence: The psychology of persuasion*. New York: William Morrow.
- Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*, 235–253.
- Davis, F. D., & Kottemann, J. E. (1995). Determinants of decision rule use in a production planning task. *Organizational Behavior and Human Decision Processes*, *64*(2), 145–157.
- Diehl, E., & Serman, J. D. (1995). Effects of feedback complexity on dynamic decision making. *Organizational Behavior and Human Decision Processes*, *62*(2), 198–215.
- Dienes, Z., & Fahey, R. (1995). Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(4), 1–15.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*(2), 179–211.
- Fudenberg, D., & Levine, D. K. (1998). *The theory of learning in games*. Cambridge: The MIT Press.
- Gibson, F. P. (2000). Feedback delays: How can decision makers learn not to buy a new car every time the garage is empty? *Organizational Behavior and Human Decision Processes*, *83*(1), 141–166.
- Gibson, F. P., & Fichman, M. (2002). When threats and encouragements are effective: The case of credit collectors. *Unpublished Manuscript*, (in revision).

- Joslyn, S., & Hunt, E. (1998). Evaluating individual differences in response to time-pressure situations. *Journal of Experimental Psychology: Applied*, 4(1), 16–43.
- Judd, C. M., & McClelland, G. H. (1989). *Data analysis: A model comparison approach*. San Diego, CA: Harcourt Brace Jovanovich.
- Kanfer, R., & Ackerman, P. (1989). Motivation and cognitive abilities: An integrative/aptitude-treatment interaction approach to skill acquisition. *Journal of Applied Psychology*, 74(4), 657–690.
- Kimbrough, S. O., Wu, D. J., & Zhong, F. (2002). Computers play the beer game: Can artificial agents manage supply chains. *Decision Support Systems*, 33, 323–333.
- Klein, G. A., Orasanu, J., Calderwood, R., & Zsombok, C. E. (Eds.). (1993). *Decision making in action: Models and methods*. Norwood, NJ: Ablex Publishing Corporation.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95(4), 492–527.
- McClelland, J. L. (2001). *The mit encyclopedia of the cognitive sciences*, Chap. Cognitive Modeling, Connectionist. Cambridge, MA: MIT Press. available on-line at <http://cognet.mit.edu/MITECS/Entry/mcclelland.html>.
- Mitchell, T. M. (1997). *Machine learning*. New York: McGraw-Hill.
- Rafaeli, A., & Sutton, R. I. (1991). Emotional contrast strategies as means of social influence: Lessons from criminal interrogators and bill collectors. *Academy of Management Journal*, 34(4), 749–775.
- Raiffa, H. (1982). *The art and science of negotiation*. Cambridge, MA: Harvard University Press.
- Rohde, D. L. T., & Plaut, D. C. (1999). Language acquisition in the absence of explicit negative evidence: How important is starting small? *Cognition*, 72, 67–109.
- Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8, 164–212.

- Rumelhart, D. E., Durbin, R., Golden, R., & Chauvin, Y. (1995). Backpropagation: The basic theory. In Y. Chauvin, & D. E. Rumelhart (Eds.), *Back-propagation: Theory, architectures, and applications*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Sengupta, K., & Abdel-Hamid, T. K. (1993). Alternative conceptions of feedback in dynamic decision environments: An experimental investigation. *Management Science*, *39*(4), 411–428.
- Stanley, W. B., Mathews, R. C., Buss, R. R., & Kotler-Cope, S. (1989). Insight without awareness: On the interaction of verbalization, instruction, and practice in a simulated process control task. *Quarterly Journal of Experimental Psychology*, *41A*(3), 553–577.
- Sterman, J. D. (1989a). Misperceptions of feedback in dynamic decision making. *Organizational Behavior and Human Decision Processes*, *43*, 301–335.
- Sterman, J. D. (1989b). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science*, *35*(3), 321–339.
- Sterman, J. D. (1994). Learning in and about complex systems. *System Dynamics Review*, *10*(2–3), 291–330.
- Sterman, J. D. (2000). *Business dynamics: Systems thinking and modeling for a complex world*. Boston: Irwin McGraw-Hill.
- Sutton, R. I. (1991). Maintaining norms about expressed emotions: The case of bill collectors. *Administrative Science Quarterly*, *36*, 245–268.
- Sweeney, L. B., & Sterman, J. D. (2000). Bathtub dynamics: Initial results of a systems thinking inventory. *System Dynamics Review*, *16*(4), 249–286.