Citation for published version:

Krause, A 2009, 'Learning and herding using case-based decisions with local interactions', *IEEE Transactions on Systems Man and Cybernetics - Part A: Systems and Humans*, vol. 39, no. 3, pp. 662-669.
<https://doi.org/10.1109/tsmca.2009.2014542>

DOI:

[10.1109/tsmca.2009.2014542](https://doi.org/10.1109/tsmca.2009.2014542)

Publication date:

2009

[Link to publication](#)

Copyright © 2009 IEEE.

Reprinted from IEEE Transactions on Systems Man and Cybernetics Part A - Systems and Humans.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Bath's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org.

By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Krause, A., 2009. Learning and Herding Using Case-Based Decisions With Local Interactions. *IEEE Transactions on Systems Man and Cybernetics Part A - Systems and Humans*, 39 (3), pp. 662-669.

Official URL: <http://dx.doi.org/10.1109/tsmca.2009.2014542>

Copyright © 2009 IEEE.

Reprinted from *IEEE Transactions on Systems Man and Cybernetics Part A - Systems and Humans*.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Bath's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org.

By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Correspondence

Learning and Herding Using Case-Based Decisions With Local Interactions

Andreas Krause

Abstract—We evaluate repeated decisions of individuals using a variant of the case-based decision theory (CBDT), where individuals base their decisions on their own past experience and the experience of neighboring individuals. Looking at a range of scenarios to determine the successful outcome of a decision, we find that for learning to occur, agents must have a sufficient number of neighbors to learn from and access to sufficiently independent information. If these conditions are not fulfilled, we can easily observe herding in cases where no best decision exists.

Index Terms—Decision making, economics, simulation.

I. INTRODUCTION

When facing a decision in unfamiliar circumstances, people tend to utilize their experiences from similar situations in the past. These experiences can be derived from their own personal experiences or those of other decision makers whose circumstances are well known to them. The degree of similarity between the circumstances will obviously have to be taken into account when analyzing these experiences.

The idea of using past experiences to inform decision making has formally been operationalized as the case-based decision theory (CBDT) in a series of papers [1]–[8]. The CBDT has received some attention in a number of applications in financial markets [9]–[11], voting [12], production theory [13], and consumer theory [14], [15]. In contrast to these applications, not much notice has been given to the dynamics of repeated decisions using the CBDT [16] or on how decision makers learn from their experiences [17], apart from chapters in [7].¹

In this correspondence, we will explore both aspects, the dynamics of decisions and whether decision makers learn the optimal decision over time. To this end, we introduce a model implementing the CBDT in Section II. Then, we determine different scenarios on how successful outcomes of a decision are determined in Section III and analyze the results of our computer experiments, which are described in Section IV, in Section V. Section VI then concludes the findings.

II. CBDT ON A NETWORK

Let us denote by $\mathcal{C}$ the set of all possible circumstances. We can then define a similarity function s as

$$s : \mathcal{C} \times \mathcal{C} \rightarrow [0; 1] \quad (1)$$

Manuscript received May 19, 2007; revised March 11, 2008 and April 22, 2008. First published February 24, 2009; current version published April 17, 2009. This paper was recommended by Associate Editor R. Popp.

The author is with the School of Management, University of Bath, Bath BA2 7AY, U.K. (e-mail: mnsak@bath.ac.uk).

Digital Object Identifier 10.1109/TSMCA.2009.2014542

¹Case-based reasoning, which is widely used in artificial intelligence, provides a broadly similar approach [18]. Our model here, however, differs in that no previous cases are explicitly selected (we select all cases available to the decision maker), and we do not revise cases but merely acknowledge the outcome of the decision made.

where for any $i, j \in \mathcal{C}$, $s(i, j) = 1$ can be interpreted as the two circumstances being identical and $s(i, j) = 0$ as being completely unrelated to each other.² From this interpretation, it is obvious that $\forall i \in \mathcal{C} : s(i, i) = 1$. A more extensive discussion of the way similarities can be measured is given in [19], but for the purpose of this correspondence, we can take the similarities as exogenously given. Katsikopoulos and Fasolo [20] discuss multiattribute models, which can also be interpreted as introducing a similarity function to aid decision making. Our network is different from Bayesian networks as, for example, used in [21], in that no Bayesian updating of beliefs is conducted, but a rather more simplified approach to decision making is taken, as we will see later in this correspondence.

A decision is a mapping from the set of circumstances into the set of all possible decisions, which is denoted by $\mathcal{D}$, i.e.,

$$d : \mathcal{C} \rightarrow \mathcal{D} \quad (2)$$

and we define an outcome function as

$$u : \mathcal{D} \rightarrow \mathcal{O}. \quad (3)$$

If we now assume that the similarity reduces over time with a constant factor $0 < \lambda < 1$, we can define the following value function for any $d_t(i) \in \mathcal{D}$, where $i \in \mathcal{C}$ and t indicates the time of the decision:

$$V(d_t(i)) = \sum_{\tau=1}^M \left(\lambda^\tau \sum_{j \in \mathcal{C}} s(i, j) u(d_{t-\tau}(j)) \mathbf{1}(d_{t-\tau}(j)) \right) \quad (4)$$

where $\mathbf{1}(d_{t-\tau}(j))$ denotes an indicator function, which is 1 if the two decisions are equivalent, i.e., $d_{t-\tau}(j) = d_t(i)$, and 0 otherwise. In every time step, the decision maker will choose the action that provides him with the highest value function, i.e.,

$$d_t^* \in \arg \max_{d_t \in \mathcal{D}} V(d_t(i)). \quad (5)$$

Using this formulation of the CBDT, we can now continue to define a simplified setup that suits our needs. First, we can restrict the set of decisions to $\mathcal{D} = \{\text{buy; sell}\}$ and let the possible outcomes be either “failure” or “success,” such that $\mathcal{O} = \{-1; +1\}$.

We can now interpret decision making as that of a number of agents making repeated decisions on similar circumstances over time, where the similarity is denoted by $s(i, j)$ if agent j shares his experiences with agent i , i.e., they are interacting with each other. If two agents are not interacting with each other, they cannot learn from their experiences, i.e., $s(i, j) = 0$. For those instances where agents are interacting, we find that $s(i, j) > 0$. We allow for the case where the interactions are asymmetric, i.e., in general, we find that $s(i, j) \neq s(j, i)$.

Using this interpretation, we can construct a matrix $\mathbf{S} = [s(i, j)]$ capturing all the interactions between agents. This matrix can be understood as an adjacency matrix of a graph representing the network of interactions between agents. We can create a random network of interactions by determining a positive entry in the matrix with

²We could also define $s : \mathcal{C} \times \mathcal{C} \rightarrow [-1; 1]$, where negative similarities would indicate that the two circumstances are opposite to each other. This simple extension does not change the arguments presented in this correspondence.

probability p , and the entry itself is randomly taken from a uniform distribution on the interval $[0; 1]$.

It has to be stressed that the CBDT as used in this correspondence deviates from the conventional approach. In the original model, agents only learn from their own past behavior, whereas we also allow learning from the behavior of other agents. While such behavior is realistic in many circumstances, it changes the dynamics of the system investigated. The experience of other agents will easily dominate an agent's own experience, even for relatively small p . While such a setup seems not to have been intended when introducing the CBDT, it is a natural extension of that framework.

We can now use the model presented here to evaluate the evolution of aggregate decisions over time, as well as learning of the best decision over time. Using different rules on how the successful outcome is determined, we can compare our results for a variety of scenarios.

III. DETERMINATION OF SUCCESSFUL OUTCOMES

We consider seven different scenarios on how the successful outcome is determined. For each scenario, we define a benchmark outcome that would be the optimal behavior of agents. Apart from the minority outcomes, the chosen benchmarks would also be the outcome of a Bayesian learning process.

In the *first* scenario, we assume that the successful outcome is randomly determined with equal probability and independently for each agent, as well as across time periods (**independent fair outcomes**). Hence, in this scenario, there is no real value in the experience of other agents or the past. Here, no best strategy exists, and any strategy of agents should be equivalent. As a benchmark in this case, we can use a scenario in which all agents behave randomly and compare the results to this benchmark case.

The *second* scenario is very similar to the first scenario except that the probability of "buy" being the successful decision is set at 75% (**independent biased outcomes**). Hence, all agents should always seek to choose "buy" as this would give them the best chances of success.

The *third* scenario is also similar to the first scenario; however, here, the successful outcome is identical for all agents in a given time period (**identical fair outcomes**). However, it remains independently across time periods, and the probabilities for either decision being successful is still 50%. In this case again, the past experience of other agents has no value, and agents can follow any strategy. Thus, the benchmark in this case would be the same as in scenario 1, i.e., all agents behave randomly.

The *fourth* scenario is based on the third scenario, and similarly to scenario 2, the probability of success for the strategy "buy" is set at 75% (**identical biased outcomes**). Thus, similarly to that case, all agents should always choose "buy."

The *fifth* scenario assumes a correlation between the successful outcomes across agents but, again, no correlation across time periods, and the *ex ante* probability of success is equal for all strategies (**correlated fair outcomes**). The correlation between the successful outcomes is determined by the adjacency matrix of the interactions, i.e., $\mathbf{S}$.³ Similarly to scenarios 1 and 3, no information can be gained from the past experience of agents, and the appropriate benchmark is again the random choice of strategy by agents.

The *sixth* scenario endogenously determines the successful outcome as the minority outcome, i.e., the decision that has been chosen by the least number of agents is the successful outcome (**minority**

TABLE I
VARIABLES FOR THE BENCHMARK CASE

Scenario	Herding	Length	Volatility
1 - independent fair outcomes	≈ 5	≈ 2	≈ 11
2 - independent biased outcomes	50	10,000	0
3 - identical fair outcomes	≈ 5	≈ 2	≈ 11
4 - identical biased outcomes	50	10,000	0
5 - correlated fair outcomes	≈ 5	≈ 2	≈ 11
6 - minority outcomes	≈ 5	≈ 2	≈ 11
7 - majority outcomes	50	10,000	0

outcomes). Such a scenario has extensively been explored using the minority game [22], where it has been applied to financial markets. In this scenario, we can use the results from minority games and find that a suitable benchmark would again be all agents behave randomly.

In the *seventh* and final scenario, the successful outcome is determined as the majority outcome, i.e., the decision that has been chosen by the largest number of agents is the successful outcome (**majority outcomes**). In this scenario, agents would ideally coordinate all their decisions and then would always be successful. Thus, a good benchmark would be that all agents make the same decision, either buying or selling.

The rationale behind choosing these scenarios was to evaluate cases in which no optimal behavior exists (scenarios 1 and 3), as well as scenarios in which it is present (scenarios 2 and 4). The scenarios of independent outcomes (scenarios 1 and 2) and identical outcomes (scenarios 3 and 4) will allow evaluating the influence that the information structure has on the outcomes. An intermediate scenario in the information structure (scenario 5) was also chosen to ensure that not only extreme cases are looked at. Apart from exogenous outcomes (scenarios 1–5), we also investigate cases in which the outcomes are endogenously determined (scenarios 6 and 7). Due to the popularity of the minority game in agent-based computational finance, we chose this scenario (scenario 6) and its opposite, i.e., the majority outcome (scenario 7), as a contrast, both requiring the coordination of the choices of agents.

Using the aforementioned seven scenarios, we can now proceed to set up the computer experiments, as described in the next section, and evaluate the outcomes.

IV. COMPUTER EXPERIMENTS

We conduct Monte Carlo simulations of the aforementioned model for a number of parameter constellations using fixed settings.

- 1) The memory length is fixed at $M = 100$, and we are investigating a system with $N = 100$ agents.
- 2) The adjacency matrix $\mathbf{S}$ is a random matrix where the entry is nonzero with probability p and the nonzero entries are taken from a uniform distribution on the interval $[0; 1]$.
- 3) Each experiment uses a new randomly created adjacency matrix.
- 4) The remaining parameters are taken from the range $p \in [0; 1]$ and $\lambda \in [0; 1]$.
- 5) The experiment is initialized with a random history of successes and failure.
- 6) Each experiment consists of 11 000 time steps, of which the first 1000 are disregarded in the analysis.
- 7) For each parameter constellation, the presented results are obtained from averaging the outcomes of 100 experiments.⁴

The parameter p controlling the number of nonzero entries in the adjacency matrix can be interpreted as a measure of the number of

³It has to be noted that the adjacency matrix does not need to be positive definite as required for a correlation matrix. This problem is not further addressed in this correspondence but does not affect the outcome of our analysis.

⁴The standard deviations across experiments for all parameters investigated are very low, such that we can conclude that the results we derived are stable.

TABLE II
OVERVIEW OF THE OUTCOME OF COMPUTER EXPERIMENTS

Scenario	Herding	Length	Volatility
1 - independent fair outcomes	increasing in interactions	increasing in memory	increasing in interactions and decreasing in memory
2 - independent biased outcomes	all buying for sufficiently high interactions - figure 3 identifies the boundary for this observation	all buying for sufficiently high interactions - figure 3 identifies the boundary for this observation	decreasing in memory and maximal at interactions of about 0.1-0.3 - figure 3 identifies the maximum more closely
3 - identical fair outcomes	as random for $p > 0.1$	as random for $p > 0.1$	as random for $p > 0.1$
4 - identical biased outcomes	as random for $p > 0.1$	as random for $p > 0.1$	as random for $p > 0.1$
5 - correlated fair outcomes	increasing in interactions and decreasing in memory	increasing in memory	decreasing in memory
6 - minority outcomes	as random for $p > 0.1$	as random for $p > 0.1$	as random for $p > 0.1$
7 - majority outcomes	as random for $p > 0.1$	as random for $p > 0.1$	as random for $p > 0.1$

interactions between agents. Although we fixed the memory at $M = 100$, the discount factor λ can be used as a good proxy for the memory length. With a low λ , past events only receive a very low impact on decisions, thus corresponding to short memory, whereas a high value of λ makes past events more relevant, thus corresponding to long memory. We can therefore use λ as a measure for the memory length of agents.

Using this setup, we will be investigating the three variables of aggregate decision making.

- 1) *Herding*. The deviation from equally balanced buy and sell decisions: $H = |\# \text{ agents buying} - 50| \in [0; 50]$. This variable allows us to measure by how much agents' decisions are biased toward one of the decisions.
- 2) *Length of herding*. The number of time steps between changes in the sign of $(\# \text{ agents buying} - 50)$, i.e., the number of time steps between a switch from a majority of agents buying to a majority selling, or vice versa. This length will positively be related to the autocorrelation of the number of agents buying. Using this variable, we can identify how prolonged any bias is until the decisions of the agents are reversed.
- 3) *Volatility*. The average absolute change between subsequent time periods in the number of agents buying, i.e., the average deviation from the mean number of agents buying. In contrast, herding measures the average deviation from a fixed benchmark of 50 agents buying. Using this volatility, it is possible to assess how well agents are coordinating their decisions over time and how persistent aggregate decisions are.

We can use these variables, averaged over the entire experiment, to evaluate the experiments. If agents are learning over time, we should observe the properties listed in Table I, which are derived from the benchmark cases discussed in the previous section. The next section will now analyze the actual outcomes of the experiments. All the figures used depict these variables relative to the case of random decision making, which has been normalized to unity for simplicity.

V. RESULTS OF THE EXPERIMENTS

In this section, we will explore the results of the computer experiments as outlined in the previous sections. Apart from the general observations of the resulting time series of aggregate outcomes, we will particularly focus on the ability of agents to learn the best decisions.

A. Some Observations

In this section, we will briefly characterize the main results that we found from our computer experiments, as summarized in Table II.

From this table and Figs. 1–8, we first observe that in the case of biased outcomes (scenarios 2 and 4) and majority outcomes (scenario 7)—where an optimal strategy exists—agents only learn the optimal strategy in the case of independent fair outcomes (scenario 2) once the interactions are sufficiently high.⁵ The observed herding is such that all agents are buying, i.e., they choose the optimal action whenever significant herding is observed. In the other two scenarios, however, agents never learn the optimal strategy and, on aggregate, behave like they were making random decisions. In all other cases where there is no best strategy, except for the case of independent fair outcomes (scenario 1), agents tend to behave as if they make random decisions. Only in scenario 1 do we observe extensive herding (for high interactions), which is long lasting (for long memory).

We furthermore observe that herding tends to become more prolonged with increasing memory in those scenarios of independent and correlated fair outcomes (scenarios 1 and 5) and a tendency to lower volatility in those cases.

We clearly see from these results that the variant of the CBDT used for the communications network, in general, does not lead to learning the optimal behavior.⁶ In the case of the minority outcomes, it only improves the coordination for agents with very long memory, and only in the case of independent biased outcomes (scenario 2) do we observe some learning of the best decision. This at first surprising result will be explained in the next section.

B. Analysis of Computer Experiments

As previously shown, agents do not commonly learn the optimal strategy; from the experiments conducted, together with the setup of the model, we can deduct the reason for this result. For agents to learn, they have to be able to adequately process information. It becomes apparent from the case of independently biased outcomes (scenario 2) that to learn, agents need a large number of interactions with other

⁵We found no evidence for a transition of the behavior of agents at the emergence of a giant component $p = 1/(N - 1)$ in the network. The apparent change in behavior for minority and majority outcomes (scenarios 6 and 7) for low p is not related to this transition; a closer analysis revealed that the change in behavior does not occur at $p = 1/99$ but at higher values of p . We also found no evidence for the behavior to change with the number of agents N . It thus suggests that the transition observed here is independent of the percolation threshold of the random graph.

⁶While we cannot exclude the possibility of a very slow learning process, we found no evidence for it. Extending the analysis for a small number of parameter constellations in various scenarios to 1 000 000 time steps and analyzing only the final 10 000 observations did not show any significant differences to the results reported here.

Outcomes independent across agents with equal probability

Outcomes independent across agents with equal probability

Outcomes independent across agents with equal probability

Fig. 1. Results for independent fair outcomes (scenario 1).

agents, i.e., they use the experience of other agents, whereas memory is of secondary significance for learning.⁷

The reason for this finding is that the additional independent experience provided by a sufficiently large number of agents provides additional information, and agents will easily realize the better of

⁷To facilitate the comparative statics in this section, the reader might want to consider a number of reference points, which can be chosen in the figures provided.

- 1) $p = 0$: Here, agents are not interacting with each other, and thus, any learning would only be based on their own experience.
- 2) $p = 1$: Agents learn from all other agents.
- 3) $\lambda = 0$: The absence of memory implies that no learning takes place,
- 4) $\lambda = 1$: Learning takes all past experiences into account at an equal rate.

Outcomes independent across agents with $\text{Prob}(\text{buy} \rightarrow \text{success}) = 0.75$ Outcomes independent across agents with $\text{Prob}(\text{buy} \rightarrow \text{success}) = 0.75$ Outcomes independent across agents with $\text{Prob}(\text{buy} \rightarrow \text{success}) = 0.75$

Fig. 2. Results for independent biased outcomes (scenario 2).

the two strategies. Once the number of interactions is sufficiently high, agents realize the optimal strategy and then consistently apply it. We observe a sharp transition from the random strategy choice to following the optimal strategy, providing evidence that the experience is not gradually built in as it increases but shows a sharp phase transition. Long memory in itself is not sufficient as the discounting of past experience renders most earlier experience useless, not providing sufficient information for learning to take place, although it affects the number of interactions required for the phase transition to occur.

In the case of identical biased outcomes (scenario 4), we do not observe any learning at all, even for a large number of interactions. The reason for this finding is that the experience of other agents does not provide any additional information for decision making as all agents receive the same outcome. As we have seen from the case of independent biased outcomes (scenario 2), only long memory is not sufficient for learning to occur; thus, in this scenario, no learning is observed, and agents behave randomly.

Fig. 3. Contour plot of the results in Fig. 2.

These two findings seem to contradict the result in [14], which would imply that for $\lambda = 1$ and $p = 0$, the agents should be learning the optimal outcome. As we found no indication of an extremely slow learning process that goes well beyond the 11 000 time periods investigated, it suggests that this result is highly nonstable for even small degrees of interactions and less than perfect memory. Further investigations of the stability of the results in [14] seem appropriate in light of our findings.

The same argument as in scenario 4 can be brought forward for the majority outcomes (scenario 7). Here, we clearly see that the reason for one strategy to dominate is irrelevant for the decision-making process. Agents are not interested in the structure of the economy but only the best outcomes, i.e., they do not attempt to learn the structure of the economy. For this reason, there is no evidence for any difference in the behavior between the identical biased and majority outcomes (scenarios 4 and 7).

In the scenarios with fair outcomes (scenarios 1, 3, and 5), no best strategy exists, and thus, we should not expect any learning. However, in the scenarios with independent and correlated outcomes (scenarios 1 and 5), we observe some significant and long-lasting herding, whereas this is not found for identical outcomes (scenario 3). The reason for this observation is that while there does not exist an optimal solution, the

Outcomes identical across agents with equal probability

Outcomes identical across agents with equal probability

Outcomes identical across agents with equal probability

Fig. 4. Results for identical fair outcomes (scenario 3).

attempts by agents to find the best strategy will lead to the following consequences: As the interactions increase, the agents become ever more homogenous, thus relying on more and more similar information. This arises from the fact that agents will have interactions with an ever-increasing number of agents that are common to other agents, too. If now, by pure coincidence arising from the random outcomes, one strategy has performed better in the past, then an increasing number of agents will choose this strategy, and the agents will tend to choose this strategy. Given that they rely on the same information, they will tend to make the same decision; thus, we obtain herding. If the memory becomes long, any such performance remains in the

Outcomes identical across agents with Prob(buy→success)=0.75

Outcomes identical across agents with Prob(buy→success)=0.75

Outcomes identical across agents with Prob(buy→success)=0.75

Outcomes correlated as from the adjacency matrix

Outcomes correlated as from the adjacency matrix

Outcomes correlated as from the adjacency matrix

Fig. 5. Results for identical biased outcomes (scenario 4).

system much longer, making decisions of agents identical over time. This is caused by the fact that new information is only slowly replacing old information, making any changes to the behavior of agents slow. This process increases the length of herding and is clearly visible in the case of independent fair outcomes (scenario 1) and, to a lesser degree, also for correlated fair outcomes (scenario 5).

In the case of identical fair outcomes (scenario 3), we do not observe any herding. Here, the reason is that the common outcome does not allow for any strategy to obtain better performance, thus not initializing any herding. The correlated fair outcomes (scenario 5) are an intermediate stage between identical and independent fair outcomes and thus show less herding than for independent outcomes, as would be expected.

Finally, for minority outcomes (scenario 6), where an optimal strategy also does not exist, we observe no herding. This can be explained by the fact that the optimal strategy is endogenously determined, and any advantage of a strategy will be eroded once herding emerges, thus reversing the decisions of anyone following it. Therefore, we do

Fig. 6. Results for correlated fair outcomes (scenario 5).

not observe any herding, and long memory can actually increase the coordination as agents tend to choose the same strategy over time, thus increasing the number of agents being in the minority. The smaller advantage of a strategy over its counterpart will avoid large swings between agents' choices.

C. Main Findings

Based on the aforementioned analysis of our computer experiments, we can now derive some general conclusions. We could first establish that with the variant of the CBDT used for communications networks,

Fig. 7. Results for minority outcomes (scenario 6).

learning is driven by the interactions between agents, with memory only accelerating this process but not being able to initiate learning. However, for any learning to be able to take place, the experience of other agents must convey some additional information, i.e., not merely be equivalent to the agent's own experience and be therefore sufficiently independent. This finding is very similar to the requirements in statistical analysis, where a sufficient number (high number of interactions) of independent (low correlation) observations is required to be able to derive sound results.

Fig. 8. Results for majority outcomes (scenario 7).

This learning process is clearly defective, and in the absence of any superior strategy, we can, depending on the structure of the economy and the determination of outcomes, observe a significant degree of prolonged herding that arises from the attempt to learn the best strategy. This herding can lead to decisions biased toward one or the other decision, thereby causing significant misallocations in the economy. We may also observe the emergence of firmly held beliefs in "what is the right way to do things," which vanish as fast as they appear, causing also significant distortions in the economy.

The origin of this result can be found in the aim of the agents. In traditional models of economic decision making, agents attempt to learn the structure of the economy and then use this information to determine their optimal strategy. With our variant of the CBDT for communications networks, in contrast, agents attempt to determine the best strategy without learning the structure of the economy. In the absence of a best strategy, the best strategy being too complex to easily be identified, or the information processed not being sufficient to identify the best strategy, we will observe either no learning or some complex dynamics resulting from the (fruitless) search process.

It has to be noticed that any herding emerging from this process is the result of an attempt to learn the best decision and not arising from a simple model of imitation as is commonly assumed [23]–[26]. While in these models imitation of the decision of others is merely the result of a desire by agents to follow the majority, in our model, this forms an essential part of the attempt to learn the best decision. Agents in our model only make the same decisions as their neighbors because the performance of their neighbors' past decisions suggested that it was superior. Thus, the same decisions are not driven by the desire to follow the same decision as others but are the result of an attempt to use all available information on the best decision. Thus, imitation of decisions is merely the result of this desire to find the optimal decision, not the end of the agents.

VI. CONCLUSION

We have evaluated a model in which agents make decisions with a variant of the CBDT using their own past experience and the experience of other agents they are interacting with. As we evaluated a number of scenarios in which different rules were applied to determine the successful outcome, we established that for learning of the best decision to occur, agents must have a sufficient large number of interactions, and the information an agent receives from his interactions must be sufficiently independent. If these conditions are not met, no learning will be observed. In the absence of a best decision, the attempt to learn this nonexistent best decision can give rise to widespread and prolonged herding. This herding is particularly prominent and long lasting in cases where there are many interactions and the memory of individuals for past experiences is long.

While we are able to demonstrate some conditions under which learning can take place and evaluate the conditions under which herding will emerge, future research will have to provide more detailed explorations of the role the interactions are playing in this process by, e.g., investigating the role of the network structure underlying the interactions, particularly in light of the results in [17], which shows that the network structure affects learning. Furthermore, the location of the phase transition and the influence of the memory on this point need to be much more precisely determined. It would finally be of interest to investigate the properties of scenarios producing herding to see whether they generate realistic time series.

Finally, our model does not allow learning from hypothetical cases, e.g., if a "buy" decision did generate an outcome of "failure," the decision maker does not conclude that a "sell" decision would have generated an outcome of "success" and uses this to adjust the value function. The use of such hypothetical cases would change the value function and, thus, might have an impact on the outcomes reported here. The results in [13] and [14] suggest that with additional information, e.g., on the demand structure, learning can be achieved. Our results suggest that such additional information seems to be critical for learning to occur. Whether the use of hypothetical cases, which would allow generating relative frequencies of successes and failures as used in [14], would have a similar effect on learning is an aspect of the model that could be addressed in future research.

ACKNOWLEDGMENT

The author would like to thank the anonymous referees and the Associate Editor, whose comments helped improve the quality of this correspondence.

REFERENCES

- [1] I. Gilboa and D. Schmeidler, "Case-based decision theory," *Q. J. Econ.*, vol. 110, no. 3, pp. 605–639, Aug. 1995.
- [2] I. Gilboa and D. Schmeidler, "Case-based optimization," *Games Econom. Behav.*, vol. 15, no. 1, pp. 1–26, Jul. 1996.
- [3] I. Gilboa and D. Schmeidler, "Act similarity in case-based decision theory," *Econ. Theory*, vol. 9, no. 1, pp. 46–61, Feb. 1997.
- [4] I. Gilboa and D. Schmeidler, "Cumulative utility consumer theory," *Int. Econ. Rev.*, vol. 38, no. 4, pp. 737–761, Nov. 1997.
- [5] I. Gilboa and D. Schmeidler, "An overview of case-based decision theory," in *Uncertain Decisions: Bridging Theory and Experiments*, L. Luini, Ed. Norwell, MA: Kluwer, 1999, pp. 215–235.
- [6] I. Gilboa and D. Schmeidler, "Case-based knowledge and induction," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 30, no. 2, pp. 85–95, Mar. 2000.
- [7] I. Gilboa and D. Schmeidler, *A Theory of Case-Based Decisions*. Cambridge, U.K.: Cambridge Univ. Press, 2001.
- [8] I. Gilboa and D. Schmeidler, "Inductive inference: An axiomatic approach," *Econometrica*, vol. 71, no. 1, pp. 1–26, 2003.
- [9] A. Guerdjikova, *Case-Based Decision Theory—An Application for the Financial Markets*. Heidelberg, Germany: Univ. Heidelberg, 2002. mimeo.
- [10] A. Guerdjikova, *Asset Pricing in an Overlapping Generations Model With Case-Based Decision Markers*. Heidelberg, Germany: Univ. Heidelberg, 2003. mimeo.
- [11] A. Guerdjikova, *On the Definition and Existence of an Equilibrium in a Financial Market With Case-Based Decision Makers*. Heidelberg, Germany: Univ. Heidelberg, 2003. mimeo.
- [12] I. Gilboa and N. Vieille, "Majority vote following a debate," *Soc. Choice Welf.*, vol. 23, no. 1, pp. 115–125, Aug. 2004.
- [13] H. Jahnke, A. Chwolka, and D. Simons, "Coordinating demand and capacity by adaptive decision making—An application of the case-based decision theory," *Decis. Sci.*, vol. 36, no. 1, pp. 1–32, Feb. 2005.
- [14] I. Gilboa and A. Pazgal, "Cumulative discrete choice," *Mark. Lett.*, vol. 12, no. 2, pp. 119–130, May 2001.
- [15] R. Meyer, T. Erdem, I. Gilboa, W. Hutchinson, S. Lippman, C. Mela, A. Pazgal, D. Prelec, and J. Steckel, "Dynamic influences on individual choice behavior," *Mark. Lett.*, vol. 8, no. 3, pp. 349–360, Jul. 1997.
- [16] A. Krause, "Herding without following the herd: The dynamics of case-based decisions with local interactions," in *Economics and Heterogeneous Interacting Agents*, A. Namatame, T. Kaizouji, and Y. Aruka, Eds. Berlin, Germany: Springer-Verlag, 2006, pp. 179–191.
- [17] M. Blonski, "Social learning with case-based decisions," *J. Econ. Behav. Organ.*, vol. 38, no. 1, pp. 59–77, Jan. 1999.
- [18] A. Aamodt and E. Plaza, "Case-based reasoning: Foundational issues, methodological variations, and system approaches," *AI Commun.*, vol. 7, no. 1, pp. 39–59, Mar. 1994.
- [19] Y. Guan, X. Wang, and Q. Wang, "A new measurement of systematic similarity," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 38, no. 4, pp. 743–758, Jul. 2008.
- [20] K. V. Katsikopoulos and B. Fasolo, "New tools for decision analysts," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 36, no. 5, pp. 960–967, Sep. 2006.
- [21] H. Tu, J. Allanach, S. Singh, K. R. Pattipati, and P. Willett, "Information integration via hierarchical and hybrid Bayesian networks," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 36, no. 1, pp. 19–33, Jan. 2006.
- [22] D. Challet, M. Marsili, and Y.-C. Zhang, *Minority Games: Interacting Agents in Financial Markets*. London, U.K.: Oxford Univ. Press, 2004.
- [23] A. V. Banerjee, "A simple model of herd behavior," *Q. J. Econ.*, vol. 107, no. 3, pp. 797–818, Aug. 1992.
- [24] S. Bikhchandani, D. Hirshleifer, and I. Welch, "Learning from the behavior of others: Conformity, fads, and informational cascades," *J. Econ. Perspect.*, vol. 12, no. 3, pp. 151–170, 1998.
- [25] A. Devenow and I. Welch, "Rational herding in financial economics," *Eur. Econ. Rev.*, vol. 40, no. 3–5, pp. 603–615, Apr. 1996.
- [26] A. Krause, "Herding behavior of financial analysts: A model of self-organized criticality," in *The Complex Dynamics of Economic Interaction: Essays in Economics and Econophysics*, M. Galegatti, A. Kirman, and M. Marsili, Eds. Berlin, Germany: Springer-Verlag, 2004, pp. 257–268.