Trade-Offs Between Information and Crowding in Sequential Decisions (Student Abstract)

Daniel Hall, Jon Kleinberg
Cornell University
dhall9734@gmail.com, kleinber@cs.cornell.edu

Abstract

A rich line of theoretical work has modeled scenarios in which a set of agents make decisions sequentially, based on observing a growing mix of public and private signals that are revealed as these decisions occur. Here, we study a second crucial dimension, which is the way in which strategies can depend on crowding. In particular, consider a setting in which agents must sequentially decide which of several options to invest in, each based on a public signal that they receive. One of these options will ultimately be revealed to be valuable; but crucially, all the agents who selected this option must divide the value that comes from it. As a result, when a given agent j goes to make a decision among the options, the decisions of earlier agents convey information about the payoff that j will receive in any eventual division of the value. When many earlier agents have chosen a specific option, the greater crowding on this option means it must be divided more finely, resulting in lower payoffs. To simulate large games when signals are public, we define a polynomial-time algorithm to compute equilibrium strategies. We show that even in this case of public signals, the interaction of crowding with informational effects leads to complex non-monotonicities in the resulting sequential decisions, with agents sometimes choosing options with lower expected levels of crowding — and hence a better split of the potential value — over options with better informational or current crowding properties.

Introduction

A long line of work in economic theory has studied models of decision-making in which a sequence of individuals evaluate a set of options over time, with information about the quality of the options emerging gradually and incrementally over the course of this sequence. This basic structure captures settings such as investment, in which potential investors in a new company encounter it sequentially over time, and information about the company's viability emerges over the same time scale; or a new product or technology, where potential consumers must make decisions about adopting it over time, again with a growing amount of available information.

Such settings have often been modeled formally via the following general framework (Banerjee 1992; Bikhchandani, Hirshleifer, and Welch 1992; Acemoglu et al. 2008): there is a set of *options* being considered; a sequence of agents must decide, one by one, which option to accept; and each agent receives a *signal* from the environment that provides probabilistic information about which option is better. We imagine that there is a "correct" or "best" option, and each agent receives a payoff if they choose this best option.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Crowding Effects Thus, when people make decisions sequentially, the accumulation of information via signals is a key effect that agents take into account. But in many settings where these questions arise, there is a second key effect that modifies payoffs: the possibility of *crowding*. In particular, if the option represents an opportunity that is potentially valuable for the agents involved, there are many cases where the opportunity is more valuable when fewer people are making use of it. As a fundamental example of this that shows up across many domains, an opportunity yielding a fixed total payoff that must be split among the people taking part in it will be more valuable if it must be split with fewer people.

We show that in many cases, the cost of crowding can discourage agents from choosing an otherwise attractive option. Surprisingly, the crowding does not need to have occurred prior to the agent's turn. An agent could predict that an option would likely increase rapidly in value to later agents and become overly popular, while another option would increase in value and therefore popularity less quickly, causing the agent to choose what is currently the worse option.

Our goal in this paper is to develop a model of these types of crowding effects in sequential decisions, and to study agent behavior in this setting. We will show that the interaction of information effects and crowding effects produces a complex set of phenomena: as agents make decisions sequentially, they must balance in subtle ways the available information (which leads them to align with what others are doing) and the loss of payoffs through crowding (which leads them to want to differentiate from others).

A Model Combining Information and Crowding In our model, an urn is chosen with equal probability to be mostly blue (MB) or mostly red (MR), and contains either a p fraction blue marbles and a $q = 1 - p$ fraction red marbles or vice versa, respectively. Each of the n players sequentially draws a marble (with replacement), observes the previously drawn marbles and guesses, and guesses that the urn is either mostly blue or mostly red. After the last player has made

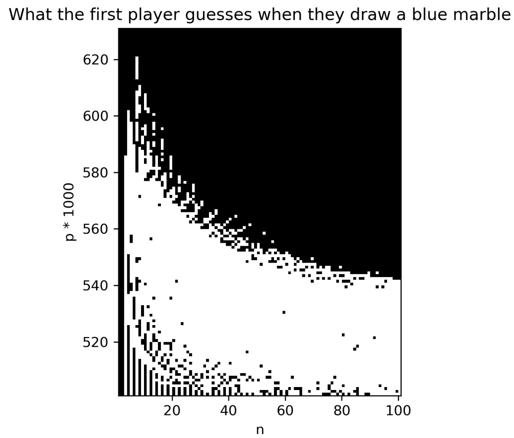

Figure 1: For games with $p \in [0.501, 0.502, \ldots, 0.629, 0.63]$ and $n \in [1, 100]$, the first player's optimal guess when they draw a blue marble is represented as a black square if MB or a white square if MR. The first player would always guess MB when $p \in [0.631, 0.632, \ldots, 0.998, 0.999]$ and $n \in [1, 100]$.

their guess, the true state of the urn is revealed and the players who guessed correctly receive a payoff of n divided by the total number of correct guesses, and those who guessed incorrectly receive nothing.

Optimal Strategies

It's reasonable to assume that the optimal strategy would be to simply guess the most frequently drawn color, and to guess the less frequently guessed color when equal numbers of blue and red marbles were drawn. It's also reasonable to assume that the optimal strategy is to guess as if the rewards will be given immediately afterwards. However, neither of these strategies are consistent with the optimal strategies computed with our dynamic programming algorithm.

Sequential decision problems with two completely symmetric options typically have the (essentially trivial) property that the first agent will decide in the direction of the signal they receive: given that the two options are identical under relabeling, and given that the first agent only has their signal, the argument for this is usually immediate. However, suppose $n = 3, p < 2 - \sqrt{2}$, and the first player draws a blue marble. According to this reasoning, the first player should guess MB. However, the optimal strategy is to guess MR. The reason for this is that p is close enough to $\frac{1}{2}$ that the cost of crowding is larger than the benefit of guessing the more likely option. We can see from simulations that this phenomenon is not restricted to the $n = 3$ case. Figure 1 shows the optimal guess for the first player for various values of n and p.

As we have just seen, increasing the number of blue marbles drawn can make the first player no longer want to guess MB. However, increasing the number of previous players who have guessed MB will never cause a player to become more likely to guess MB. More formally, we have:

Theorem 1

The probability of guessing MB is a non-increasing function of the number of previous MB guesses.

$$\forall j \in [2, n], u \in [0, j], x \in [0, j-2] :$$
$$P\left(\text{player j guesses MB}; U_j = u, X_{j-1} = x\right)$$
$$\geq P\left(\text{player j guesses MB}; U_j = u, X_{j-1} = x + 1\right),$$

where X_j is the number of MB guesses and U_j is the number of blue marbles drawn after the j-th player.

While the crowding mechanism can cause players to guess what is less likely to be correct, we can prove that once $O(\log n)$ guesses have been made, players will guess correctly with high probability.

Theorem 2

Let

$$J(n) = \beta \frac{\log_{\frac{p}{q}}(n)}{(p-q)},$$

where $\beta > 1$ is any real constant. Then $\forall j \geq J(n)$:

$$\lim_{n \to \infty} P(\text{player j guesses correctly}) = 1.$$

Further, we can prove that the total number of incorrect guesses after $O(\log n)$ guesses have been made goes to 0 in expectation as n becomes large.

Conclusion

In this work, we have considered a model of sequential decisions in which an agent must balance between the *informational effects* of prior agents' decisions and the *crowding effects* of having to share payoffs with them. We have shown that the optimal solution exhibits subtle trade-offs between the different effects, including cases where agents disregard their information in anticipation of future crowding.

There are a number of interesting questions left open by this work. First, the complex behavior particularly of the first few agents in the process is something that our algorithm can determine optimally, and where we have given some basic analytical results, but it would be interesting to try to more fully characterize the alternation in decisions that these agents exhibit.

Acknowledgments

The authors thank David Easley for his advice on research directions.

References

Acemoglu, D.; Dahleh, M. A.; Lobel, I.; and Ozdaglar, A. 2008. Bayesian Learning in Social Networks. Technical Report 2780, MIT Laboratory for Information and Decision Systems (LIDS).

Banerjee, A. 1992. A simple model of herd behavior. 107: 797–817.

Bikhchandani, S.; Hirshleifer, D.; and Welch, I. 1992. A theory of fads, fashion, custom and cultural change as information cascades. 100: 992–1026.