Social Learning from Online Reviews with Product Choice

Stefano Vaccari
Dipartimento di Economia e Finanza,
LUISS - "Guido Carli" University
Rome, Italy
svaccari@luiss.it

Costis Maglaras
Columbia Business School, Columbia University
New York, USA
c.maglaras@gsb.columbia.edu

Marco Scarsini
Dipartimento di Economia e Finanza,
LUISS - "Guido Carli" University
Rome, Italy
marco.scarsini@luiss.it

ACM Reference Format:
Stefano Vaccari, Costis Maglaras, and Marco Scarsini. 2018. Social Learning from Online Reviews with Product Choice. In *NetEcon '18: Economics of Networks, Systems and Computation, June 18, 2018, Irvine, CA, USA.* ACM, New York, NY, USA, 1 page. https://doi.org/10.1145/3230654.3230662

It is common for consumers to rely on online reviews when they plan to purchase a new product or a new service. In fact, consumers, who may not be accurately informed about the product attributes, often refer to information aggregator websites that compile product reviews, such as TripAdvisor, Yelp, etc., to quickly gather information and estimate the quality of each alternative. As consumers observe and report reviews over time, they engage in what is called a social learning process. Namely, consumers share their experiences about the different products with other consumers, who, in turn, use this information to estimate the quality of the different alternatives to make better-informed purchase decisions.

However, given that users in such online marketplaces typically choose among many alternatives, what matters for the selection is not only the amount of information about a particular product but also the amount and level of information about its close substitutes. As a result, learning transients are correlated across products: information accumulates faster for products that consumers perceive as more appealing, as these products attract most of the observations. This paper investigates the interplay between the information dispersion effects induced by the presence of product choice and the learning dynamics of a population of consumers who communicate through the mechanism of online reviews.

In a bit more detail, we study a model of a marketplace where consumers arrive sequentially and decide whether to buy one of the available products or to take an outside option. Consumers have heterogeneous preferences towards the observable features of the products, and, although initially uninformed about the quality of the different products, they estimate it by observing the binary online reviews reported by other purchasers prior to their arrival. Given their estimates, their prior quality beliefs, their idiosyncratic preferences, and the prices, consumers choose the product that maximizes their expected utility. Purchasers of a given product experience its true quality plus a small random perturbation, and report a "like" review if their ex-post utility exceeds the utility they expected before the purchase, and a "dislike" otherwise. The consumer learning model is naïve, i.e., consumers take the ratio of likes over the total number of reviews as a proxy for quality: if 80% of the reviews are positive, then consumers perceive that the value of the quality is 80% of the maximum possible quality.

We analyze the learning process of consumers and we establish that, even if perfect learning fails (consumers' quality estimates do converge almost surely, albeit not to the correct quality vector), (a) consumers do learn to rank products correctly according to their intrinsic quality, and (b) the bias that the naïve rule induces in the asymptotic quality estimates can be made arbitrarily small under some mild conditions that we provide.

Then, to characterize the impact of product choice on the quality of consumers' decisions, we study the dynamics of the learning paths in a deterministic large market (fluid) approximation, and we provide tight bounds for the rate of information accumulation for the different products. In particular, we prove that the time-to-learn is (roughly) inversely proportional to the number of available options, and further quantify its dependence on the dispersion of product prices, the prior quality estimates, and the distribution of true product qualities.

Finally, we address the placement problem of the platform that displays the online reviews. We assume that consumers perceive an additional "search cost," which is an increasing function of the ranking in which the product is displayed by the platform. The presence of these search costs, allows the platform to influence consumer choice and learning transients, by choosing the ordering in which the products are displayed to arriving consumers.

First, we show that even if the platform has no influence on what consumers learn in the long run, she can have a dramatic impact on the speed of learning. For instance, if search costs are linear and if the platform adopts a static ranking policy that displays products in alphabetic order, the learning process for bottom-ranked products will experience a slowdown that is exponential in the number of products. Then, we motivate two variants of a platform control problem where the platform focuses on maximizing discounted cumulative revenue on one hand, and on minimizing the time-to-learn on the other. While finding the optimal dynamic ranking policy is intractable, we study a family of heuristics that either greedily maximize next period revenue (exploit), or maximize the speed of the next period information acquisition (explore), or some variants of these two. We resort to a set of numerical simulations to demonstrate that ranking policies that allow for a suitable amount of exploration may lead to higher discounted profits compared to a full-exploit policy. And, perhaps surprisingly, in some cases the platform can achieve a faster overall time-to-learn than even with respect to a system where search costs are absent. The extended version of the paper can be found at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3101747.

NetEcon '18, June 18, 2018, Irvine, CA, USA

ACM ISBN 978-1-4503-5916-0/18/06.
https://doi.org/10.1145/3230654.3230662