

Learning in the Repeated Secretary Problem

DANIEL G. GOLDSTEIN, Microsoft Research
R. PRESTON MCAFEE, Microsoft Corporation
SIDDHARTH SURI, Microsoft Research
JAMES R. WRIGHT, Microsoft Research

In the classical secretary problem, one attempts to find the maximum of an unknown and unlearnable distribution through sequential search. In many real-world searches, however, distributions are not entirely unknown and can be learned through experience. To investigate learning in such a repeated secretary problem we conduct a large-scale behavioral experiment in which people search repeatedly from fixed distributions. In contrast to prior investigations that find no evidence for learning in the classical scenario, in the repeated setting we observe substantial learning resulting in near-optimal stopping behavior. We conduct a Bayesian comparison of multiple behavioral models which shows that participants' behavior is best described by a class of threshold-based models that contains the theoretically optimal strategy. In fact, fitting such a threshold-based model to data reveals players' estimated thresholds to be surprisingly close to the optimal thresholds after only a small number of games.

For the full version of this paper see: http://www.dangoldstein.com/papers/Goldstein_McAfee_Suri_Wright_Learning_Repeated_Secretary.pdf.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). *EC'17, June 26–30, 2017, Cambridge, MA, USA.* ACM ISBN 978-1-4503-4527-9/17/06.

© 2017 Copyright held by the owner/author(s).
DOI: <http://dx.doi.org/10.1145/3033274.3085112>