Segmenting Retail Advisors and Optimizing Coverage Model

William McEntee[A]; Chinmay Jha[A]; Dimitris Bertsimas[B]; Ryan Cory – Wright[C]; Nadine Kawkabani[D]; Brian Shaw[D]; Brendan Mannix[D]

[A]MIT MBAn 2018; [B]Faculty Advisor, MIT; [C]PhD Student – mentor, MIT; [D]Mentor, MFS Investment Management, Boston, US

Abstract

Clients of Massachusetts Financial Services' (MFS) US Retail business include 300,000 financial advisors spread across the US. With a salesforce of 150 representatives, MFS can only service 7.5% of all the financial advisors effectively.

First, we explore how accurately we can predict the transactions from a financial advisor across various MFS funds in the next six months. Second, we address the problem of optimal resource allocation using an optimization framework to prescribe interaction levels for every advisor using the predictive model.

Third, we identify new approaches for MFS to grow its business by identifying new funds to recommend to advisors. Finally, we propose an extension to the slice recovery algorithm to recommend funds to new advisors.

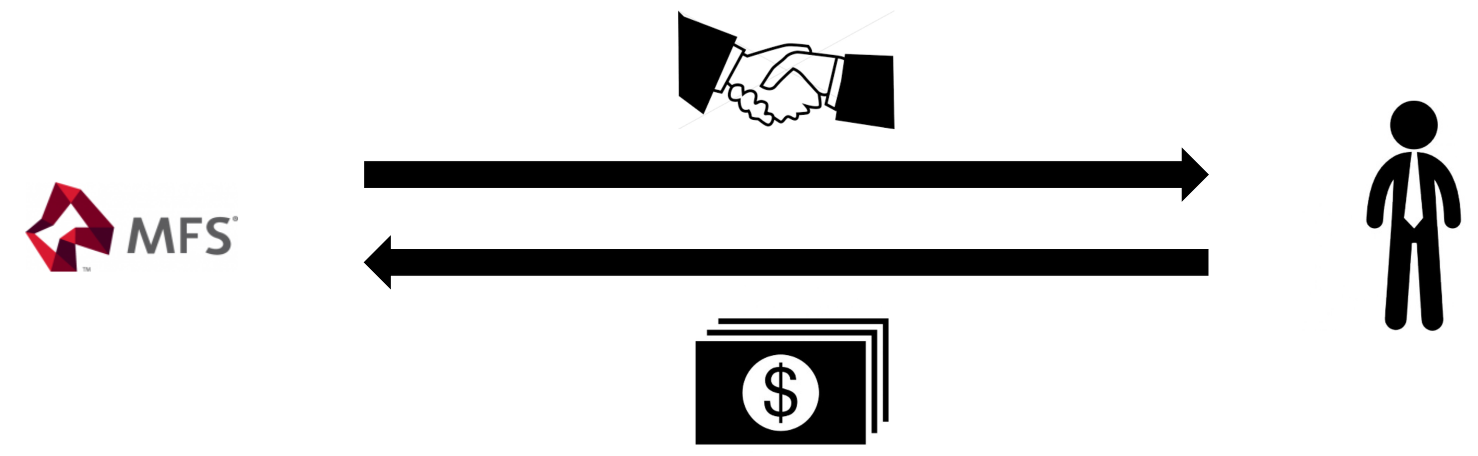

Market Share & Performance	Activities	Transactions	Advisor Information
MFS's market share across **~13000** client offices	**~1 million** activities(emails, meetings, phone calls) from 2013-16	**~15 million** transactions across MFS funds between 2014-17	Information on **~19000** advisors

- January – *Matched with MFS*
- February – *Developed a deeper understanding of the problem*
- March – *Got access to data and initial exploratory analysis*
- April – *Continued exploratory analysis; First meeting with MFS senior management to set project scope*
- May – *Data preprocessing and preliminary predictive modeling*
- June – *Incorporated feedback on models; Presented preliminary results to MFS President's Council*
- July – *Expanded modeling to prescription and recommendation*
- August – *Delivered final results; Presentation at MIT*

Prediction

How accurately can we predict flows from advisors in the next six months?

- Data: transactions, advisor-specific information, activities, and fund performance
- Methods: regression trees, boosted trees, optimal trees, and classify-then-predict
- Evaluation metric: R^2, mean absolute error (MAE) compared against mean absolute deviation (MAD)

Prescription

Which interactions should we prescribe for an advisor based on the predictive model?

- Data: transactions, advisor-specific information, activities, and fund performance
- Methods: optimal trees and optimization formulation for prescriptive approach
- Evaluation metric: % lift over predicted flows

Recommendation

Which new funds should we recommend to existing advisors?

- Data: purchase history observed across time slices of six months
- Methods: slice recovery, user-based collaborative filtering, item-based collaborative filtering, and matrix factorization
- Evaluation metric: % of new funds purchased which were correctly recommended

Extrapolation

Which new funds should we recommend to new advisors?

- Data: purchase history observed across time slices of six months, advisor-specific information
- Methods: slice recovery and nearest neighbors approach
- Evaluation metric: % of new funds purchased which were correctly recommended

Optimal trees' out-of-sample R^2 is at par with boosted trees

Prescription approach gives lifts over the predicted flows

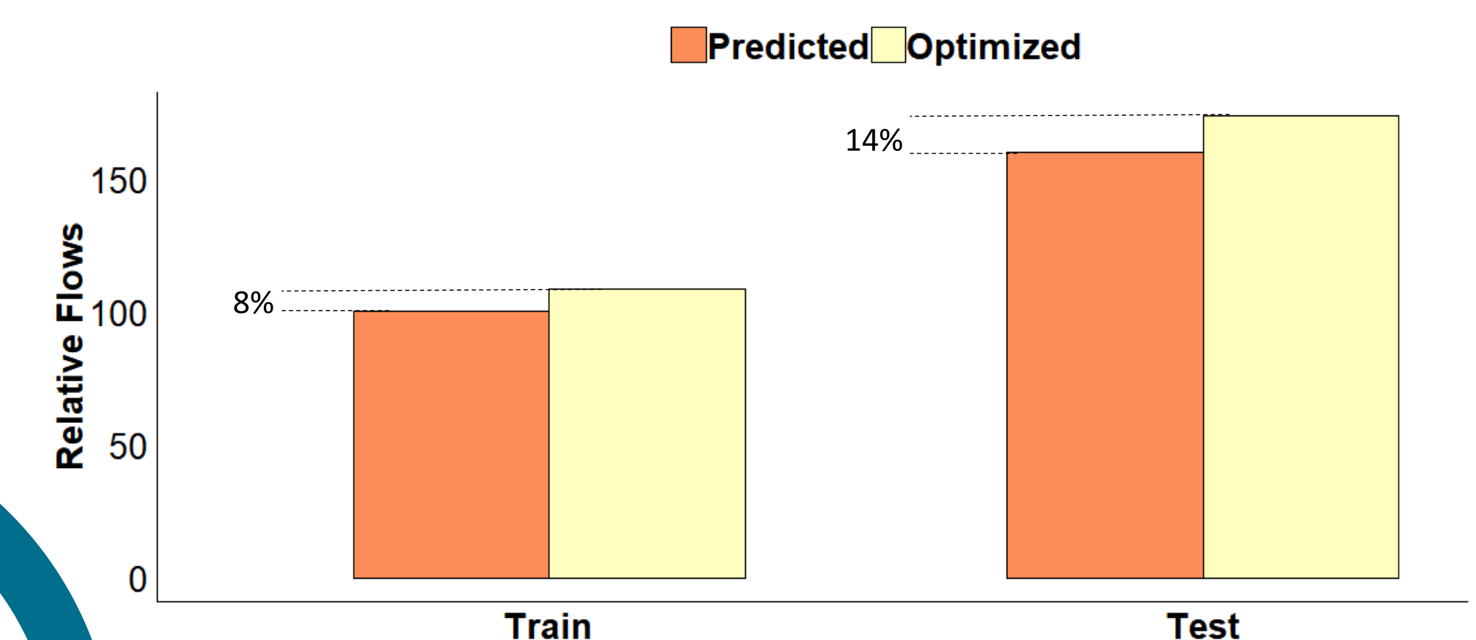

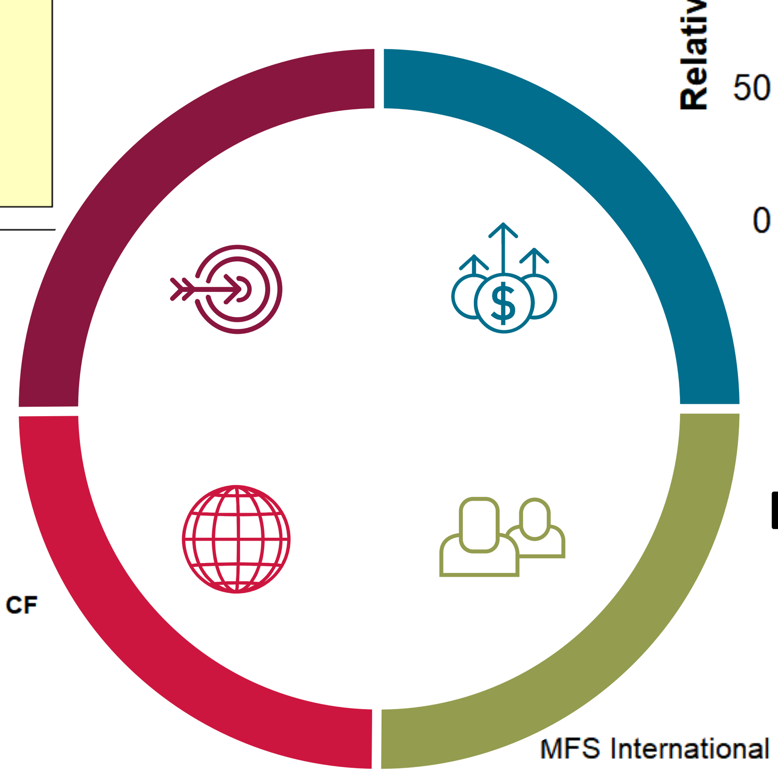

Slice recovery beats incumbent approaches

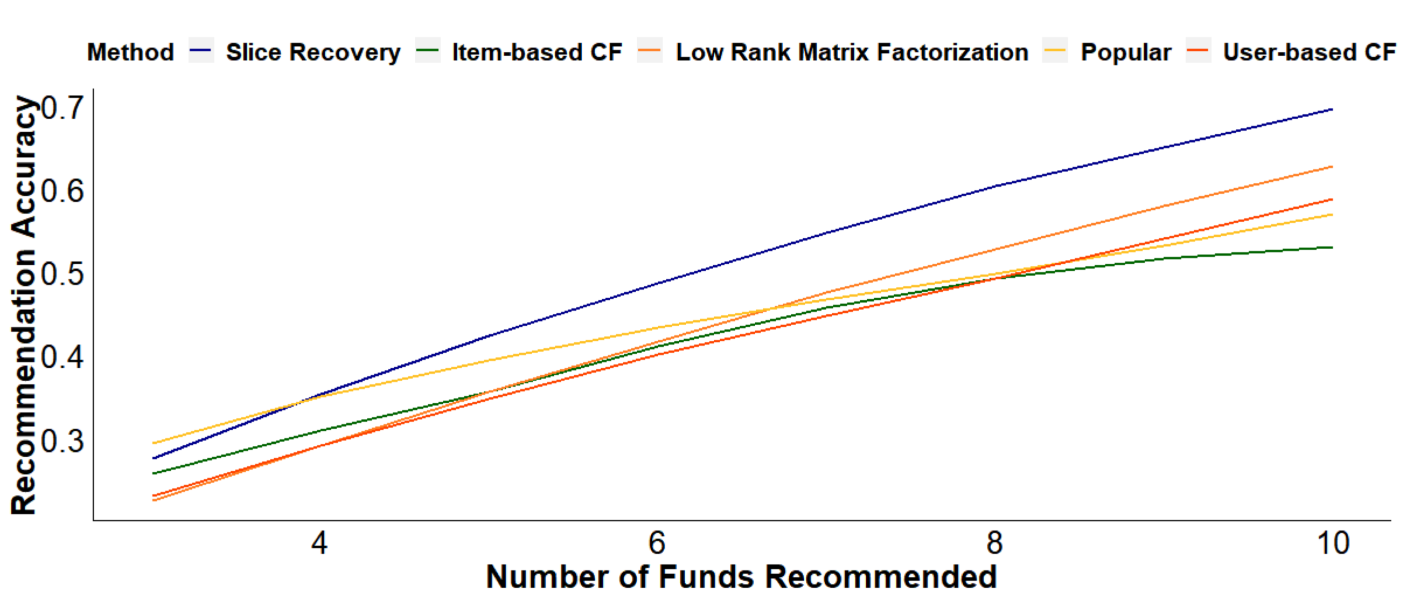

Extrapolation has consistent performance across funds

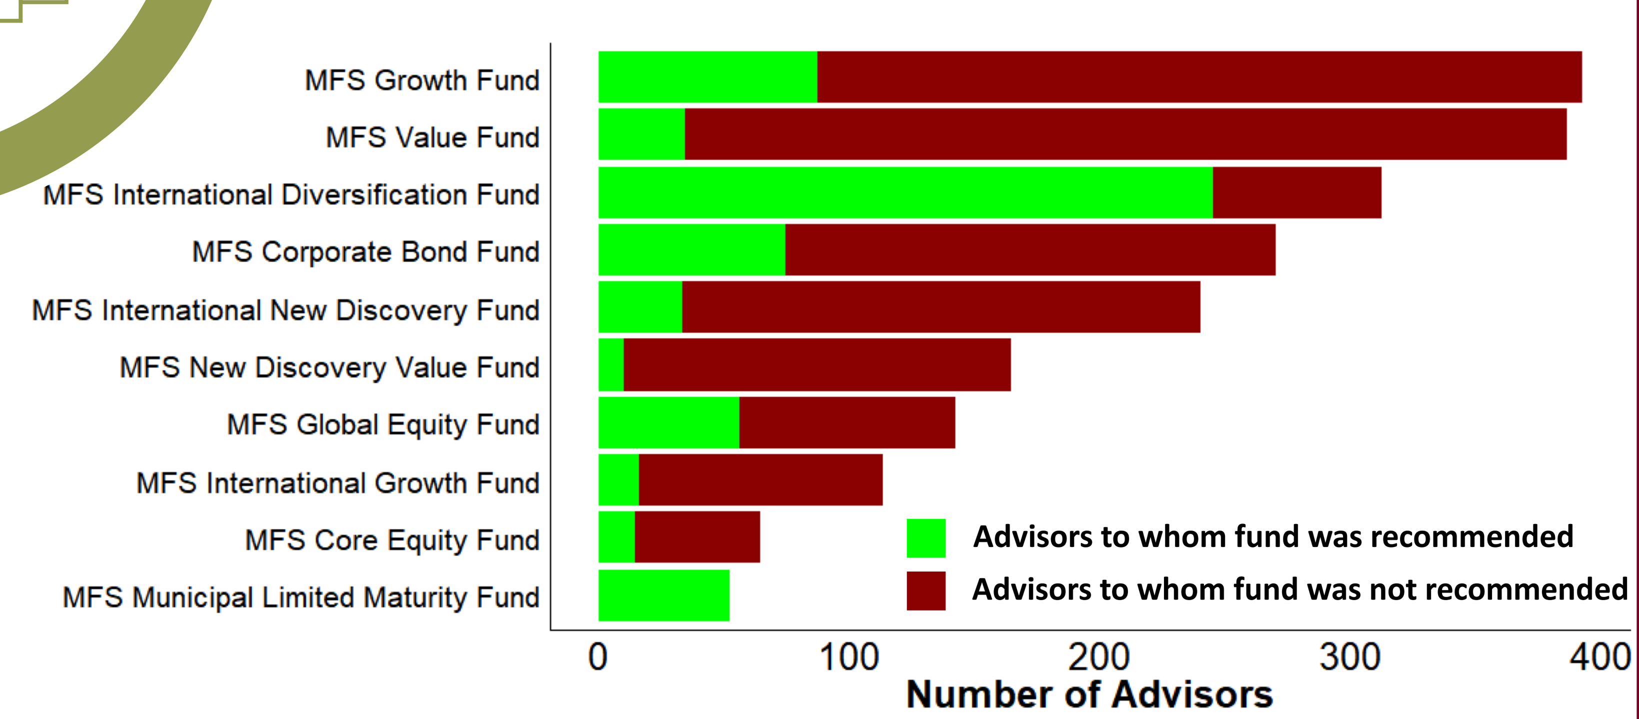

Classified 80% advisors correctly as high or low-value

8% lift over predicted flow levels

Recommended 70% of new funds purchased

30% recommendation accuracy for new advisors

1. D. Bertsimas, J. Dunn. **"Optimal Classification Trees"**. *Mach. Learn., 2017*
2. A. Li, V. Farias. **"Learning Preferences with Side Information"**. *Mang. Sci. (to appear)*
3. D. Bertsimas, N. Kallus. **"From Predictive to Prescriptive Analytics"**. Mang. Sci. (under review)