MARKETING COLLOQUIA
Fall 2026 - Abstract

Kevin Lee
Assistant Professor of Marketing
University of Michigan, Ross School of Business

Thursday, October 15, 2026
Where: 741 JMHH
When: 12:00 PM to 1:20 PM

IDEAS WORTH TESTING: GENERATIVE AI AND THE RETURNS TO EXPERIMENTATION

Predicting which version of marketing content (promotional content, email copy, push notifications, etc.) will perform best is difficult. Randomized experiments can resolve this uncertainty but are used sparingly in practice, as the gains from testing the available alternatives may not justify the potential loss of engagement from showing customers poorly performing content. We propose using generative AI to increase the returns to experimentation by creating better content to test. We use a predictive model not to choose the best content directly, but rather as a guide for generative AI to produce improved variants of human-authored content. Across three large field experiments, we find substantial improvements in the upper tail of the performance distribution without deterioration in the lower tail. The best AI-generated content outperforms the best human-authored content by 36 to 47 percent, while average performance rises modestly by 5 percent. Despite the predictive model's modest accuracy--it cannot reliably identify the high performers in advance--it helps generate a set of candidates that contains high performers. These hard-to-predict gains raise the experimentation premium, the within-campaign gain from testing the candidates rather than choosing one before launch. In simulations calibrated to the main experiment, an adaptive design that shifts traffic toward better-performing variants achieves engagement 33 percent above the human-authored average. Four percentage points of this lift come from the higher average quality of the combined human and AI candidate set and can be captured without testing; the remaining 29 points are the experimentation premium. In comparison, the premium on the human-authored content alone is 1 point. A theoretical model shows how prediction can substitute for experimentation when used to select content, yet complement it when used to guide generation. More broadly, holding the predictive model fixed, our approach can make experimentation a worthwhile route to better campaign performance.