Summary of Research Contributions and Research Plan - Shan Huang

My research examines how emerging technologies enable innovative marketing strategies and solutions at the intersection of marketing and information technology, connecting to the fields of information systems, operations, and computer science.

I advance the marketing literature both empirically—through large-scale digital field experiments and computational approaches—and methodologically—by developing experimentation methods for data- and AI-driven decision making. A defining principle of my research is **independence and relevance**: independence breeds originality, as reflected in my role as lead author on nearly all projects, where I chart distinctive research directions and initiate collaborations by defining the agenda; relevance grounds my work in practice, demonstrated by long-term partnerships with leading companies where my methodological innovations have been implemented in real-world decision systems.

I began by examining how technologies reshape marketing strategies through large-scale empirical studies. I have conducted and analyzed field experiments involving millions of users, generating gold-standard causal evidence on how emerging technologies—such as new social media platforms (Huang, Aral, Hu, and Brynjolfsson, 2020; Huang and Lin 2024[1]) and algorithmic recommendations (Huang and Ji 2024)—influence user engagement with advertisements and online information. Complementing these experiments, I have conducted large-scale field analyses of how information is generated and diffused in social networks from a computational social science perspective (Yu, Huang, Liu, and Tan 2024; Huang and Yu 2024; Chen, Hu, and Huang 2019).

A signature contribution of this stream was one of the first—and largest—field experiments ever conducted on a major social media platform. In close collaboration with industry practitioners, I designed and implemented the earliest large-scale experiment on WeChat in 2015, reaching over 37 million users. This project not only led to the establishment of WeChat's first A/B testing infrastructure but also formed the foundation of my doctoral thesis. It produced two articles published in top journals: Huang, Aral, Hu, and Brynjolfsson (2020) in Marketing Science and Huang and Lin (2024) in the Journal of Marketing.

The first paper provides the earliest large-scale, cross-product experimental evidence on how displaying social cues—friends' endorsements, a defining feature of social media—enhances advertising effectiveness, and shows how this effect varies across product characteristics. The second paper extends these insights by uncovering how behavioral characteristics shape advertising effectiveness, demonstrating that public responses (likes) and private responses (clicks) differentially drive advertising outcomes. Huang, Aral, Hu, and Brynjolfsson (2020) has been cited by scholars at leading institutions such as Harvard, Stanford, and MIT, ranks as the second-most cited article in Marketing Science's special issue on field experiments, and has influenced not only marketing's top journals but also broader scientific outlets such as Science and PNAS as a representative example of studying social media using designed field experiments.

Beyond experimentation, I have also analyzed massive-scale observational datasets to study how information is produced and diffused in digital platforms. This research has resulted in two top-journal articles: Chen, Hu, and Huang (2019) in the Journal of Management Information Systems and Yu, Huang, Liu, and Tan (2025) in Information Systems Research. The first paper investigates how monetary incentives shape content creation, using data from Seeking Alpha, a leading investment platform, showing that incentives increase content output and community engagement but do not enhance stock recommendation quality. The second paper analyzes 387,486 articles from WeChat Official Accounts, using large-scale natural language processing to measure eight discrete emotions embedded in content and evaluate their impact on information

[1] The full references of my papers are listed in my CV; due to space constraints, I will not reiterate them here.

cascade dynamics—including size, depth, breadth, and structural virality—while also identifying how demographics and tie strength influence online content diffusion.

Through sustained collaboration with industry practitioners, I came to recognize a persistent paradox: although experiments (A/B tests in industry settings) are widely deployed, their results often fail to meaningfully inform managerial decisions. This gap presents a distinctive opportunity for business researchers to design methods that strengthen the external validity of experiments—enhancing their generalizability and managerial relevance, and ultimately improving decision quality. Since joining HKU, I have launched a research stream dedicated to developing experimentation methods for data- and AI-driven decision making in marketing. This agenda builds on my expertise in causal inference and digital experimentation, as well as my long-term collaborations with China's technology industry. My work thus extends beyond empirical examination to methodological innovation, addressing fundamental barriers that prevent experiments from shaping practice effectively.

I began my methodological innovations by addressing a core challenge in digital experimentation: enabling short-term experimental results to inform long-term decision making. Working with my Ph.D. students, junior colleagues, and industry collaborators, I pioneered some of the first methods to systematically bridge this gap. Our approach integrates causal inference with machine learning to address two validity concerns: temporal validity, since user responses to treatments evolve over time; and population validity, since the composition of participants often shifts over time due to staggered adoption and selective attrition, threatening representativeness.

To address temporal validity, we developed a longitudinal surrogate framework that decomposes long-term effects into functions of user attributes, short-term metrics, and treatment assignments, supported by new identification assumptions, estimation strategies, and validation techniques. To address population validity, we designed a dynamic sampling framework that adapts to evolving participation, segmenting experiments into stages of differing generalizability and employing survival analysis to detect stage transitions and create stage-specific estimators. We validated these approaches not only through two long-term field experiments but also via real-world deployment in platform settings, encompassing more than 600 large-scale A/B tests. This methodological stream has produced one top-journal article (Huang, Wang, Yuan, Zhao and Zhang 2025) forthcoming in Management Science and two conference papers (Huang, Wang, Yuan, Zhao and Zhang 2025; Wang, Han, and Huang 2024) published at the ACM Conference on Economics and Computation, a leading computer science conference. Beyond academia, they have been implemented at Tencent and ByteDance, and released as open-source software to ensure broad accessibility and impact.

Looking ahead, my research continues to focus on ensuring that emerging technologies are deployed responsibly and effectively to advance marketing strategy and decision making. Building on my independent track record of pioneering experimentation methods, I am now extending this agenda to the era of artificial intelligence. A key stream of my current agenda examines how large language models (LLMs) can enhance digital experimentation and data-driven decision processes by improving interpretability, generalizability, and cost efficiency.

First, in collaboration with Tencent, my team is designing LLM-powered agents that act as automated experimental subjects, enabling causal inference at unprecedented scale and offering new ways to test product interventions before costly live rollouts. Second, working with small merchants, I am developing hybrid models that integrate LLMs with deep learning to support product selection and strategic decision making in resource-constrained environments, thereby broadening access to advanced AI tools beyond large firms. Third, in partnership with ByteDance, I am constructing digital twin frameworks that predict and interpret user responses to digital content prior to wide-scale distribution, providing managers with interpretable, pre-market diagnostics that improve both efficiency and user experience. Collectively, these projects highlight how marketing research can both shape and critically evaluate the responsible use of AI in practice.