Ke-Cheng Hsu (Clio)
Ph.D. Candidate, Department of Economics, HKUST
(852) 6840 9753 | khsu@connect.ust.hk
cliohsu1997.github.io | github.com/cliohsu1997

November 10, 2026

Professor Yiming Zhang
School of Management and Economics
The Chinese University of Hong Kong, Shenzhen
Re: Post Doctoral Fellow — Zhang Yiming Project (PR2026/049/01)

Dear Professor Zhang,

I am writing to apply for the postdoctoral position on your platform-operations team. I am a Ph.D. candidate in Economics at HKUST (expected June 2027), a native Mandarin speaker fluent in English. I work in empirical industrial organization and digital economics. My interest is how people decide under information uncertainty: a first-stage choice anticipates costly later adjustment, so the earlier decision and the later adjustment are determined together. With Professors Kohei Kawaguchi and Kei Kawai, I clean and link large bookstore point-of-sale records and estimate how sellers set prices and hold inventory when the retail price is committed in advance and demand is uncertain. In my job market paper, I build and estimate a structural model of purchase, return, and whether to buy return-shipping insurance on Taobao apparel, comparing automatic coverage of return-shipping insurance (RI) with optional coverage (CI), ranking policies by welfare, and decomposing the loss into an ex-ante purchase loss and an ex-post keep-or-return loss. The main finding is simple: a stricter return policy curbs inefficient ex-post returns, but the resulting lower return rates make higher product prices more attractive for sellers, leading to possible exclusion of ex-ante efficient purchases. A menu pairing a discounted price with a medium restocking fee as penalty often improves on both RI and CI by better balancing purchase and return efficiency.

I am strongest at structural estimation end to end. I clean large private datasets, write tests, build estimation pipelines in Python, R, and C++, run them on SLURM and AWS, and deliver figures, tables, and short reports on fixed deadlines. My GitHub (github.com/cliohsu1997) holds the code I am comfortable sharing. Every number should be recoverable from the code and the tests. I believe I can contribute to the behavioral modeling, structural estimation, experiments, and policy simulation your ride-hailing collaboration requires, and to any other project the team needs me to take on.

Your work on wait-time information (WTI) is what draws me to the group. In Yu, Zhang, and Zhou (Management Science, 2022), both the initial magnitude of WTI and its update frequency shape abandonment. Your dissertation also studies how the granularity of WTI (point estimates versus intervals) changes beliefs and waiting cost–reward ratios. On top of supporting the ongoing platform project and any other work required, I am interested in two related topics.

The first topic is the *intertemporal path* of WTI: how the *rate of change* of displayed waiting time updates beliefs. Whether a rider stays depends not only on the literal number on the screen, but also on the *information value* of the next minute in the queue, the option to see a more favorable update. Holding the actual wait fixed, that wait is like a budget for designing the signal path: on average the rider's posterior about remaining wait should still be right, but a fast drop into the one- or two-minute range then a long stay there is not the same experience as a slow drop early and a fast drop late.

The design question is how to choose the rate of change under that budget to maximize retention, with updates just strong enough when needed. Because of sunk-cost bias, building confidence early with rapid improvements and then exploiting that confidence with flatter updates can raise retention, at the cost of

not screening out impatient consumers early and of hurting trust. My guess is that the right rate depends on the rider's current prior. When the prior on remaining wait is already low, the rider mainly wants to stay, so the platform should avoid bad signals that would worsen that prior. When the prior is still long, the rider mainly wants to quit, so the platform needs a strong positive shock to reverse that prior, for example a drop from forty minutes to twenty. Empirically, I would like to recover a *required waiting-time curve*, the downward-sloping path of believed wait that a rider needs to stay, see how it differs across consumers (for example, blue-collar workers, white-collar riders with an urgent trip, and casual travelers), and then design a WTI policy that keeps riders' beliefs above it. If possible, I would also like to see the effects on firm profit and consumer surplus. Holding riders in the queue raises platform profit if long-term trust is ignored, but not screening out impatient consumers can hurt consumer surplus; balancing these two should also enter the design of the waiting-time policy.

The second topic is WTI at the *competition stage*, before the rider commits to a trip. Platforms typically show detailed waiting time after an order is placed, so rivals compete mainly on price. Offering precise WTI at the choice stage would force a later commitment: the platform must live with the information it used to attract the rider. That tension is close to the attraction-versus-persuasion trade-off in Au and Whitmeyer (Journal of Political Economy, 2023): more information attracts visits, but later the firm wants just enough information to keep the conversion.

I have attached my CV and job-market paper. Two recommendation letters will be sent separately before the November 1 deadline. Thank you for your consideration.

Sincerely,

Ke-Cheng Hsu (Clio)