ECON 339: Behavioral Economics

Take-Home Quiz 1: Behavioral Econ and AI

---

1. Large language models (e.g., ChatGPT, Claude) are useful for summarizing academic papers and helping you quickly learn about behavioral economics. GenAI should not, however, be doing the actual writing or thinking for you! AI is often incorrect, through hallucinations or errors. You are responsible for your own work, and learning, and this quiz is designed to help you understand why you should be wary of relying on ChatGPT for your classes.

   Below is a copy-pasted AI summary of a recent paper (*"Safer Driving for a Price: Evidence on Behavior and Habits in Kenyan Minibuses"* by Schönholzer, Lane, and Kelley.) I fed ChatGPT the instructions on Paper Summaries from the syllabus, and asked it to write the summary for me. Your task is to critically evaluate AI's work (using the paper provided on BlackBoard) on its response. Does it answer the questions correctly? Does it get any details incorrect? **You may respond in a bulleted list or in paragraph form. You should include at least 4 comments on what AI does right/wrong.**

   AI Summary: "Safer Driving for a Price: Evidence on Behavior and Habits in Kenyan Minibuses" (Schönholzer, Lane & Kelley)

   Research question. The paper asks whether raising the cost of unsafe driving can improve road safety in settings where traffic laws are weakly enforced. The study is set in Mombasa, Kenya, and focuses on the informal minibus sector—vehicles known locally as "matatus"—which account for a disproportionate share of passenger casualties. The central worry is not just whether an incentive works while it is active, but whether it produces genuine safety gains rather than simply shifting unsafe driving to unmonitored times, and whether any improvements last.

   Data and method. The authors run a randomized controlled trial with 203 minibus drivers, assigning 126 to a treatment group and 77 to a control group. Every study vehicle is fitted with a GPS device that logs alerts for speeding, sharp braking, over-acceleration, and sharp turning. After a calibration period used to establish each driver's personal baseline, treated drivers face a pay-for-safety scheme: for each unsafe maneuver above their baseline, the driver is charged an out-of-pocket fine, with each day's driving scored around 10 p.m. and communicated to the driver by SMS that evening. Drivers start each day with a 300 KES bonus (roughly 3 USD).

To test whether safe driving becomes habitual, the authors randomly assign treated drivers to either one or two months of the incentive and compare what happens after it ends.

Main findings. The intervention meaningfully improves driving safety, raising a composite safety index by 0.96 standard deviations. This is driven by a 13 percent reduction in speeding and a 29 percent reduction in harsh braking. Effects on other margins are modest: drivers increase their labor supply somewhat but show no meaningful change in kilometers driven, revenue, or take-home pay, so the safety gains are achieved while leaving earnings intact. Examining behavior throughout the day, the authors find substantial displacement—drivers shift a large portion of their unsafe driving to after the 10 p.m. scoring cutoff, partially offsetting the daytime safety gains.

Conclusions. The authors conclude that directly pricing unsafe driving is a viable policy tool even where enforcement is weak, and that it appears cost-effective under conservative assumptions. Notably, the safety improvements persist well after the incentive ends, especially for drivers with longer exposure, which the authors interpret as evidence of strong habit formation—suggesting that even a temporary campaign can generate durable changes in driving behavior.