ALISHER BATMANOV

alisher-batmanov.github.io | abatmanov@ucsd.edu

Research Statement

My research is primarily in experimental and behavioral economics, with a focus on studying systematic biases in decision-making and the factors that impede learning from feedback diagnostic of optimal behavior. In my primary strand of work, I use laboratory experiments to design controlled environments that isolate specific sources of suboptimal behavior, such as neglect of informative signals, failure to disregard irrelevant contingencies, and susceptibility to how information is framed. I explore how exposure to others' decisions can facilitate learning, and how the structure and presentation of tasks can make environments that require contingent reasoning easier to navigate. By exogenously varying key features of the decision environment, I study when improvements in behavior reflect deeper understanding of the underlying decision rules as opposed to superficial imitation of external cues, and I test what changes to the learning environment can foster more effective decision-making. In addition, I study miscalibrated beliefs and misperceptions related to mental health stigma among university students, and how targeted information interventions can correct them in a field setting.

1. Learning from Feedback and Social Exposure

One of the central questions in behavioral economics is why biases in decision-making do not self-correct in the long run, even when individuals have ample opportunity to learn from repeated feedback. In my job market paper, titled ***"Fragile Learning From Others"***, I examine whether exposure to others' decisions can mitigate the factors that impede learning from feedback, and characterize the nature and persistence of the resulting learning effects beyond the period of exposure. In a repeated hiring task, participants choose between two workers whose observable education signals screen ability differently across groups. Despite a clear theoretical benchmark, baseline behavior exhibits systematic neglect of the informative signal, with participants instead relying on prior group averages. I examine whether exposure to another participant's decisions can correct, or exacerbate, these misconceptions. Importantly, observed choices do not convey new payoff-relevant information; instead, they serve as an external reference that may prompt reconsideration of one's interpretation of feedback. I show that exposure to optimal decisions substantially improves performance relative to a control group that learns independently, whereas exposure to suboptimal decisions leaves performance statistically indistinguishable from the control group. The asymmetry suggests that learning is not driven by imitation alone, but by the informational content embedded in observed outcomes. Crucially, I document substantial heterogeneity in learning. Roughly one-third of participants mechanically replicate observed choices and revert once exposure ends; another third maintain improvements only while the environment primitives remain unchanged; and the remaining third exhibit deeper learning that transfers to a modified environment. Brief written advice explaining *why* the optimal rule works subsequently doubles optimality among participants whom observation left unmoved. These findings highlight the dual nature of social learning from exposure: it can accelerate improvement, but it can also generate fragile learning that fails to generalize.

My ongoing joint work, ***"Learning to Ignore Irrelevant Contingencies: An Experiment"*** (with Jeongbin Kim and Emanuel Vespa), complements and extends this agenda. Whereas my job market paper focuses on neglect of informative signals, this project examines the opposite problem: failure to disregard payoff-irrelevant contingencies. Participants face an inference task in which optimal behavior requires conditioning only on contingencies that affect outcomes. In a static decision environment, participants must reason through potential contingencies without observing which one they are in. We then introduce a dynamic learning stage in which participants observe the realized contingency and receive

feedback across many rounds. Our main treatment variation concerns participants' experience during this learning stage: in one condition, irrelevant contingencies are excluded by design and participants only encounter the payoff-relevant event; in the other, irrelevant contingencies occur with positive probability but never affect payoffs. The endline results show that participants who never encountered irrelevant contingencies during the learning phase are significantly more likely to exhibit optimal behavior when they are back in the static environment. Importantly, this effect persists after controlling for learning in the dynamic environment and reported beliefs about the environment's primitives, suggesting that the experience of engaging with irrelevant events itself hinders the ability to disregard them. Analysis of incentivized written advice further shows that participants in the zero-probability treatment demonstrate a considerably higher understanding of the optimal approach when explaining the task to a future participant. We are currently finalizing the working paper draft.

Together, these two projects demonstrate that learning depends critically on how individuals represent the informational structure of their environment, as suboptimal behavior may persist when individuals exhibit confidence in incorrect models they adopt, or fail to identify which contingencies matter.

2. Persuasion and Information Interpretation

In many economic settings, individuals receive advice or recommendations from others before making decisions. A key question is what makes such communication persuasive: is it the sender's informational advantage, or the way the sender frames and interprets available evidence? These two channels of persuasion have largely been studied separately. Classic models of cheap-talk communication in sender-receiver frameworks emphasize settings with information asymmetry, in which a sender observes the underlying state and transmits a message to a less-informed receiver. Meanwhile, a more recent literature on narrative persuasion studies how a sender's interpretation of the same information can sway a receiver's actions, even absent any informational advantage. In the joint project with fellow PhD student Bridget Galaty titled ***"Experiment on Narratives and Information in Persuasion"***, we study how these two channels interact in a unified experimental persuasion setting. In our inference task, receivers observe a dataset of input-output pairs generated by an unknown data-generating process (DGP) and must predict the output for a new, out-of-sample observation. Before making their prediction, receivers observe a recommendation from a sender who is either better informed by observing the true DGP or equally informed by possessing the same data as the receiver, and who communicates either through a simple recommendation or a recommendation augmented with an elaborate explanation (narrative). By varying both the informational advantage and the communication format, we distinguish the effects of new information from those of new interpretation, and by varying whether the sender's and receiver's preferences are aligned or misaligned, we assess how each channel operates under strategic tension. In a working paper, we find that receivers' predictions land roughly midway between the truth and the sender's recommendation, and that both channels increase persuasive power, but information plays a much stronger role. A sender with an informational advantage moves predictions approximately 39% closer to their recommendation, whereas adding a narrative under equal information yields a smaller and statistically insignificant shift of roughly 9%. This asymmetry is stable across levels of task complexity (the number of input variables and the conditional structure of the DGP), and while aligned preferences amplify both channels, information continues to dominate.

3. Descriptive Complexity in Matching Mechanisms

My third line of research applies behavioral insights to the study of matching mechanisms. Although strategyproof matching mechanisms such as deferred acceptance (DA) and serial dictatorship (SD) guarantee that straightforward reporting is a dominant strategy, empirical evidence documents widespread

deviations from straightforward behavior. This raises a fundamental question: are deviations driven by strategic incentives, or by misunderstanding of the mechanism's informational structure? In my ongoing project ***"Descriptive Simplicity in Strategyproof Matching Mechanisms"***, I examine whether alternative descriptions of mechanisms can reduce cognitive burden and improve straightforward reporting. Rather than altering the underlying structure, I focus on how mechanisms are explained to participants. Specifically, I focus on interventions that highlight counterfactual outcomes under alternative rank-order lists, enabling participants to observe for themselves that strategizing cannot improve their payoff. By emphasizing the logic of dominance through concrete examples, the intervention aims to make strategyproofness more transparent. This project builds naturally on my other work on learning. Just as individuals may neglect informative components of their environment or engage with irrelevant contingencies, participants in matching mechanisms may misinterpret how their reported preferences translate into outcomes. By studying how descriptive complexity affects behavior, I aim to contribute to the growing literature on complexity and cognitive constraints in market design. I am currently in the process of designing the experiment.

4. Belief Distortions and Mental Health Care Use

While my primary work relies on laboratory experiments, I also study miscalibrated beliefs and information interventions in field settings. In ***"Beliefs, Information Sharing, and Mental Health Care Use Among University Students" (Journal of Development Economics 2026)***, with fellow PhD student Ida Grigoryeva, Master's students Bruno Calderon and Roberto González, and field collaborator Alejandro Guardiola, we examine why nearly half of students in distress at a large private university in Mexico do not utilize freely available mental health services. We document a substantial treatment gap: roughly half of students exhibiting moderate to severe symptoms of depression or anxiety have not received professional support in the past 12 months, despite the availability of free on-campus counseling. Importantly, this gap is correlated with widespread misperceptions: students underestimate the prevalence and effectiveness of therapy, overestimate stigma, and three out of four incorrectly believe that students in mental distress perform worse academically. To address these distortions, we implement a facts-based information intervention that targets beliefs along three dimensions: therapy effectiveness, the profile of typical therapy users, and the relationship between distress and academic performance. The intervention produces a clear hierarchy of effects along behavioral margins. Low-cost, outward-facing behaviors respond most strongly: treated participants are more likely to share mental health resources with peers and to recommend on-campus counseling to a friend in distress. However, the intervention has more limited effects on high-cost personal decisions, such as initiating therapy or willingness to pay for private counseling. These findings suggest that factual belief corrections can shift peer-support behaviors, while changing individuals' own help-seeking may require addressing deeper barriers related to perceived social norms and stigma. This project extends my work on belief formation beyond the laboratory and demonstrates how informational structure shapes real-world behavior in a developing-country setting. I am excited to continue building on this line of work, exploring how targeted information interventions can address belief distortions that affect consequential decisions in other applied contexts.

Together, my research agenda across laboratory and field settings shows that persistent suboptimal behavior often reflects a misunderstanding of the environment's informational structure rather than a failure of incentives, and that these misunderstandings are systematic, diagnosable, and correctable. Going forward, I plan to keep studying the environments that reveal not only *whether* an intervention improves behavior, but *whom* it reaches and *how deep* the understanding runs, and use them to turn brittle behavioral change into durable understanding.