David S. Ahn

statement of research interests

The core of my research examines the foundations of decision theory and game theory. Experiments and introspection highlight the importance of behavior excluded by the standard economic paradigm. My work attempts to develop theoretical foundations that incorporate such behavior and allow for rigorous applied models.

One stream of research studies behavior facing ambiguous or imprecise probabilities. My primary job market paper, *Ambiguity without a State Space*, introduces a new theory of decisions involving both ambiguous probabilities and intractable states of the world. The domain of von Neumann–Morgenstern theory assumes that probabilities are precise, while the domain of Savage theory assumes that states are completely specified. This paper analyzes a third domain: sets of consequential lotteries. This approach, which questions existing foundations, contrasts with the standard approach to ambiguity, which proposes alternative utility functions or more specific state spaces within the basic framework of subjective expected utility. That framework assumes that the decision maker can access a comprehensive list of possible states. In the canonical example of ambiguity, the Ellsberg paradox, such a list is simple to construct; but in many actual decisions, the relevant states are difficult to comprehend. In fact, ambiguity is often *caused* by an inability to understand such complicated states. The domain of this paper allows for ambiguity without invoking any state space. Its relatively simpler theoretical primitives are perhaps closer to the cognitive fundamentals used by real decision makers.

These primitives still yield to clean axiomatic analysis. I characterize a utility representation on sets that integrates a transformed expected utility function over single lotteries with respect to a second order measure. The representation provides two mechanisms for capturing ambiguity aversion. The first is the curvature of the transformation, which resonates the theory of risk and produces a natural Arrow–Pratt measure of quantitative ambiguity aversion. The second is the weighting of the second order measure, which reflects the decision maker's sense of her own "luck" and improves the extreme pessimism implied by maxmin utility. I propose a general definition of comparative ambiguity aversion and uniquely characterize ambiguity neutrality. I finally present applications to reinsurance and asset pricing which I plan to develop as separate papers.

Having established foundations for a single decision maker, my paper *Hierarchies of Ambiguous Beliefs* studies settings with multiple agents. The beliefs of interacting agents are subject to the same ambiguity as those of a solitary decision maker, but to a much higher order: even if a particular agent has a precise belief about the natural uncertainty, she can still hold imprecise beliefs about others' beliefs, about others' beliefs regarding her beliefs, and so on. Meanwhile, the formulation of standard Bayesian games precludes any ambiguity. I consider agents that have ambiguous beliefs about others' ambiguous beliefs. The main result shows that common knowledge of a consistency condition on these hierarchies of ambiguous beliefs produces a universal type space: each type identifies a set of beliefs about others' types. Furthermore, common knowledge of precise beliefs carves out the standard universal type space as a subset of this more general space. Invoking

the space of ambiguous types, I formulate Bayesian games of both incomplete and ambiguous information. This shows decision theories with imprecise beliefs can generalize to multiple agents without sacrificing the tractability of the implicit Harsanyi construction.

A second stream of work studies altruism. If each agent's utility for an outcome depends on others' utilities, is there a stable utility profile? In *Stability of Interdependent Utility*, I examine this question. Each agent has a reaction function that takes a profile of others' utility functions to a personal utility function. A utility profile is stable if it is a fixed point of the profile of reaction functions. The paper's novel observation is that these reactions are structurally identical to social welfare functions, and can similarly be discussed axiomatically. If each reaction function is cardinally Paretian (or altruistic), there exists a stable utility profile. I then consider interdependence with incomplete information, where each agent is uncertain how much value the others place on a particular outcome, and offer a notion of stability for this setting. A form of stochastic dominance, which generalizes altruism to probabilistic environments, suffices to produce stable utilities.

Another paper, *Stability of Interdependent Preference*, conducts a similar exercise on preference orders. The analysis is completely ordinal. Each agent now has an Arrovian preference aggregator that takes a profile of preferences to a personal preference order. One innovation of the paper is an ordinal definition of altruism. The ordinal Pareto axiom is far too strong; I suggest an analog of Maskin monotonicity. If each aggregator is altruistic in the suggested sense, then the profile of aggregators has a fixed point.

Beyond my research in theoretical foundations, my applied interest in political economy dates back to my undergraduate training in political science. In *Repeated Elections with Noisy Policy-making* (in progress), politicians implement one-dimensional policies with some idiosyncratic shock. Each candidate's ideal point and the distribution of her shock is private information. A single voter decides each period to keep the current incumbent or elect a random challenger. If the voter is risk averse, the model generates a natural notion of competence as less noise. Similar models without noise predict: stationary voting depends only on the policy enacted in the previous term; and incumbents are either defeated immediately or reelected indefinitely. In contrast, my model forces the optimal voting rule to condition on the entire history of the incumbent's policies, since the voter learns about competence over time. In addition, incumbents are sometimes reelected multiple terms before eventually being defeated.

A final project, *Learning Self-Control* (in progress), studies the self-awareness of a temporally inconsistent agent. Existing models stress the distinction between a naive agent who is unaware of her self-control problem and a sophisticated agent who correctly anticipates her future actions; they also assume that the agent's self-awareness is fixed over time. This model allows the agent to learn about her inconsistency by creating a subjective statistical link between the self-control of past and future time identities. Under mild statistical assumptions, the agent updates her beliefs regarding future identities until she realizes the level of her self-control problem.