Faheem H. Gilani

RESEARCH STATEMENT

Department of Economics
University of Texas at Austin
fhg3@utexas.edu
faheemgilani.com

I am a microeconomic theorist working in public economics, dynamic games, and decision theory. My job market paper studies taxation when firms can reduce their liabilities by underreporting their tax-relevant bases or dividing into smaller entities. With my co-advisor Maxwell Stinchcombe, I model decision making when the precise distribution of outcomes is not known (*ambiguity*). With my co-advisor Caroline Thomas, I study how agents explore and exploit opportunities with uncertain payoffs in the presence of free-riding and pre-emption motives (*strategic experimentation*).

My long-term goal is to bridge the gap between theoretical models and their implementation. I plan to extend my job market paper into a richer theory of size-based rules for agents who can divide their activity into smaller units, expressing optimal policy in quantities authorities can measure. I am also pursuing a tractable and applicable theory of updating under ambiguity.

When developing models, I try to identify reductions or push for assumptions that lead to computational tractability rather than just theoretical generality. This approach is guided by my background in applied mathematics, which equipped me with a holistic understanding of the interplay between theory and computation. In my job market paper, the main traction came from scaling all manipulation costs by a single private parameter, which let me translate proportional tax changes into changes in effective manipulation costs and derive an accounting identity linking tax receipts to firms' costs. Together with a "compliance at the top" assumption and restrictions on the cost distribution and welfare weights, this yielded a concave formulation of the policy problem. In my work on ambiguity, I helped identify a tractable simplicial domain and a stochastic dominance axiom which led to an additive utility representation that allowed us to decompose risk and ambiguity attitudes and obtain bounds on relative ambiguity aversion in terms of risk aversion.

JOB MARKET PAPER

Optimal Taxation When Firms Can Underreport and Split

With few exceptions, theories of optimal taxation with avoidance and evasion treat taxpayers as single entities, despite evidence of tax-motivated firm division. In the Netherlands, Massenz (2025) finds a persistent increase in companies per owner once the tax advantage of separate assessment was perceived as permanent and minimum capital requirements were abolished. The paper finds no corresponding effect on employment or total assets, consistent with a reorganization of legal entities rather than an expansion of productive activity. My job market paper examines how such splitting affects tax revenue, welfare, and optimal policy.

The model. A tax authority designs a marginal tax schedule, a fee per entity, and enforcement against tax manipulation. Firms then manipulate their taxes by choosing sequences of underreports and splits. For example, a firm could reduce its operations, divide its remaining operations among newly created entities, illegally conceal part of the new entities' tax bases, and so on. The baseline cost of tax manipulation differs across firms and rises linearly with enforcement levels. Enforcement is costly for the tax authority. For example, the authority could spend administrative resources on enacting stricter laws, increasing scrutiny, or increasing audits.

A Ramsey rule for entities. Optimal marginal rates satisfy a Ramsey rule (Ramsey, 1927; Diamond, 1975). Each reporting threshold plays the role of a good, and its *entity base*—the number of entities reporting above it—plays the role of quantity demanded. This means, roughly, that at the optimum a one-percent proportional cut in the entire tax schedule, including the per-entity fee and holding enforcement fixed, changes the entity base at each interior rate by a percentage determined by how the authority values a dollar retained by the owners represented in that base relative to an additional dollar of public revenue. Greater welfare weight on these owners implies a smaller percentage response. Under the model's restrictions on costs and welfare weights, and assuming that the highest-cost firms comply under every admissible policy ("compliance at the top"), an accounting identity gives a concave formulation despite jumps in reports. The joint optimality conditions for tax rates, the entity fee, and enforcement, together with a binding revenue constraint, are then sufficient for a global optimum.

The threat of splitting. Under certain conditions, an optimal progressive tax schedule is flat at the top. Splitting can also increase burdens on small firms: larger firms pay each lower rate once per entity, making those rates more attractive to the tax authority. In the paper's three-type example, with enforcement fixed at zero, allowing splitting raises small firms' average burden by 38 percent and lowers large firms' by 9 percent. Yet only 2.6 percent of firms split at the new optimum, compared with 13 percent under the schedule that is optimal without splitting. The observed incidence of splitting can therefore understate its influence on tax design. Subadditivity, the standard remedy for multiple purchases under nonlinear pricing (Wilson, 1993), need not deter splitting: when smaller entities can conceal income more cheaply, firms may split even under a flat tax.

What to measure. The *coverage ratio* at a threshold is the number of entities reporting above it divided by the number of firms whose potential tax base exceeds it. With a common power-law distribution of private cost parameters, this ratio determines the revenue raised per dollar of additional burden by a marginal rate increase. Under equal welfare weights, coverage ratios therefore rank local revenue-raising rate changes without requiring knowledge of firms' manipulation technologies. Tax returns provide entity counts; random audits and ownership links could help estimate potential firm bases.

> **Future research.** I plan to build on this project by giving firms' production decisions a more explicit role. The current paper captures concealment and possible reductions in activity through reduced-form costs. I plan to begin with a turnover-tax model in which an owner chooses employment and output alongside reporting and legal structure. This would clarify when small-business relief expands production and when it mainly creates additional entities. Comparing responses to reforms at the entity and owner-group levels, using tax returns linked to payroll and audit information where available, could help distinguish these mechanisms and inform the model's cost parameters.

Strategic Experimentation with Multidimensional States

I have two papers on strategic experimentation: *Hopeful Experimentation* and *Strategic Experimentation with Evolving Resource Quality*. The workhorse model of strategic experimentation is the exponential bandit model of Keller, Rady, and Cripps (2005). Players allocate resources between a safe activity and a risky activity that is either worthless or generates lump-sum payoffs at a known rate if productive. Conditional on productivity, the reward process is fixed. My papers relax different features of this benchmark. In *Hopeful Experimentation*, a productive activity has a fixed but unknown arrival rate: it may produce large payoffs rarely or small payoffs frequently, with the same expected flow return. In *Strategic Experimentation with Evolving Resource Quality*, the probability that an opportunity pays off recovers over time and resets whenever an opportunity arrives. Both models involve multidimensional state variables. Solving these models directly would therefore require analyzing the partial differential equations (PDEs) defined by the Hamilton–Jacobi–Bellman equation for the values. It is difficult to solve these PDEs in general and to apply smooth pasting and value matching to the solutions. In the resource model, we simplify the analysis by transforming beliefs and the recovery process. In *Hopeful Experimentation*, I use cumulative experimentation to reduce the learning problem before a first success to a single variable.

Strategic Experimentation with Evolving Resource Quality *joint with Caroline Thomas*

The following is a stylized setting for the model in this paper. Before firewalls became more sophisticated, scammers would target economics graduate students with phishing scams at the beginning of every fall semester, usually by pretending to be the head of the department and asking students to send them money. Each attempt, if it makes it through cybersecurity, has some probability of success that depends on how cautious the students are. A target population could be completely resilient to scams (its cybersecurity could be of very high quality, or all of the students might be exceptionally wary), or there could be some chance that scams are successful. In the latter case, an encounter—someone taking the bait, whether or not they ultimately pay the scammer—can trigger reports and department-wide warnings that raise everyone's caution for some time. If someone in the target population responds to such a phishing attempt, this tells scammers that the population is vulnerable, and they will then try to exploit individuals in that population after waiting for a bit to lower the caution levels. But phishing is a competitive activity, and rivals can raise the caution level of the population, either through a successful payout or a reported encounter, thereby making it harder for all other scammers to exploit the population. When the probability that the population is vulnerable is low, scammers may have free-riding incentives (they prefer to learn from other scammers' attempts rather than expend resources trying to exploit the population themselves). Once that probability is high, there is a tragedy of the commons: scammers want to wait for caution levels to go down before attempting scams but prefer exploiting the population before others raise the caution level. This paper analyzes where each of these two forces dominates.

The model. Players divide effort between a safe activity and harvesting a potentially productive resource. Effort generates payoff opportunities only if the resource is productive. These opportunities translate into payoffs with a probability that grows deterministically in the absence of opportunity arrivals and resets to zero following an arrival (whether or not it translates into a payoff). Every opportunity is public, so it conclusively reveals that the resource is productive. We solve the cooperative problem (the planner's solution) and construct a symmetric Markov-perfect equilibrium (MPE) under conditions that hold in sufficiently large groups. Confident players resume exploitation before recovery makes it optimal for the group (the tragedy of the commons).

When the prior belief about the resource's productivity is high, players cumulatively experiment too little (free riding). The two distortions are connected: future competition erodes the returns from discovering a vulnerable population, weakening the incentive to produce information even though everyone benefits from it.

Method. Strategies depend on two state variables: beliefs about productivity and current resource quality. We express beliefs in log odds and measure quality with a clock that advances uniformly during recovery. Dividing a player's payoff above the safe return by the probability that the resource is unproductive turns the game before discovery into a deterministic differential game. Along paths of waiting or full effort, the marginal gain from experimentation satisfies a linear ordinary differential equation. Integrating this equation and imposing indifference where effort is interior lets us construct the equilibrium region by region. We can then match values and derivatives at the boundaries and verify that no player gains by deviating.

Hopeful Experimentation

An author might be unproductive and never produce work, or they may be productive and produce many modest works frequently, or they may produce high-quality works infrequently. Suppose these two productive types have the same expected flow return to a publisher. Continued silence from the author makes it more likely that they are unproductive but, conditional on their being productive, makes it more likely that they are the type of author to produce high-quality works infrequently. How would a publisher divide resources between an established author offering a known return and a new author with a higher return if productive, but uncertain viability and output frequency?

The model. Players allocate resources between a safe return and a project with unknown viability and, if viable, an unknown arrival rate but a fixed expected flow payoff. In the single-player setting, compared with a benchmark where the arrival rate is known, this uncertainty shortens experimentation without success for less promising projects and lengthens it for more promising ones. Yet whenever experimentation is worthwhile, it lowers value and raises the chance of abandoning a viable project. In the unique symmetric MPE, sufficiently optimistic players experiment fully, then gradually scale back as each becomes indifferent between experimenting and free riding. Without success, effort approaches zero and cumulative experimentation converges to the single-agent stopping level, below the efficient amount. Near that boundary, even the group's combined effort falls below that of a lone experimenter.

Method. The state variables are beliefs about viability and the arrival rate; with two possible positive rates, these beliefs are two-dimensional. Before the first arrival, the full posterior is determined by cumulative experimentation. I use the Laplace transform of the arrival-rate distribution to calculate the probability that a viable project remains silent after any amount of effort. Expressing payoffs in this experimentation clock gives scalar equations for the stopping rule and equilibrium effort, even with a general distribution of arrival rates. In the strategic model, the clock accumulates everyone's effort, making explicit how each player's experimentation contributes to information available to the group.

> **Future research.** I plan to develop a model of strategic experimentation with multidimensional beliefs in which the first successful outcome does not establish a project's commercial viability. The setting I have in mind is geothermal energy production near natural steam vents, with firms using wells to tap a shared reservoir. Successful short tests can leave sustained profitability unresolved, as the Department of Energy's discussion of extended testing at FORGE illustrates (DOE,

2026). I would allow several productivity types, including ones that produce positive output but yield less than an operator's outside option, so later evidence could lead a firm to abandon a previously successful project. Building on the inconclusive-news model of Keller and Rady (2010), I would study continued learning alongside pressure depletion and recovery. Pressure interference among operators is documented at The Geysers by Sanyal and Enedy (2011). In the proposed model, public test results would supply information to competitors while withdrawals reduce their immediate production opportunities. Starting with fixed productivity types and observed pressure, I would combine analytical benchmarks with numerical equilibrium solutions to ask whether encouraging results induce further experimentation that benefits the pioneer through learning or whether the resulting withdrawals undermine that incentive. Comparing independent and coordinated testing would clarify how withdrawal limits affect both learning and resource use.

The von Neumann–Morgenstern (vNM) Approach to Ambiguity

with Maxwell Stinchcombe

A gamble is *risky* when its outcomes have a known probability distribution and *ambiguous* when that distribution is known only to lie in a set. A bet on a fair coin is risky; a bet on a coin whose probability of heads is known to lie between one-half and three-quarters is ambiguous. Subjective models of risk derive expected utility through axioms over acts: functions from a subjective set of states to outcomes, as in Savage's (1954) *The Foundations of Statistics*, or to objective lotteries, as in Anscombe and Aumann (1963). The classic axioms yield a subjective expected utility representation that values acts by averaging utility under a single distribution (the prior) over the state space. Models of ambiguity usually modify these axioms to allow for multiple priors. In the maxmin representation of Gilboa and Schmeidler (1989), uncertainty aversion and the other axioms imply that each act is evaluated by its lowest expected utility over the admissible priors. The decision maker chooses the act with the highest such value; the least favorable prior can differ across acts.

The objective approach to ambiguity instead places axioms directly on preferences over sets of outcome distributions (Olszewski, 2007; Ahn, 2008). State-space models offer a rich domain of acts. Unless the axioms require otherwise, preferences can distinguish acts that induce the same set of outcome distributions; maxmin expected utility treats such acts alike. A separate issue is which sets a fixed collection of priors can generate. Stinchcombe (2016) shows that this descriptive range can be restricted. For example, with acts mapping states directly to heads or tails, admitting every prior over states allows either a certain outcome or complete ambiguity between the two. It cannot represent the narrower probability interval in the coin example. The objective approach makes the domain of outcome-distribution sets explicit from the outset.

Following Dumav and Stinchcombe (2013), our work pursues a von Neumann–Morgenstern analogue of expected utility for ambiguous prospects. In *Choice over Sets of Distributions*, we identify a simplicial domain whose extreme points are sets of distributions that specify the possible outcomes while leaving their probabilities unrestricted. Every prospect in this domain has a unique representation as a mixture of these extreme sets. They play the role that sure outcomes play in expected utility: under the mixture-space axioms of Herstein and Milnor (1953), a prospect's utility averages their values using its mixture probabilities.

For example, let A be the set of coin distributions assigning probability $p \in [1/2, 3/4]$ to heads (H) and the remaining probability to tails (T). Writing δ_H and δ_T for the distributions concentrated on

heads and tails, respectively, and $\Delta(\{H, T\})$ for the set of all distributions over heads and tails, we can verify that

$$A = \frac{1}{2}\{\delta_H\} + \frac{1}{4}\{\delta_T\} + \frac{1}{4}\Delta(\{H, T\}).$$

Here, addition is Minkowski addition: $\alpha X + Y$ is the set of all $\alpha x + y$, where $x \in X$ and $y \in Y$. The mixture assigns one-half weight to certain heads, one-quarter to certain tails, and one-quarter to complete ambiguity between the two. Each extreme set is a definite object in the domain, even when it leaves the realized outcome uncertain. The simplicial structure makes these mixture weights unique.

The domain includes ordinary lotteries and standard Ellsberg bets. A stochastic dominance axiom yields an additive representation in the lowest and highest outcomes of each extreme set. This representation separates risk and ambiguity attitudes and allows ambiguity attitudes to vary with outcome levels. With smooth utility indices, second-order dominance imposes a testable link between the two: risk aversion constrains how quickly the local ambiguity coefficient can change with outcomes. In particular, risk neutrality over an interval requires a constant ambiguity coefficient there. These restrictions can be assessed using choices over risky lotteries and ambiguous intervals. Our applications show how decreasing ambiguity aversion can generate wealth-dependent portfolio participation, why ambiguity aversion alone does not determine whether wider income bounds raise saving, and how stochastic dominance of outcome bounds gives treatment rankings that are robust across the preferences in our class.

In *Updating Sets of Distributions*, we use the unique decomposition to extend this approach to learning. The prior assigns probabilities to extreme sets, and a finite-valued signal has a known likelihood conditional on each set. Bayes' rule updates these probabilities; mixing the extreme sets with their posterior weights gives the updated set of outcome distributions. We show that this procedure agrees with conditioning each prior in the state-space representation we construct. For finite menus of actions mapping possible-outcome sets into consequence sets and plans evaluated by averaging continuation values, we obtain dynamic consistency: a plan is optimal before learning exactly when it prescribes an optimal action after every possible signal. More informative experiments weakly increase the optimized value. Learning can therefore improve decisions while leaving ambiguity unresolved: identifying the relevant outcome set need not determine the probabilities of outcomes within it.

> **Future research.** I plan to build on these foundations with a dynamic programming application that can be confronted with data. A promising setting is irreversible investment, where a firm chooses whether to invest, wait for public information, or buy a diagnostic test. I would begin with a finite-horizon model with a small number of possible return ranges and tests of known accuracy about which range applies. Probabilities within a range would remain unresolved. The first task is to establish a recursive formulation and use it to study how persistent ambiguity affects investment timing and demand for information. Connecting these choices to elicited beliefs about return ranges and uncertainty about probabilities would help distinguish the model's implications from those of ordinary risk.