Research Statement

Can Urgun

I am mainly a microeconomic theorist with wide interests in dynamic decision making. I broadly work on search theory, dynamic contracts, stochastic games. My broader interests search and contracts have also led me to engage in experimental work in search theory and some theoretical work in mechanism design. In the following sections I will describe my published work, ongoing work and future research interests on broadly search and dynamic contracts and games.

1 Search

My work in search theory focuses primarily on search in correlated environments, where a decision-maker proceeds through a sequence of observations until deciding to conclude the search. In the simplest form, the searcher is allowed to return to previous observations. This basic discovery process is integral to many economic decisions, ranging from online shopping to canonical models in search theory. Beginning with landmark studies such as those by Stigler (1961), McCall (1970), and Mortensen (1970), traditional search models assume independent search outcomes. Independence greatly simplifies search analysis: in basic models, the distribution of possible discoveries is assumed to be fixed. Even in sequential search, recall plays a limited role, with the agent only needing to consider one step ahead. Thus, the agent effectively solves a static optimization problem.

However, in many real-world settings, search follows a path, and prior discoveries not only serve as valuable fallbacks but also influence future discoveries. My work seeks to capture search environments where current observations provide information about future ones. Specifically, in a series of papers, I examine search over a prescribed path of Brownian motion, where the time variable corresponds to the index of the search alternative. As in Callander

(2011), Brownian motion governs the value of each ordered discovery, meaning that the realized Brownian path maps each point in time—representing an ordered discovery—to its value. This framework naturally introduces intertemporal correlation between discoveries: one observation informs expectations for future ones.

Without the technical convenience of independence, models that allow for correlation often assume short-lived or myopic agents with a fixed number of alternatives explored per unit of time, as in Callander (2011) and Garfagnini and Strulovici (2016). In Urgun and Yariv (2021), we first compare simple search over Brownian motion with searching over independent and identically distributed (i.i.d.) normal variables. In the i.i.d. case, search ends when the searcher reaches a satisfactory observation, and recall is not used. In the correlated case, however, the searcher stops when observations become sufficiently bad, and recall is always employed.

In Cetemen et al. (2023), we focus on group search, where search intensity is a choice variable. A group of searchers jointly determines their search intensity over a Brownian path paying an individual cost that depends on their effort choice. Each searcher observes the search outcomes but not the effort the other searchers exert. Additionally each searcher can choose to conclude their own search implementing the best outcome that the team has found so far. We explore how searchers conclude their search and how future expectations influence current search intensity and free-riding incentives. The optimal search intensity of each searchers turns out to be not forward looking and is a best response to the current group of searchers. The optimal order of searchers concluding is essentially pre-determined forming exit waves. Notably, the time when the wave happens depends on the particular path and is therefore random but the order of the searchers concluding their search is deterministic and can be identified from the primitives of model.

In Urgun and Yariv (2025), we address the full individual search problem, where a decision-maker controls both search intensity and the timing of the search's conclusion. In this model, the decision-maker may also face discounting in addition to search costs. We characterize the optimal search intensity and how it determines the optimal stopping rule, which follows a "drawdown" strategy: the search concludes when the current observation falls below a threshold determined by past observations and search costs. Our sharp characterization also allows us to embed the search model into a contracting environment, where a principal hires a searcher on commission. We fully characterize both the searcher's behavior and the optimal commission contract the principal should offer.

In ongoing theoretical work with Theo Durandard and Leeat Yariv, we extend the analysis in Urgun and Yariv (2025) to a multi-dimensional setting. Here, the searcher faces multiple potential avenues, each represented by an independent Brownian motion. At any given time, the searcher chooses both which path to explore and the search intensity in the chosen path, with instantaneous costs dependent on both the path chosen and the intensity. The search culminates in the searcher selecting the highest value observed across all paths. The prioritization of paths is determined by an index, and within each path, the searcher optimizes the search intensity as if only that path were available. The index of each path takes a remarkable form: it is the maximum of the value associated with the single path found in Urgun and Yariv (2025) and the historical maximum along that path. Notably, compared to the scenario where each path is analyzed in isolation, only the ultimately chosen path reaches its drawdown point, while other paths are abandoned.

Our precise characterizations of correlated search models and the significant behavioral differences compared to i.i.d. models have prompted us to explore whether these theoretical predictions hold in laboratory settings. In collaboration with Leeat Yariv, Sevgi Yuksel, and Clement Herman, we are conducting experiments at PExL to compare search behavior in i.i.d. environments, as outlined in Urgun and Yariv (2021), with correlated environments, as explored in Urgun and Yariv (2025).

Another project in search theory, "When to Decide: Timing of Choice in Parallel Search," is a joint work with Pietro Ortoleva. In this paper, we consider a decision-maker who must choose between two options that are ex-ante identical. The decision-maker holds a prior belief about their values and can gather a constant stream of information about each option (e.g., normal signals in continuous time), but incurs a cost for the time spent. When should the decision-maker stop evaluating the options and make a choice? This is a fundamental problem in decision-making, traceable to classical models of hypothesis testing. To our knowledge, a general solution remains unknown, except for special cases such as when the value difference between options is known ex-ante, as in the Drift Diffusion Model Fudenberg et al. (2018). Such models are prevalent in neuroscience (Tajima et al. (2019)), but the lack of an analytical solution often requires simultaneous estimation of optimal behavior and model parameters. In our study, we provide a complete characterization of the optimal solution, including the functional form of the stopping boundary when waiting costs are convex, the items are drawn from the same normal distribution, and the signals have equal noise levels.

The broad area of search both theoretic and experimental is exciting to me and I plan to keep working on this research agenda.

2 Dynamic Games-Contracts

My earlier work on dynamic games examines how decision-makers should optimally determine when, if ever, to initiate a destructive confrontation with an adversary. In Sandroni and Urgun (2018), we explore models of ordinary conflicts, such as divorce, price wars, and commercial litigation, as a stopping game, analyzing the interplay between patience, aggressiveness, and strength. Specifically, we model an imminent conflict as a sequence of opportunities, with each party independently deciding whether to initiate a confrontation after observing the opportunity they face. If neither party starts the conflict, both receive zero payoffs and the game continues. If at least one party decides to initiate the conflict, it is resolved based on the opportunity available in that period, with one party receiving a positive payoff (the winner) and the other a negative payoff (the loser). A critical aspect of the model is that conflicts are ultimately destructive, with the sum of payoffs being negative. We show that when opportunities are fair and independent, the unique equilibrium as parties become more patient converges to immediate conflict, forming what we call an anti-folk theorem. Conversely, if one side holds an intrinsic advantage and the parties learn who has this advantage by observing opportunities, the advantaged party becomes less likely to initiate conflict as their patience increases, while the disadvantaged party becomes more inclined to do so immediately. The reasoning behind patience leading to faster confrontation lies in the fact that in a strategic environment with fully rational and forward-looking players, the incentive to preempt overrides potential gains from waiting for better opportunities. This highlights the subtle distinction between a stopping game and a repeated game, despite the apparent similarities between the two settings.

In Sandroni and Urgun (2017), we explore a similar environment to Sandroni and Urgun (2018), with one key difference: one player is not strategic. Notably, even if the non-strategic player follows a fixed aggressive strategy, we assume they do not best respond to the rational player's strategy. This assumption is sufficient to overturn the incentive to preempt entirely. Regardless of the non-strategic player's aggressiveness or the process generating opportunities, a more patient rational player will always be inclined to postpone conflict until a suitably favorable opportunity arises. Thus, increased patience leads to longer delays in confrontation, with infinitely patient players refraining from engagement unless they are almost certain of victory.

In Cetemen et al. (2023), we analyze a dynamic moral hazard problem with a fixed finite horizon, where both the principal and the agent exhibit generic non-exponential discounting and are able to renegotiate. Since the parties are time-inconsistent, their preferences over the contract change, making renegotiation critical. While the finite horizon complicates recursive approaches, it allows for the application of backward stochastic differential equations to characterize the principal's optimal renegotiation-proof contract. This characterization is accomplished through an extended Hamilton-Jacobi-Bellman system rather than a single equation. A notable technical result of independent interest is the existence of solutions to such systems. To our knowledge, this remains the only result proving the existence of stochastic optimal control with time inconsistency. We also derive the contract in closed form for simple types of non-exponential discounting, such as quasi-hyperbolic discounting or anticipatory utility.

My paper, "Restless Contracts" analyzes dynamics in relationships between a principal and multiple agents, where agents improve (or worsen) when employed by the principal and deteriorate (or recover) when not. The principal can utilize these improvement opportunities as additional incentives, but agents in advantageous positions may use their status as a threat to exit the relationship and collect additional rents. The problem resembles a restless bandit model (where unemployed agents deteriorate rather than remaining static), but with added complexity due to agents' exit threats. It turns out that the contracting environment and the presence of exit threats allow for a more tractable solution to the principal's problem compared to general restless bandit problems, which suffer from significant tractability issues, including challenges related to indexability. Specifically, a Whittle index, constructed through incentive conditions, is optimal for the principal and is identified in closed form.

Finally, a paper with a different focus, "Costly Verification and Money Burning," co-authored with Rohit Patel, addresses resource allocation within a firm, though unlike most of my previous work, the interaction is one-shot. We examine the allocation of a scarce resource—such as computing equipment, personnel, or a research grant—to one of many agents, who are heterogeneous in their productivity gains from the resource, which they privately know. The principal can employ two instruments, previously studied in isolation but frequently used together in practice: costly verification and money burning. We identify the optimal allocation protocol, which takes one of two forms: either both money burning and costly verification are used as complements, or money burning is not used at all. When used together, these instruments achieve allocative efficiency, with each type fully differentiated and the most productive agent receiving the resource. In contrast, costly verification alone results in significant pooling, where types below the verification cost are treated identically despite differing productivities. Money burning alone can achieve allocative efficiency, but at a large surplus loss, as each productivity increase requires progressively higher levels of money burning. When both instruments are available and used together, money burning is used for lower types, increasing in severity until productivity reaches the verification cost. Higher types burn the same amount of money, with verification probabilities rising, thereby avoiding excessive surplus loss.

I am broadly interested in dynamic decision making whether at an individual level or within a contractual framework. Some of my work on search opens up some interesting environments such as commission contracts for search which we partially explored in Urgun and Yariv (2025) but I intend to explore further dynamic contractual settings with discovery and search in the future.

References

Callander, S. (2011). Searching and learning by trial and error. *American Economic Review 101*(6), 2277–2308.

Cetemen, D., F. Z. Feng, and C. Urgun (2023). Renegotiation and dynamic inconsistency: contracting with non-exponential discounting. *Journal of Economic Theory 208*, 105606.

Cetemen, D., C. Urgun, and L. Yariv (2023). Collective progress: Dynamics of exit waves. *Journal of Political Economy 131*(9), 2402–2450.

Fudenberg, D., P. Strack, and T. Strzalecki (2018). Speed, accuracy, and the optimal timing of choices. *American Economic Review 108*(12), 3651–3684.

Garfagnini, U. and B. Strulovici (2016). Social experimentation with interdependent and expanding technologies. *Review of Economic Studies 83*(4), 1579–1613.

McCall, J. J. (1970). Economics of information and job search. *Quarterly Journal of Economics*, 113–126.

Mortensen, D. T. (1970). Job search, the duration of unemployment, and the phillips curve. *American Economic Review 60*(5), 847–862.

Sandroni, A. and C. Urgun (2017). Dynamics in art of war. *Mathematical social sciences 86*, 51–58.

Sandroni, A. and C. Urgun (2018). When to confront: the role of patience. *American Economic Journal: Microeconomics 10*(3), 219–252.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy 69*(3), 213–225.

Tajima, S., J. Drugowitsch, N. Patel, and A. Pouget (2019). Optimal policy for multi-alternative decisions. *Nature neuroscience 22*(9), 1503–1511.

Urgun, C. and L. Yariv (2021). Constrained retrospective search. *AEA Papers and Proceedings 111*, 549–53.

Urgun, C. and L. Yariv (2025). Contiguous search: Expoloration and ambition on uncharted terrain. *Journal of Political Economy*.