# Research Statement

My research is focused on providing new ways for understanding and computing optimal decisions and strategies in complex operational settings, in ways that not only support decision-making, but also render quantitative models more pertinent to real-world settings. My work has had two main thrusts: a more methodological one, aimed at developing new tools and algorithms for dynamic optimization under uncertainty and risk, and a more qualitative one, aimed at enhancing our understanding of the interplay between key operating decisions and important financial and contractual considerations that firms face.

Some of my early work on dynamic optimization under uncertainty has followed the Robust Optimization (RO) paradigm, a methodology developed during the last 20 years to analyze and optimize the performance of complex systems under limited information.[1] My research started with [1], where I investigated a classical model for inventory management under limited demand information, for which I showed that a new class of replenishment policies, depending affinely on the history of observed demands, is actually optimal. The result is practically relevant: the optimal policies have an intuitive interpretation (new orders are assigned to partially meet all backlogged demands), and are easy to find, by solving a single linear optimization problem of small size. Interestingly, in [3], I showed that such "simple" policies remain optimal even if the cost structure becomes "complex" (e.g., arbitrary convex ordering costs), when typical approaches such as Dynamic Programming (DP) become intractable or prescribe very complex policies. This allows adapting the approach to more realistic capacity planning problems under uncertainty, such as the design of flexible retailer-supplier contracts with pre-commitments and penalties [3]. From a theoretical perspective, these results were the first of their kind, and the proof techniques were entirely novel, combining continuous and discrete optimization to construct the optimal strategy.[2]

These results were enabled by a simple—yet critical—observation that my work pointed out: namely, that *multiple* solutions or policies typically exist in RO that achieve the same worst-case performance, but only some of which also satisfy the well-known Bellman optimality principle [4]. In this sense, my work in [1, 2, 3] showed how this degeneracy can help: among these policies, some might have a "simple" structure, and could be obtained by solving "simple" problems, without resorting to potentially intractable DP recursions. To argue the existence and to retrieve such policies though, one cannot entirely rely on the usual recursive DP arguments, and new techniques must be developed.

Viewed differently, the same observation also highlighted a potentially serious flaw of the RO paradigm: namely, that the degeneracy may cause severe inefficiencies and suboptimal performance when non-worst-case outcomes materialize. To this end, in [5], I introduced and

[1]The framework has a long history in various disciplines, including economics, engineering, and statistics, and is built around the premise that an exact probability distribution is often unavailable in a decision process due to, e.g., a lack of ample historical data or inherent non-stationarities. The standard approach is to assume that only some partial information is available (e.g., support, moments, etc.), and to seek feasible decisions that optimize the worst-case (expected) performance.

[2]In [2], I also examined more complex supply networks, where I showed how optimal policies can be obtained by solving a hierarchy of tractable, convex optimization problems. The approach performs well in simulation, and places little burden on the decision maker, who can adjust the performance-computation trade-off through a single parameter.

formalized the concept of Pareto efficiency in RO. I showed that decisions generated using classical RO methods need not have this property, and that particular algorithms/software used in RO problems are *guaranteed* to produce inefficient outcomes. I also provided simple procedures for testing whether a given solution is Pareto-efficient, and for generating Pareto-efficient outcomes. This methodology was a strict improvement of the RO framework, incurring no extra computational cost, and significantly reducing conservativeness.

This line of research allowed me to recognize some of the benefits and pitfalls of RO first-hand, and lead me to explore in more depth its connections with axiomatic approaches for decision-making under limited information and ambiguity developed in economics, decision analysis, and mathematical finance.[3] I remain convinced that RO carries benefits, both in generating computational solutions to complex problems—at an individual, firm, or societal level—but also in a more *qualitative* role, by providing motivation for certain "simple" solutions that "work well" (in keeping with the established principle that the more tailored solutions are to existing information, the more sensitive they become to slight changes). At the same time, I have also come to believe that naïvely applying RO methods and ignoring important features—e.g., new information, preferences beyond worst-case scenarios, etc.—can generate inconsistencies, inefficiencies, and even misleading results, which can reduce the framework's appeal.

Some of my on-going research strives to build on this philosophy. In [14], I consider the problem of determining when to monitor and possibly alter the course of several processes governing future outcomes. This is a generic question with many applications, from healthcare—where physicians treating chronic conditions[4] choose the frequency of office visits or tests—to collateralized lending or procurement contracting—where a firm's cashflows and assets are evaluated to assess whether an outstanding loan is "safe," or the firm's operations are inspected to determine environmental violations. In a model where monitoring is potentially costly but can reduce the uncertainty about future process values, I provide general conditions when pre-committing to a set of equi-distant monitoring times—routinely done in practice—is just as good as adjusting the times dynamically, and I discuss tractable computational procedures for finding monitoring schedules with good performance. Similarly, in [15], I study a firm's procurement strategies under scale economies and limited information about deeper supply tiers, and show how simple plans—such as uniformly splitting orders among direct suppliers—can actually be optimal.

As some of these examples hint at, I believe that a natural set of problems where robustness can play a key role is the design of (long-term) relationships and contracts, where limited and/or asymmetric information, dynamics, moral hazard, contractual incompleteness, and learning are critical features to account for. This also brings me to the second main focal point of my research so far, which has been to understand how a firm's "nexus of contracts"—and particularly the material and financial flows—can be designed or improved

---

[3] Related research includes [6], where I propose approximating a decision maker's static risk preferences (which may be easier to state, but dynamically inconsistent) through axiomatically justified dynamic risk measures.

[4] In the related work in [7], I develop adaptive treatment policies for chronic illnesses based on dynamically monitoring a new patient's state in order to reduce the uncertainty related to the patient's response to treatment.

when accounting more precisely for salient operational features.[5]

Some of my research here focuses on the inefficiencies—or agency costs—that arise when inventory-heavy firms under leverage use their operating capabilities to engage in risky strategies that benefit equity-holders at the expense of debt-holders. My primary aim was to *quantify* the degree to which common debt contracts can mitigate such costs through a more "operations-centric" design. My work finds substantial distortions in operating policy caused by such risk-shifting incentives: firms tend to liquidate less/replenish more inventory [9], and they apply (increasingly) smaller price markdowns that compound over time [10]. These distortions, which are consistent with some observed behavior, lead to (increasingly) substantial agency costs when debt is fairly priced. However, I find that practical terms and contingencies commonly encountered in practice—e.g., amortization schedules, financial covenants, or borrowing base covenants combined with a suitable appraisal of the collateralized inventory—can either fully alleviate inefficiencies [9] or at least substantially reduce them [10]. For such contracts to work though, inventory-heavy firms with enhanced flexibility and/or selling in better markets must also be willing to abide by more and/or stricter requirements [9]. That such contracts rely on dynamic monitoring mechanisms that are quite sensitive to the available information further motivates some of our ongoing work in [14].

In several related research projects, I also examine how procurement contracts can be designed with a more precise view of the firm's supply base. This work was motivated to a large degree by discussions with industry partners. It includes [11]—where I examine the sourcing strategies of a firm that cannot directly contract with its more distant and potentially unreliable suppliers, but can influence the outcomes through contracts offered to its immediate suppliers—as well as [12]—where I study how a firm sourcing from many small suppliers (as is typical in agricultural supply chains) can incentivize better management practices and create shared value by changing sourcing channels and/or payment schemes. The findings confirm certain observations reported in the media and by our industry partners, such as the diamond-shaped supply chains that Toyota uncovered in its internal investigation following the Tōhoku earthquake [11], and the success/failure of simple contracting schemes in the wool supply chain [12]. While quantifying the "sticks" [11] or "carrots" [12] that can improve efficiency, the results also showcase the importance of certain operational and structural details in the contract design (e.g., the degree of overlap in the deeper supply tiers [11], their segmentation [12], their relative cost-effectiveness [11, 12], etc.), and the importance of limited or asymmetric information, which motivates our ongoing work in [15].

Looking ahead, I find myself increasingly excited by research questions aimed at understanding how firms with a globalized presence should coordinate their policies and strategies, with a long-term view towards creating *sustainable* and *shared* value. As a step in this direction, in [16], I study how large firms within the palm oil supply chain can design preferred payments offered to mills and/or farmers in rural communities in Indonesia, while aiming to reduce inefficiencies, improve livelihoods, and translate bold sustainability commitments—such as a "no-deforestation" policy—into concrete progress on the ground.

[5] Related research includes [8]—where I show how portfolio managers in charge of multiple client accounts could allocate transaction costs due to market impact and adjust portfolios to balance fairness and efficiency, and [13]—where I show how financial accounting considerations can turn a loyalty program into a tool for hedging against uncertainty.

## References


[1] D. Bertsimas, D. A. Iancu, and P. Parrilo, "Optimality of Affine Policies in Multi-stage Robust Optimization," ***Mathematics of Operations Research****, vol. 35, No. 2, pp. 363-394, 2010.*

[2] "A Hierarchy of Near-Optimal Policies for Multi-stage Adaptive Optimization," D. Bertsimas, D. A. Iancu, and P. Parrilo, ***IEEE Transactions on Automatic Control****, vol. 56, No. 12, pp. 2809-2824, 2011.*

[3] D. A. Iancu, M. Sharma, and M. Sviridenko, "Supermodularity and Affine Policies in Dynamic Robust Optimization," ***Operations Research****, vol. 61, No. 4, pp. 941-956, 2013.*

[4] E. Delage and D. A. Iancu, "Robust Multi-stage Decision Making," ***TutORials in Operations Research****, pp. 20-46, 2015.*

[5] D. A. Iancu and N. Trichakis, "Pareto Efficiency in Robust Optimization," ***Management Science****, vol. 60, No. 1, pp. 130-147, 2014.*

[6] D. A. Iancu, M. Petrik, and D. Subramanian, "Tight Approximations of Dynamic Risk Measures," ***Mathematics of Operations Research****, vol. 40, No. 3, pp. 655-682, 2015.*

[7] D. Negoescu, K. Bimpikis, M. Brandeau, D. A. Iancu, "Dynamic Learning of Patient Response Rates: An Application to Treating Chronic Diseases," *Forthcoming in* ***Management Science****.*

[8] D. A. Iancu and N. Trichakis, "Fairness and Efficiency in Multiportfolio Optimization," ***Operations Research****, vol. 62, No. 6, pp. 1283-1301, 2014.*

[9] D. A. Iancu, N. Trichakis, and G. Tsoukalas, "Is Operating Flexibility Harmful Under Debt?," *Forthcoming in* ***Management Science****.*

[10] O. Besbes, D. A. Iancu, and N. Trichakis, "Dynamic Pricing under Revenue Targets: Incentives and Spiraling," *Under second round of review in* ***Management Science****.*

[11] E. Ang, D. A. Iancu, and R. Swinney, "Disruption Risk and Optimal Sourcing in Multi-tier Supply Networks," *Forthcoming in* ***Management Science****.*

[12] J. de Zegher, D. A. Iancu, and H. Lee, "Designing Contracts and Sourcing Channels to Create Shared Value," *Forthcoming in* ***Manufacturing & Service Operations Management****.*

[13] S-Y. Chun, D. A. Iancu, and N. Trichakis, "Setting the Value of Loyalty Points," *working paper.*

[14] D. A. Iancu, N. Trichakis, and D.-Y. Yoon, "Optimal Monitoring Policies Under Limited Information," *in preparation.*

[15] H. Bidkhori, D. A. Iancu, and R. Swinney, "Managing Disruption Risk in Supply Chains with Limited Visibility," *in preparation.*

[16] J. de Zegher, D. A. Iancu, H. Lee, and E. Plambeck, "Operational Strategies to Achieve No-Deforestation Commitments in Palm Procurement," *in preparation.*