

Three Essays on Sustainable Operations: Renewable Energy
Procurement, Battery Storage, and Equitable Work Scheduling

Xiaoxuan Hou

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Shi Chen, Chair

Michael R. Wagner, Chair

Kamran Moinszadeh

Program Authorized to Offer Degree:
Business Administration

©Copyright 2026

Xiaoxuan Hou

University of Washington

Abstract

Three Essays on Sustainable Operations: Renewable Energy Procurement, Battery Storage, and Equitable Work Scheduling

Xiaoxuan Hou

Co-Chairs of the Supervisory Committee:
Professor of Operations Management Shi Chen
Foster School of Business

Professor of Operations Management Michael R. Wagner
Foster School of Business

This dissertation consists of three essays on sustainable operations management. The essays study how operational decisions shape both environmental outcomes and social outcomes in settings characterized by uncertainty, strategic interaction, and competing stakeholder objectives. Together, they examine three domains central to sustainable operations: renewable energy procurement, battery storage dispatch, and equitable workforce scheduling.

The first essay studies collective power purchase agreements for corporate renewable electricity procurement. Although power purchase agreements have become an important vehicle for financing renewable energy projects, small and medium-sized corporate buyers often lack the demand scale, expertise, or bargaining power to contract independently. The essay develops a generalized Nash bargaining framework to compare separate negotiation, consortium-based joint negotiation, and agent-based negotiation. The analysis shows that, under risk neutrality, all three mechanisms induce the same renewable project size, so the choice of negotiation architecture primarily affects surplus allocation and buyer participation rather than capacity expansion. When buyers are risk averse, project size may increase or decrease depending on the balance between hedging benefits and market-risk exposure, while renewable energy certificate banking increases the marginal value of capacity. The results provide guidance for when joint negotiation should serve as the participation-safe

default and when agent-based negotiation may be preferred for weak or asymmetric buyer coalitions.

The second essay studies battery dispatch in electricity markets with negative prices. It develops a tractable dynamic programming framework that embeds the forecasting state directly into the dispatch problem and solves the state-of-charge dimension exactly on an endogenous event grid. Using matched synthetic experiments and out-of-sample tests in California’s real-time electricity market, the essay compares stochastic dynamic programming with open-loop certainty-equivalent planning and rolling deterministic reoptimization. The results show that stochastic control creates substantial value relative to open-loop planning, but its advantage over deterministic reoptimization is selective rather than universal. In particular, the value of stochastic look-ahead depends on forecast-error dependence, market uncertainty, and storage duration. The essay further identifies operational signals that can support a switching rule between stochastic and deterministic controllers.

The third essay turns to social sustainability in service operations by studying disparities in the temporal quality of workers’ schedules. Using granular shift, transaction, and customer-review data from an American restaurant chain, the essay examines scheduling sufficiency and predictability across gender and racial groups. The analysis shows that women and minority workers receive fewer initially scheduled hours and more last-minute shift additions than White male workers, even after accounting for role, tenure, absenteeism, and multidimensional measures of worker ability. Further analyses suggest that these disparities are not fully explained by worker preferences or availability; instead, they point to systematic under-scheduling and unequal exposure to unpredictable work.

Across the three essays, the dissertation shows that sustainable operations require more than efficiency-oriented optimization. In renewable procurement, sustainable outcomes depend on mechanisms that align incentives and preserve participation. In storage operations, they depend on dispatch policies that use information effectively under uncertainty. In workforce scheduling, they depend on managerial processes that allocate operational flexibility fairly. By combining analytical modeling, dynamic programming, and empirical analysis,

this dissertation contributes tools and evidence for designing operations that are environmentally responsible, economically viable, and socially equitable.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	iv
Chapter 1: Introduction and Overview	1
1.1 Sustainable Operations as Operational Design	2
1.2 Essay 1: Collective Renewable Energy Procurement	3
1.3 Essay 2: Battery Dispatch in Electricity Markets	4
1.4 Essay 3: Workforce Equity in Service Operations	6
1.5 Cross-Essay Themes	7
1.6 Organization of the Dissertation	8
Chapter 2: Collective Power Purchase Agreement Negotiation	10
2.1 Introduction	10
2.2 Literature Review	14
2.3 The Base Model	18
2.4 Risk-Averse Buyers and Price Correlation	37
2.5 Extension: REC Banking	46
2.6 Concluding Remarks	52
Chapter 3: Battery Dispatch in Electricity Markets with Negative Prices: Forecast- Embedded Dynamic Programming	54
3.1 Introduction	54
3.2 Literature Review	57
3.3 Model	59
3.4 Electricity-Market Setting and Numerical Experiments	73
3.5 Discussion, Managerial Implications, and Conclusion	85
Chapter 4: Explain Worker Schedule Quality Disparities: Evidence from an Amer- ican Restaurant Chain	88

4.1	Introduction	88
4.2	Literature Review	93
4.3	Data	97
4.4	Schedule Disparities - The Role of Human Capital	105
4.5	Schedule Disparities - The Role of Worker Preference and Availability	108
4.6	Discussions	118
4.7	Concluding Remarks	122
	Bibliography	125
Appendix A: Supplementary Material for Chapter 2: Collective Power Purchase Agreement Negotiation 144		
A.1	Equilibrium Outcomes Under Base Model	144
A.2	Alternative Bargaining-Power Assumptions	144
A.3	Kalai-Smorodinsky Negotiation Framework	149
A.4	Equilibrium Outcomes with Risk-Averse Buyers	151
A.5	Heuristic REC Banking Policies for Risk-Averse Buyers	152
A.6	Proof of Statements	154
A.7	Numerical Study	166
Appendix B: Supplementary Material for Chapter 3: Battery Dispatch in Electricity Markets with Negative Prices 173		
B.1	Proofs of Propositions and Corollaries	173
B.2	Holt-Winters model embedding	183
B.3	Deterministic LP benchmark (certainty-equivalent and perfect foresight).	186
Appendix C: Supplementary Material for Chapter Chapter 4: Worker Schedule Quality Disparities 187		
C.1	Robustness Analysis	187
C.2	Estimating Worker Skill Levels	190
C.3	Earning Potentials By Week-Hour	194
C.4	Sales Performance Analysis Procedure	195
C.5	Discordance with Manager Background	197

LIST OF FIGURES

Figure Number		Page
2.1	Illustration of a Virtual PPA Contract	11
2.2	Illustration of Negotiation Mechanisms	23
2.3	Mechanism Viability and Buyer-Group Preference Map	36
3.1	Bridge-to-real via residual block bootstrap	78
3.2	<i>Ex ante</i> regime analysis with recomputed duration-specific deterministic bench- mark values	82
3.3	Profit impact of two approximation choices relative to the proposed DP. . . .	84
A.1	Electricity and REC Price Series	168
A.2	Equilibrium Project Size Difference Relative to RN No-Banking Setting . . .	169
A.3	Mechanism Viability and Buyer-Group Preference Maps	171
C.1	Expected Hourly Tip Earning Potential (\$/Worker) for Day-of-Week and Hour-of-Day Combinations.	195

LIST OF TABLES

Table Number		Page
3.1	Baseline DP under independent one-step shocks versus LP benchmarks. . . .	77
3.2	Main real-data comparison: dependence-aware DP versus LP benchmarks. . .	79
3.3	Computational footprint of the dependence-aware DP on CAISO real data. .	80
3.4	Rolling CART switching rule on real data.	83
4.1	Summary Statistics of the Variables at Worker-Store-Week Level.	104
4.2	Sufficiency and Last-minute Addition Disparities.	107
4.3	First-Stage Regression Results.	111
4.4	Productivity for Different Shift Types.	113
4.5	Differences in Hourly Earning Potential for RTA and SNA Shifts Relative to Regular Shifts.	115
4.6	Actually Worked Sufficiency.	118
4.7	Schedule Disparities Using 3-yr Full Sample	121
A.1	Equilibrium outcomes under different negotiation mechanisms	145
A.2	Equilibrium outcomes with risk-averse buyers	152
A.3	Annual Price Series Summary Statistics	168
C.1	Schedule Disparities Among Server-Only Workers.	188
C.2	Schedule Disparities Using Namsor Algorithm.	190
C.3	Schedule Disparities Using NamPrism Algorithm.	191
C.4	Schedule Disparities Using 3-yr Full Sample	192
C.5	Schedule Disparities Using 3-yr Full Sample.	193
C.6	Moderating Effects of Gender and Race Discordance with Manager.	199

ACKNOWLEDGMENTS

I am deeply grateful to my co-chairs, Professor Shi Chen and Professor Michael Wagner, for their guidance, patience, and unwavering support throughout my PhD. They have shaped not only this dissertation, but also the way I think about research and about myself as a researcher. Through their insight, encouragement, and example, I learned how to ask better questions, navigate ambiguity, work through challenges, and keep going even when the path forward was unclear. I am especially thankful for the confidence they placed in me at moments when I struggled to find it in myself.

I would also like to express my sincere gratitude to Professor Qiuping Yu and Professor Kamran Moinzadeh, who have been far more than committee members and collaborators. I have had the privilege of working closely with all four of them, and I have learned from each of them what it means to do thoughtful, rigorous, and meaningful research. Their questions, feedback, and perspectives strengthened this work in important ways, and their mentorship has left a lasting influence on how I approach research.

I am particularly grateful to Shi, Mike, Qiuping, and Kamran for their kindness and understanding throughout this journey. Their support made it possible for me to navigate the dual roles of PhD student and mother with far more encouragement and grace than I could have expected. At many points, they reminded me that I belonged here, that the work mattered, and that it was possible to continue even when graduate school and life felt overwhelming. I feel very fortunate to have learned from and been supported by them.

I am fortunate to have learned from and worked alongside many colleagues and friends at Foster. Thank you for the discussions, technical help, encouragement, and moments of humor that made this journey more meaningful. I am especially grateful to Haonan Zhang for his constant camaraderie, thoughtful conversations about research, emotional support, and countless moments of shared commiseration and laughter. I would not have made it

through this journey without you alongside me.

I would also like to thank Beau Kirkeby and Nuzulita Budhiari for their tremendous administrative support throughout my PhD. From teaching assignments and funding reimbursements to the many other administrative details that arise along the way, they handled everything with care, patience, and generosity. Because of their help, I never felt burdened by administrative challenges, and their warmth and support made Foster feel like home.

I gratefully acknowledge the support of the Foster School of Business at the University of Washington. This work would not have been possible without the resources, opportunities, and community that Foster provided.

Finally, I thank my family for their unwavering love, patience, and belief in me. Their encouragement sustained me through the challenges of graduate school and continually reminded me of what matters beyond the dissertation. I am especially grateful to my husband, Ran Fu, for being my steady source of emotional support, especially in difficult moments, and for carrying so much of our family life when I was unable to.

I am also deeply grateful to my parents-in-law, Ping Gao and Hong Fu, and to my parents, Aiping Zhu and Bingrang Hou, for stepping in to care for Ken and to help keep our home running. Completing a PhD while becoming a new mother was not easy, and their generosity - with their time, energy, and love - gave me the space to breathe, think, and continue doing research.

I am also deeply grateful to my son, Ken. Even when he needed me most at the very moments I was trying to work, his laughter brought a kind of relief and joy that nothing else could. Watching him grow from a baby into his own little person has been the most emotionally and intellectually inspiring experience of my life. My family has reminded me that, as fascinating and meaningful as research can be, life is far richer and larger than research alone.

DEDICATION

To Ran and Ken,
for filling my life with love, laughter, and meaning beyond this dissertation.

To my parents and parents-in-law,
for their unwavering support, generous care, and constant belief in me.

Chapter 1

INTRODUCTION AND OVERVIEW

Sustainability has become a central concern for firms, workers, regulators, and society. In operations management, sustainability is not only a question of long-term goals or high-level commitments. It is also shaped by the everyday design of procurement contracts, the use of information in operational control, and the allocation of flexibility across stakeholders. A firm may commit to renewable electricity, but whether that commitment leads to new renewable capacity depends on how electricity is procured and how risks and bargaining power are allocated among contracting parties. A power system may invest in battery storage, but the value of storage depends on how batteries are dispatched under uncertain and sometimes negative electricity prices. A service firm may use flexible labor scheduling to match staffing with demand, but the consequences of that flexibility depend on whether unstable schedules are distributed equitably across workers.

This dissertation studies these issues through three essays on sustainable operations management. The first two essays focus on environmental sustainability in energy and electricity markets. The first essay examines how multiple corporate buyers can jointly procure renewable energy through collective power purchase agreements. The second essay studies how batteries should be dispatched in electricity markets with negative prices when operators rely on forecasts that evolve over time. The third essay turns to social sustainability in service operations by studying gender and racial disparities in the temporal quality of workers' schedules.

Although the three essays examine distinct empirical and institutional settings, they share a common theme: sustainable outcomes depend on operational mechanisms that align incentives, use information effectively, and allocate risk and flexibility appropriately. Environmental and social goals do not implement themselves. They must be translated into operational decisions, such as who participates in a renewable procurement contract,

how a storage asset responds to uncertainty, and how managers distribute work hours and last-minute schedule changes. Across the three essays, I study how such decisions affect capacity expansion, market performance, and equity.

1.1 Sustainable Operations as Operational Design

Operations management traditionally studies how firms design processes, allocate resources, and make decisions under uncertainty. Sustainable operations expands this perspective by asking how those same decisions affect environmental and social outcomes. In environmental settings, operational decisions influence resource use, emissions, renewable energy adoption, and the integration of low-carbon technologies. In social settings, they influence worker well-being, fairness, income stability, and access to opportunity. These outcomes are often shaped not by a single decision maker, but by interacting stakeholders whose objectives and constraints differ.

This dissertation emphasizes three recurring frictions in sustainable operations. The first is strategic alignment. Many sustainable investments require coordination among parties that may benefit differently from the same operational arrangement. Collective renewable procurement, for example, can make power purchase agreements accessible to smaller buyers, but only if the chosen mechanism gives every participant a reason to join. The second friction is uncertainty. Sustainable technologies and practices often operate in volatile environments, such as electricity markets with uncertain and negative prices or service systems with fluctuating demand and worker availability. The third friction is allocation. Operational flexibility can create value, but the costs and benefits of flexibility may not be distributed evenly. In workforce scheduling, the ability to adjust labor close to the time of service may help firms manage demand uncertainty, but it may also impose income instability and unpredictability on workers.

The essays in this dissertation address these frictions using complementary methods. The first essay develops analytical bargaining models to study renewable energy procurement. The second essay develops a dynamic programming framework and evaluates it through synthetic experiments and real electricity-market data. The third essay uses granular field data from a restaurant chain to empirically examine schedule quality disparities and their

potential mechanisms. Together, these essays contribute to a broader understanding of how operational decisions can support sustainability when firms face strategic interaction, uncertainty, and competing stakeholder objectives.

1.2 Essay 1: Collective Renewable Energy Procurement

The first essay, “Collective Power Purchase Agreement Negotiation,” studies how small and medium-sized corporate buyers can procure renewable electricity through collective power purchase agreements. Corporate decarbonization pledges have increased demand for renewable electricity, and long-term power purchase agreements have become an important mechanism for financing renewable energy projects. Under a power purchase agreement, a buyer commits to purchasing renewable energy from a project developer at a fixed price over a long horizon, often ten to fifteen years. This contractual commitment can provide the revenue stability needed to finance new renewable capacity.

However, signing a power purchase agreement can be difficult for small and medium-sized corporate buyers. These buyers may not have enough electricity demand to support a full project on their own, may lack the expertise needed to negotiate a complex energy contract, or may have limited bargaining power relative to a project developer. Collective power purchase agreements have emerged as a potential solution. In these arrangements, multiple buyers aggregate their demand and jointly support a renewable project. Yet collective procurement introduces its own governance problem: the arrangement must be individually attractive to each buyer, and the buyers must agree on how to negotiate with the seller.

This essay compares three procurement architectures. In the separate-negotiation benchmark, each buyer negotiates independently with the renewable project developer. In joint negotiation, the buyers form a consortium and negotiate collectively with the seller. In agent-based negotiation, one buyer acts as an agent or anchor tenant, first contracting with the other buyer and then negotiating with the seller on behalf of the group. The essay develops a generalized Nash bargaining framework to compare these mechanisms in terms of project size, contract prices, surplus allocation, and buyer participation.

The analysis shows that, under risk neutrality, all three negotiation mechanisms induce the same equilibrium renewable project size. This result implies that, in the base model, the

choice of negotiation mechanism does not change renewable capacity expansion. Instead, it changes how the surplus created by the project is allocated among the buyers and the seller. This distinction clarifies the central challenge in collective renewable procurement: the key issue is not only whether the collective agreement creates value in aggregate, but whether the value is allocated in a way that makes every buyer willing to participate.

The essay then characterizes when the two collective mechanisms are viable relative to separate negotiation. Joint negotiation pools buyer bargaining power before the seller-facing negotiation and tends to be the broader participation-safe mechanism. Agent-based negotiation has a tighter implementability requirement because the agent must carry the remaining buyer-side surplus into the seller-facing negotiation and must therefore receive enough seller-facing leverage to participate. However, when agent-based negotiation is viable, it can outperform joint negotiation from the buyer group’s perspective. Its sequential structure allows the follower to lock in part of the surplus before the seller-facing negotiation, shielding that share from seller extraction. This protection is valuable when buyers are weak or asymmetric, but it comes at the cost of delegation loss because only the agent bargains with the seller.

The essay also studies operational extensions. When buyers are risk averse, the equilibrium project size may increase or decrease relative to the risk-neutral benchmark. A power purchase agreement can hedge renewable energy certificate price risk and, when project-market prices co-move with local retail electricity prices, can also hedge electricity-cost risk. At the same time, a larger contracted quantity exposes buyers to additional wholesale and renewable energy certificate price volatility. The direction of the project-size effect depends on which force dominates. Allowing buyers to bank renewable energy certificates further increases the marginal value of renewable output and weakly increases equilibrium project size. These results show that operational contract features, risk profiles, and governance mechanisms should be designed jointly rather than treated as separate decisions.

1.3 Essay 2: Battery Dispatch in Electricity Markets

The second essay, “Battery Dispatch in Electricity Markets with Negative Prices: Forecast-Embedded Dynamic Programming,” studies how battery storage should be operated in

electricity markets where prices can be negative and forecasts evolve over time. Battery storage is widely viewed as an important technology for integrating renewable energy into power systems. By charging when electricity is abundant and discharging when it is scarce, batteries can help absorb renewable generation, reduce curtailment, and support system flexibility. However, realizing this value requires effective dispatch decisions.

Electricity prices are volatile, and in markets with high renewable penetration they can become negative. Negative prices create both opportunities and challenges for battery operators. A battery may profit by charging when prices are negative, but its ability to do so depends on its state of charge, efficiency, storage duration, and expectations about future prices. In practice, operators often use rolling deterministic reoptimization with updated point forecasts. This approach is computationally convenient and easy to implement. Stochastic dynamic programming, by contrast, can explicitly account for uncertainty but is often viewed as too cumbersome for realistic storage settings.

This essay develops a tractable dynamic programming framework that embeds the forecasting state directly into the dispatch problem. The method solves the state-of-charge dimension exactly on an endogenous event grid, so approximation is confined to the exogenous price-information process. This structure makes it possible to compare stochastic control with simpler deterministic approaches in a way that is both analytically grounded and operationally relevant.

The essay evaluates the proposed policy using matched synthetic experiments and out-of-sample tests in California’s real-time electricity market. It compares the dynamic programming policy with open-loop certainty-equivalent planning and receding-horizon deterministic reoptimization. The results show that stochastic control creates substantial value relative to open-loop planning. However, its advantage over rolling deterministic reoptimization is selective rather than universal. In real data, performance depends critically on modeling dependence in forecast errors. When this dependence is ignored, deterministic reoptimization remains difficult to beat. When dependence is incorporated, dynamic programming becomes more competitive, especially for longer-duration batteries.

The essay further identifies operational conditions under which stochastic look-ahead is most valuable. In particular, stochastic control is more useful when the deterministic value

signal is weak and forecast uncertainty is high. These same signals can support a switching rule that chooses between stochastic and deterministic controllers and outperforms always using either one. Thus, the essay does not make a blanket argument that more sophisticated stochastic control is always better. Instead, it identifies when the additional modeling effort is operationally justified.

This essay contributes to sustainable operations by showing how information and control affect the realized value of energy storage. Investment in storage capacity is important, but capacity alone is not enough. The environmental and economic value of storage also depends on dispatch policies that can respond effectively to uncertainty. By clarifying when stochastic look-ahead matters, the essay provides guidance for operators deciding how much modeling sophistication is warranted in practical battery dispatch.

1.4 Essay 3: Workforce Equity in Service Operations

The third essay, “Explaining Worker Schedule Quality Disparities: Evidence from an American Restaurant Chain,” studies social sustainability in service operations. Many service firms use just-in-time scheduling to match labor capacity with fluctuating demand. Such scheduling practices can improve labor flexibility for the firm, but they may create unpredictable and insufficient work hours for employees. For hourly workers, schedule quality is an important dimension of job quality because it affects income stability, work-life coordination, childcare arrangements, and the ability to hold additional employment.

This essay examines whether schedule quality is distributed equitably across workers. Prior research on labor market disparities has focused extensively on hiring, wages, benefits, and promotions, but the temporal quality of work schedules has received less attention. This gap is important because scheduling is a core operational decision in service firms, and because workers may experience inequality not only in how much they are paid, but also in when and how reliably they are given work.

The essay uses large-scale field data from an American casual restaurant chain. The data include nearly half a million offered and worked shifts for more than 2,400 customer-facing workers, along with transaction records and customer reviews. The analysis focuses on two dimensions of schedule quality. The first is sufficiency, measured by the number of hours

workers are offered. The second is predictability, measured by the frequency of last-minute additions, including short-notice shift additions and real-time extensions.

The empirical analysis compares scheduling outcomes across gender and racial groups while accounting for a rich set of human capital and operational factors. These include role, tenure, absenteeism, store and time effects, and multidimensional measures of worker ability based on sales performance, service speed, and customer satisfaction. After accounting for these factors, the essay finds that women and minority workers receive fewer initially scheduled hours and are more likely to experience last-minute additions compared to White male workers. Human capital explains part of the observed disparities, but sizable gaps remain.

The essay then investigates whether these remaining gaps can be explained by worker preferences or availability. Using transaction-level data, the analysis compares performance and earning potential across regular shifts and last-minute additions. The results suggest that workers do not prefer last-minute additions over regular shifts; real-time additions are associated with lower sales performance, and last-minute shifts tend to occur during hours with lower earning potential. The analysis of offered and worked hours further suggests that lower availability explains part of the gap for White female workers but does not explain the gaps for minority workers. Overall, the evidence points to systematic under-scheduling and unequal exposure to unpredictable work.

This essay contributes to sustainable operations by highlighting the social consequences of operational flexibility. Flexible scheduling can help firms respond to demand uncertainty, but the costs of that flexibility may be borne disproportionately by certain groups of workers. By documenting disparities in both schedule sufficiency and predictability, and by examining mechanisms that are difficult to observe in survey data, the essay shows how operational data can be used to study workplace equity.

1.5 Cross-Essay Themes

The three essays are connected by several common themes.

First, sustainable operations require attention to participation and implementation. In the renewable procurement setting, a collective power purchase agreement must make every

buyer at least as well off as separate negotiation. In the battery dispatch setting, a stochastic control method must create enough value to justify its modeling and computational effort relative to simpler reoptimization. In the scheduling setting, a flexible labor system must be evaluated not only by its ability to match staffing with demand, but also by whether the resulting schedules are acceptable and equitable for workers. Across all three settings, a sustainable solution must be operationally implementable for the stakeholders involved.

Second, sustainability depends on how uncertainty is managed. In renewable procurement, uncertainty in wholesale electricity prices, renewable energy certificate prices, and retail electricity costs affects the value and risk of power purchase agreements. In battery dispatch, forecast uncertainty and forecast-error dependence determine the value of stochastic look-ahead. In workforce scheduling, demand uncertainty and attendance uncertainty motivate last-minute schedule adjustments, but these adjustments can impose instability on workers. The essays show that uncertainty is not merely a background condition; it shapes which operational mechanisms are effective and who bears the consequences.

Third, the allocation of flexibility is central to sustainable operations. In the PPA essay, collective procurement reallocates bargaining flexibility and surplus among buyers and the seller. In the battery essay, storage flexibility is allocated over time through charging and discharging decisions. In the scheduling essay, labor flexibility is allocated across workers through offered hours and last-minute additions. In each case, flexibility creates value only when it is governed appropriately. Poorly designed flexibility can lead to participation failures, suboptimal operational performance, or inequitable outcomes.

Finally, the essays illustrate the value of combining analytical and empirical approaches. Analytical models clarify mechanism design tradeoffs in collective renewable procurement. Dynamic programming and numerical experiments identify when stochastic battery dispatch is valuable. Granular field data reveal disparities in scheduling practices that would be difficult to observe through aggregate data or surveys. Together, these methods provide complementary evidence on how sustainable operations can be designed and evaluated.

1.6 Organization of the Dissertation

The remainder of the dissertation is organized as follows.

Chapter 2 studies collective power purchase agreement negotiation. It develops a generalized Nash bargaining model to compare separate negotiation, joint negotiation, and agent-based negotiation for corporate renewable electricity procurement. The chapter characterizes equilibrium project size, buyer-wise viability, and mechanism preference, and studies how risk aversion, price correlation, and renewable energy certificate banking affect the results.

Chapter 3 studies battery dispatch in electricity markets with negative prices. It develops a forecast-embedded dynamic programming framework and compares its performance with open-loop certainty-equivalent planning and rolling deterministic reoptimization. The chapter identifies when stochastic look-ahead creates operational value and proposes signals that can guide switching between stochastic and deterministic controllers.

Chapter 4 studies worker schedule quality disparities in service operations. Using detailed shift, transaction, and customer-review data from a restaurant chain, it documents gender and racial disparities in schedule sufficiency and predictability. The chapter further examines whether human capital, worker preferences, or availability explain these disparities and discusses implications for equitable labor scheduling.

Together, these chapters show that sustainable operations management requires more than optimizing efficiency under a narrow objective. It requires designing operational mechanisms that expand environmental benefits, use information effectively, and distribute operational flexibility fairly. By studying renewable energy procurement, battery storage dispatch, and workforce scheduling, this dissertation contributes to a broader understanding of how operations management can support environmental and social sustainability.

Chapter 2

COLLECTIVE POWER PURCHASE AGREEMENT NEGOTIATION

This chapter is based on joint work with Shi Chen and Kamran Moinzadeh.

2.1 Introduction

Corporate decarbonization pledges are accelerating demand for renewable electricity, yet many firms-especially small and mid-sized buyers-struggle to contract new supply. Programs such as RE100 document hundreds of companies committing to 100% renewable electricity [RE100, 2023], but many procurement strategies still rely on instruments such as renewable energy certificates (RECs) or energy attribute certificates (EACs) that do not necessarily finance new generating capacity. Direct ownership of generation assets, by contrast, is capital intensive and operationally complex. Corporate Power Purchase Agreements (PPAs) have therefore emerged as a middle path. Under a PPA, the buyer commits to a fixed unit price for a fixed term, and the project developer uses the contracted revenue stream to secure financing and build capacity. In practice, most corporate buyers use a virtual PPA (vPPA), a financial contract-for-difference that settles against wholesale market prices and allocates the associated RECs to the buyer [Kansal, 2018]. Figure 2.1 illustrates this cash-flow logic. Throughout, we use “PPA” to refer to virtual PPAs.

Large buyers increasingly execute PPAs independently, but smaller buyers often lack the demand scale, internal expertise, or balance-sheet capacity to support a standalone transaction. This has motivated *collective* procurement structures in which multiple buyers aggregate demand and share the output of a larger renewable project. In practice, these arrangements typically take one of two forms: an agent-based (or anchor-tenant) structure, where one buyer contracts with the others and negotiates with the seller on behalf of the group; or a joint (consortium) structure, where buyers form a coalition to negotiate collectively with the seller [Zanchi et al., 2017, PV Magazine, 2021, Greenbiz, 2022]. While these




Figure 2.1: Illustration of a Virtual PPA Contract

Note: The buyer pays the seller the contract price and receives the market value and RECs generated by the project. The buyer purchases physical electricity from its local utility at the local market price. Because the renewable electricity does not need to be physically delivered to the buyer, the buyer and seller can be connected to different electricity grids. If the project generates more (less) RECs than needed, the buyer can sell (purchase) RECs in the REC market. Source: Kansal [2018]

structures theoretically make PPAs accessible to smaller buyers, their real-world adoption has lagged significantly behind single-buyer PPAs.

In this paper, we argue that the bottleneck to collective PPA adoption is a challenge of strategic alignment in twofold. First, the collective arrangement must be individually rational, meaning every buyer must weakly prefer the coalition over negotiating separately; and second, the buyers must agree on the optimal procurement mechanism for the group. Neither of these requirements is trivially satisfied in practice.

We structure our analysis to directly address these two questions using a generalized Nash bargaining framework featuring two buyers and a single renewable project developer. Our model compares three distinct procurement architectures: separate negotiation (the baseline), joint negotiation, and agent-based negotiation. By jointly determining the project

size, contract prices, and surplus distribution, we isolate how the structural rules of each mechanism dictate both the equilibrium project size of the PPA and the strategic preferences of the individual buyers.

Our analysis is organized around three questions. First, do negotiation mechanisms change the equilibrium renewable project size? Second, when do joint and agent-based negotiation make both buyers better off than separate negotiation? Third, conditional on buyer-wise viability, when should buyers use agent-based negotiation versus joint negotiation? We then examine how operational factors including buyer risk aversion, serial correlation in energy-related prices, and REC banking modify these conclusions.

Our first contribution is a unified bargaining framework for comparing collective PPA mechanisms. We establish that under risk neutrality, separate, joint, and agent-based negotiations all yield the identical equilibrium project size. This capacity-neutrality result ensures that all negotiation structures generate the same total surplus, differing only in how that surplus is allocated among the buyers and the seller. Consequently, the mechanism design problem is fundamentally about surplus allocation rather than altering the efficient project size.

Our second contribution characterizes the buyer-wise viability of the two collective mechanisms and compares their performance. We define a collective mechanism as viable only if both buyers weakly prefer it to separate negotiation. Under joint negotiation, the buyers first pool their bargaining power to extract a collective surplus from the seller, and subsequently divide this retained surplus internally. Consequently, the mechanism is viable if and only if this internal allocation provides both buyers with a payoff that exceeds their respective separate-negotiation alternatives. Agent-based negotiation, by contrast, faces a higher implementability hurdle due to its asymmetric role assignment. Because the follower locks in their surplus internally before facing the seller, the agent must bear the friction of carrying the remaining buyer-side surplus into the final negotiation. Consequently, the agent must secure enough seller-facing leverage to justify serving as the anchor, making the viability conditions for agent-based negotiation strictly tighter. However, when both collective mechanisms are viable, agent-based negotiation can outperform joint negotiation. Its sequential structure allows the follower to shield their locked-in share from seller extraction.

Yet this protection comes at the cost of a delegation loss; specifically, because only the agent bargains with the seller, the group’s effective bargaining leverage may be lower than its fully pooled potential as in joint negotiation. Thus, agent-based negotiation is preferred precisely when the value of shielding the follower’s share outweighs this delegation loss, which typically occurs when the buyers are weak or have imbalanced bargaining powers. Conversely, joint negotiation dominates when the coalition’s pooled bargaining power is strong, particularly for balanced, high-power buyer groups.

Our third contribution is to show how operational features of renewable procurement alter project size while preserving the mechanism-choice logic. With mean-variance preferences, risk aversion can either *increase or decrease* project size. This is because on the one hand, a PPA can hedge REC-price risk and, when wholesale market prices correlate with local retail costs well, can hedge electricity-cost risk; on the other hand, larger contracted quantities expose buyers to additional wholesale and REC price volatility. Serial correlation in prices changes the variance and covariance terms in the risk-adjusted utility but preserves the same bargaining structure. In contrast, REC banking unambiguously increases project size by contributing a positive marginal banking value to the buyers’ payoffs. Numerical experiments calibrated to market data validate these project-size results and confirm that the mechanism-viability and preference maps exhibit largely the same structural patterns as in the main model results.

Our results offer several managerial implications for corporate buyers, project developers, and procurement intermediaries. First, collective PPAs should not be evaluated only by their aggregate value or contracted capacity; buyers must also verify whether the proposed mechanism improves each participant’s payoff relative to separate negotiation. This participation requirement is especially important for anchor-tenant structures, where the buyer serving as agent may bear a disproportionate share of the seller-facing negotiation burden. Second, the choice between joint and agent-based negotiation should be guided by the buyers’ bargaining-power profile. Joint negotiation is the more robust structure when buyers have sufficiently strong and balanced bargaining positions, while agent-based negotiation is more attractive when buyers are weaker or asymmetric. Third, operational contract features should be incorporated at the mechanism-design stage rather than treated as afterthoughts.

Risk aversion can either expand or shrink the efficient project size depending on the strength of the PPA’s hedging value, while REC banking systematically increases the marginal value of capacity. Thus, buyers and intermediaries can improve collective PPA implementation by matching buyers with compatible valuation and risk profiles, choosing the governance structure that satisfies buyer-wise participation, and incorporating REC banking provisions when feasible.

2.2 Literature Review

Our work connects to two main streams of literature in operations and supply chain management, namely socially and environmentally sustainability and group buying.

2.2.1 Sustainability and Renewable Energy

Our study falls within the broad field of socially and environmentally sustainable supply chain management. Notable examples include investigating the impact of supply chain transparency (Kraft et al. 2018, Kalkanci and Plambeck 2020, Chen et al. 2019), encouraging responsible sourcing (Chen and Lee 2017, Agrawal and Lee 2019), and focusing on climate-related strategies (Caro et al. 2013, Sunar and Plambeck 2016, Gopalakrishnan et al. 2021, Dong 2021). For an extensive review of this broad literature, see Lee and Tang [2018] and Sunar and Swaminathan [2022].

Our focus on collective PPA contracts intersects with sustainability through the lens of renewable energy procurement. The existing literature examines this topic typically from two viewpoints: the policy/market designer and the market participants.

(1) Policy makers and market designers are concerned with eliminating barriers and incentivizing as much economically sensible renewable energy procurement as possible. For instance, Siddiqui et al. [2016] demonstrate that the optimal renewable portfolio standards (RPS) target of a perfectly competitive electricity market is higher than that of a centrally planned market, but an overly high RPS target may lead to an economic distortion that excessively compensates renewable energy. Khazaei et al. [2017] investigate how the REC price should be set by using a stochastic dynamic programming model to simulate market players’ behavior. Boomsma et al. [2012] compare the effectiveness of Feed-in-Tariff (FIT)

and REC trading by modeling renewable energy investment timing and capacity choices with a real options approach. Zhou et al. [2016] and Zhou et al. [2019] investigate optimal firm trading strategies for storage-renewable pairs and show that allowing for negative market price (or not) significantly affects the financial performance of renewable resources. Additionally, carbon pricing or taxation is another mechanism widely studied and debated in both the industry and academia (see, e.g., Aflaki and Netessine 2017). In addition to examining policies directly targeting renewable procurement, another research stream investigates how market designs affect renewable procurement indirectly. Wu and Kapuscinski [2013] find that allowing renewable curtailment decreases system costs and emissions. K  k et al. [2018] find that flat and peak pricing can provide different incentives to various resources. Sunar and Birge [2019] assert that implementing an under-supply penalty enhances renewable energy commitment but reduces system reliability. Sunar and Birge [2019] and Ghiassi-Farrokhfal et al. [2023] investigate market participants’ responses to the prevalent day-ahead versus real-time market design in many competitive electricity markets.

(2) Unlike policy and market designers who focus on social welfare, market participants’ objective is to maximize profits when contemplating investments in renewable energy assets. In a survey, Agrawal and Y  cel [2021] divide renewable energy market participants into three categories: (i) utility firms, (ii) commercial and residential consumers, and (iii) corporations. The majority of existing research focuses on the optimal investment strategies for utility firms and consumer buyers. For instance, K  k et al. [2020] model a utility firm’s decisions as a two-stage problem where the utility first determines the level of investment and then the operations of the generating capacity. Sunar and Swaminathan [2021] examine how distributed renewable energy resources, such as residential rooftop solar panels, affect utilities’ financial performance. When considering the investment strategies of consumers, the focus shifts to whether a specific renewable project is financially viable (see, e.g., Hu et al. 2015, Goodarzi et al. 2019).

Corporate buyers blend the characteristics of utility firms, which directly negotiate with generation suppliers, and end-use consumers who emphasize the financial results of renewable projects. Unlike utilities, these corporate buyers often operate without the burden of extensive regulatory constraints, thereby enjoying more flexibility in contracting mech-

anisms. Therefore, corporate renewable procurement holds significant promise to channel private sector investment and accelerate the adoption of renewable energy. Interestingly, despite the rapid increase in renewable energy procurement by corporate buyers and its great potential, scant attention has been paid to their sourcing strategies for renewable energy (Agrawal and Yücel 2021).

Notable exceptions include Trivella et al. [2023], with a focus on how corporations can meet their renewable targets by evaluating the dynamic trade-off between purchasing RECs from the spot market and engaging in PPA contracts, reassuring the great potential of PPA contracting for corporate buyers. As another rare example, Gao et al. [2023] investigate the optimal timing for a firm to sign a PPA contract, suggesting that tax reduction policies may backfire and reduce the future renewable capacity addition. Finally, a few more examples come from the field of energy research. For instance, Mendicino et al. [2019] discuss variations of corporate PPA structures and demonstrate how a corporate PPA can be priced based on a detailed breakdown of cost factors. O’Shaughnessy et al. [2021] discuss how corporate renewable procurement can impact grid operation and propose pathways to incorporate corporate procurement into grid planning.

The fragmented research on corporate renewable procurement underscores a notable gap in exploring the challenges confronted by corporate buyers. In particular, relatively little research has been done concerning the collective procurement strategies of multiple buyers under the PPA contracts. Therefore, we aim to fill in the gap by centering our research on corporate buyers who negotiate PPA contracts collectively.

2.2.2 Group Buying

Our work relates to several streams in the group-buying literature.

A first stream studies joint procurement among competing downstream buyers that source from a common upstream supplier and then interact in a product market (e.g., Chen and Roma 2011, Keskinocak and Savaşaneril 2008, Fu et al. 2024). These models typically assume that the supplier posts an exogenous price or quantity-discount schedule, and buyers play a noncooperative quantity game. The analytical challenge comes from downstream

product-market competition rather than from uncertainty in input prices. Consequently, while this literature sheds light on how joint purchasing affects strategic output and pricing, it largely abstracts from multilateral buyer–seller bargaining and treats supplier pricing as fixed.

A related set of noncooperative models considers strategic buyers without downstream competition. These papers emphasize bargaining and pooling incentives but still rely on exogenous supplier pricing. For example, Hu et al. [2012] study posted quantity discounts combined with group purchasing organizations (GPO)-set on-contract prices and membership fees; Hu et al. [2013] show that pooling may reduce buyer surplus when a sole-source supplier extracts rents via type-dependent contracts; and Chae and Heidhues [2004] show how alliances across independent markets improve buyers’ fallback positions. Despite highlighting informational and bargaining frictions, this stream does not model a joint determination of contract terms by multiple buyers and the seller.

Another stream employs cooperative frameworks to examine coalition formation, stability, and cost allocation in buyer groups and GPOs (e.g., Nagarajan and Sošić 2007, 2009, Nagarajan et al. 2010, Slikker et al. 2001, Elomri et al. 2012, Zhou et al. 2019). Supplier terms are encoded in an exogenous value function, and the analytical focus lies on internal coordination rather than on buyer–seller negotiation. Even when bargaining is modeled explicitly, as in Hsu et al. [2017], where a designated leader negotiates a wholesale price with the supplier while buyers later choose quantities, the representation remains agent-based, not one in which multiple buyers participate directly as bargaining players with respect to the seller or bargain jointly over both price and quantity. In addition, this literature focuses on leader-based collaboration and does not perform a unified comparison against a symmetric joint-form coalition. Moreover, although some cooperative or hierarchical models incorporate demand uncertainty, they do not capture correlated multi-market price risks or the hedging motives central to renewable-energy procurement. Closest in spirit is Wang et al. [2021], who compare direct versus agent sourcing in a three-player bargaining model, but their setting involves a single buyer, does not link project feasibility to buyer participation, and centers on demand rather than multi-market price volatility.

Across these streams, several structural simplifications become consequential in the cor-

porate PPA context: (i) supplier pricing is taken as exogenous rather than endogenously bargained; (ii) the supplier’s disagreement payoff does not depend on whether buyers coordinate; and (iii) uncertainty arises from demand or inventory, not from jointly stochastic and correlated electricity, REC, and retail prices.

Our contribution is to fill these gaps by modeling collective corporate PPAs as a multi-lateral bargaining problem in which buyers and the project developer jointly determine both contract price and project size. Because projects are often not financially viable without buyer participation, buyer coalitions in PPAs are project-enabling, generating bargaining dynamics fundamentally different from those in standard group-buying models where the supplier can supply unilaterally. Our framework also incorporates correlated market-price uncertainty and accommodates both risk-neutral and risk-averse buyers, allowing us to show how hedging motives shape coalition formation and negotiation outcomes. Within a unified structure, we compare three negotiation architectures—separate, agent-based, and joint coalition bargaining—and characterize how each affects project size, contract prices, and surplus allocation. Together, these results provide the first comprehensive theoretical foundation for collective corporate PPAs, a setting outside the scope of existing supply-chain and group-purchasing literatures.

Finally, our work is also related to the literature on supply chain risk mitigation through contract design (Kamrad and Ritchken 1994, Kamrad and Siddique 2004, Hou et al. 2021, Dong et al. 2022) and financial hedging (see, e.g., Kouvelis et al. 2019 and the references therein).

2.3 The Base Model

In this section, we present our base model focusing on risk-neutral buyers negotiating with a seller under three distinct negotiation mechanisms, namely separate negotiation, joint negotiation, and agent-based negotiation. We begin in Section 2.3.1 by formulating each player’s bargaining gain and outside option. Section 2.3.2 then introduces the mechanisms and formalizes the corresponding bargaining problems. We subsequently summarize the closed-form equilibrium solutions and highlight the main analytical insights in Section 2.3.3.

2.3.1 Base Formulation for Individual Entities

To set the foundation, we first derive the expected profits of a buyer i (he) and a seller s (she) both with and without a PPA contract. Since buyer i and the seller may hold different beliefs about these random variables, we let $\mathbb{E}_i[\cdot]$ denote expectation under buyer i 's subjective belief and $\mathbb{E}_s[\cdot]$ denote expectation under the seller's subjective belief.

Suppose buyer i and seller s sign a PPA contract of output size q_i per period at contract price v_i per unit of output for T periods. Consider period t during the contractual period, $t \in \{1, \dots, T\}$. The project's unit revenue (i.e., the actual price of the energy market where the project is located) in period t is a random variable p_t , thereby leading to a profit of $(p_t - v_i)q_i$ from the energy market for the buyer. Suppose that buyer i has a renewable energy target level g_{it} , which is met by claiming the RECs associated with the project. When there is a surplus of REC generation in period t , i.e., $q_i > g_{it}$, the buyer can sell the extra RECs to the REC market and make a profit of $r_t(q_i - g_{it})$, where r_t is the unit price of REC in period t . On the contrary, if the project's RECs are insufficient, i.e., $q_i < g_{it}$, then $r_t(q_i - g_{it})$ is negative, implying that he needs to buy RECs from the REC market to fill the gap. The buyer pays his local utility $c_{it}d_{it}$ in period t for his electricity demand, d_{it} from the local energy market at a retail price of c_{it} . In sum, the annual cash flow in period t for buyer i during the contractual period can be written as $(p_t - v_i)q_i + r_t(q_i - g_{it}) - c_{it}d_{it}$.

When buyer i is not under a PPA contract, he buys electricity from his local utility and meets the entirety of his renewable target through the REC market. Thus, his cash flow in period t is simply $-c_{it}d_{it} - r_tg_{it}$. To sum up all factors, the Net Present Value (NPV) of buyer i 's cash flow over the contract horizon is characterized as

$$\pi_i(q_i, v_i) = \sum_{t=1}^T \gamma_i^t [(p_t - v_i)q_i + r_t(q_i - g_{it}) - c_{it}d_{it}] = \sum_{t=1}^T \gamma_i^t (p_t + r_t - v_i)q_i - \sum_{t=1}^T \gamma_i^t (c_{it}d_{it} + r_tg_{it}), \quad (2.1)$$

where γ_i is the buyer's financial discount factor. In contrast, if buyer i does not engage in a PPA contract at all, the NPV of his cash flow reduces to

$$\pi_i(0, 0) = - \sum_{t=1}^T \gamma_i^t (c_{it}d_{it} + r_tg_{it}). \quad (2.2)$$

Buyer i 's risk-neutral bargaining gain from entering the contract is his subjective ex-

pected gain relative to no contract:

$$\begin{aligned}\Delta_i(q_i, v_i) &:= \mathbb{E}_i[\pi_i(q_i, v_i) - \pi_i(0, 0)] \\ &= \mathbb{E}_i\left[\sum_{t=1}^T \gamma_i^t (p_t + r_t - v_i) q_i\right] = (Q_i - G_i v_i) q_i,\end{aligned}\tag{2.3}$$

where

$$Q_i := \sum_{t=1}^T \gamma_i^t (\bar{p}_{it} + \bar{r}_{it}), \quad \bar{p}_{it} := \mathbb{E}_i[p_t], \quad \bar{r}_{it} := \mathbb{E}_i[r_t],$$

is buyer i 's subjective expected present unit value of the project output during the contractual period, before the PPA payment, and $G_i := \sum_{t=1}^T \gamma_i^t$ is buyer i 's cumulative discounting factor. In this context, $G_i v_i q_i$ is the NPV of the contract payment from buyer i 's perspective. Notice that the buyer's risk-neutral gain in (2.3) is not affected by his local electricity demand d_{it} or local retail price c_{it} , because these terms are incurred both with and without the PPA and therefore cancel out. The buyer's gain depends only on his subjective valuation of the project's wholesale-market and REC-market output, net of the contract payment.

Remark 2.1. The cross-time and cross-market dependence of (p_t, r_t, c_{it}) does not affect (2.3), since only the period means enter the risk-neutral profit calculation. This property, however, no longer holds once we introduce a mean-variance-based risk-averse buyer utility function, as discussed in Section 2.4.

The seller receives the fixed PPA payment $v_i q_i$ in each period during the contractual horizon. After the contract expires in period T and until the project's physical retirement in period L , the seller continues to operate the project and is exposed to post-contract market prices. The seller's realized NPV from the project is

$$\pi_s(q_i, v_i) = \sum_{t=1}^T \gamma_s^t v_i q_i + \sum_{t=T+1}^L \gamma_s^t (p_t + r_t) q_i - C(q_i),\tag{2.4}$$

where $C(q_i)$ is the upfront investment cost for a project of size q_i .

In the absence of a PPA contract, we assume that the seller does not undertake the project because she lacks the buyer commitment needed to support project financing. Thus, $\pi_s(0) = 0$. The seller's risk-neutral bargaining gain is her subjective expected gain:

$$\Delta_s(q_i, v_i) := \mathbb{E}_s[\pi_s(q_i, v_i) - \pi_s(0)] = (G_s v_i + Q_s) q_i - C(q_i),\tag{2.5}$$

where

$$G_s := \sum_{t=1}^T \gamma_s^t, \quad Q_s := \sum_{t=T+1}^L \gamma_s^t (\bar{p}_{st} + \bar{r}_{st}), \quad \bar{p}_{st} := \mathbb{E}_s[p_t], \quad \bar{r}_{st} := \mathbb{E}_s[r_t].$$

Thus, Q_s is the seller's subjective expected present unit value of the project after the PPA contract expires. The seller's gain consists of the contract payment from the buyer, plus the seller's subjective post-contract market value of the project, minus the upfront project cost.

Further, we make the following assumption:

Assumption 2.2. *The project's upfront cost $C(q)$ is continuously differentiable, increasing, and strictly convex in project size q .*

We acknowledge that there could be an alternative assumption that the cost function is concave, which may be suitable for capturing economies of scale. A concave cost function, however, would lead to an unrealistic result that suggests building the maximal project size (due to the decreasing marginal cost of the concave cost function). In practice, the cost function can be first concave for small project size q because of the economies of scale, whereas it turns convex as q increases beyond a certain point because of the increasing marginal cost as production level rises. Indeed, we find that the true optimal point always lies within the curve's convex range, implying that our assumption of convexity is without loss of generality.

2.3.2 Three Bargaining Mechanisms

We compare three procurement architectures. The benchmark is *separate negotiation*, in which each buyer negotiates bilaterally with the seller without coordinating with the other buyer. We then study two collective mechanisms observed in practice [Bird, 2018]. In *joint negotiation*, the buyers form a buyer group, determine their internal allocation, and then bargain collectively with the seller [Zanchi et al., 2017]. In *agent-based negotiation*, an anchor tenant (which we call the “*agent*”) first contracts with a follower and then negotiates with the seller on behalf of both buyers [Tamarindo, 2019, Zeigo, 2023]. Figure 2.2 illustrates these three mechanisms.

Our base model considers two buyers, denoted by 1 and 2, and one seller, denoted by s . The parties negotiate over PPA contract prices and project size. Depending on the mechanism, the buyers may receive distinct contract prices (v_1, v_2) and project sizes (q_1, q_2) , or a common total project size q that is allocated internally. The three mechanisms are detailed as follows.

- **Separate Negotiation:** Buyers engage in concurrent bilateral negotiations with the seller to determine their respective (v_1, q_1) and (v_2, q_2) . Buyers negotiate with the seller separately while knowing the other buyer is also engaged but without coordinating with each other.
- **Joint Negotiation:** Buyers first form a buyer group by deciding on allocation fractions (z_1, z_2) , after which they jointly negotiate with the seller over a uniform price v and total project size q . This approach mirrors practical applications where buyers establish a consortium, agree internally on project output distribution, and then negotiate collectively with the seller Zanchi et al. [2017].
- **Agent-based Negotiation:** One buyer assumes the role of the agent while the other acts as the follower. The negotiation proceeds in two stages: (i) the agent (assume to be buyer 1 for illustration) first reaches an agreement with the follower (signing a PPA contract specifying project size q_2 and contract price v_2), and (ii) given this agreement, the agent negotiates with the seller (on behalf of both parties), setting the project's total size $q = q_1 + q_2$ and the contract price v_1 . This structure reflects industry practices in which an anchor tenant (the agent) contracts with sub-level buyers (the followers) before negotiating a larger agreement with the seller [Tamarindo, 2019].

To determine the equilibrium PPA contract prices and project sizes under different negotiation mechanisms, we apply the generalized Nash bargaining framework widely used in operations management research (Nagarajan and Bassok 2008, Feng and Lu 2013, Feng et al. 2015, Hsu et al. 2017, Wang et al. 2021). Following this stream of literature, we adopt the notion of relative bargaining power of different stakeholders as follows. For separate




Figure 2.2: Illustration of Negotiation Mechanisms

negotiation, let $\alpha_{1s} \in [0, 1]$ and $\alpha_{2s} \in [0, 1]$ denote buyers 1 and 2's respective bargaining power against the seller. Let $\alpha_{12} \in [0, 1]$ denote buyer 1's bargaining power against buyer 2. In joint negotiation, let $\alpha_{bs} \in [0, 1]$ denote the buyer group's bargaining power against the seller. In agent-based negotiation, let $\alpha_{as} \in [0, 1]$ denote the acting agent's bargaining power against the seller. Since the buyers and the seller are assumed to be risk neutral in the base model (see Section 2.4 for the risk-averse setting), each party maximizes its expected bargaining gain from reaching agreement.

The Separate Negotiation Baseline

We begin with separate negotiation as the mechanism selection benchmark. This baseline is useful because the collective mechanisms will later be evaluated relative to whether they improve upon separate bilateral negotiation with the seller. Buyers 1 and 2 negotiate bilaterally and concurrently with the seller, without coordinating with each other. The corresponding bargaining gains are:

$$\Delta_1 = (Q_1 - G_1 v_1) q_1, \quad \Delta_2 = (Q_2 - G_2 v_2) q_2, \quad \Delta_s = G_s (v_1 q_1 + v_2 q_2) + Q_s (q_1 + q_2) - C(q_1 + q_2). \quad (2.6)$$

The equilibrium outcomes are determined by simultaneously maximizing the Nash products Ω_{1s}^S and Ω_{2s}^S in Equations (2.7) and (2.8). This method aligns with the approach utilized

by Wang et al. [2021].

$$\max_{v_1, q_1} \Omega_{1s}^S = \Delta_1^{\alpha_{1s}} \Delta_s^{1-\alpha_{1s}}, \quad \text{subject to } \Delta_1 \geq 0, \Delta_s \geq 0; \quad (2.7)$$

$$\max_{v_2, q_2} \Omega_{2s}^S = \Delta_2^{\alpha_{2s}} \Delta_s^{1-\alpha_{2s}}, \quad \text{subject to } \Delta_2 \geq 0, \Delta_s \geq 0. \quad (2.8)$$

The objective functions in Equations (2.7) and (2.8) depend on all four decision variables because the seller's gain reflects the total project size. When addressing (2.7), we treat v_2 and q_2 as given and determine $v_1^S(v_2, q_2)$ and $q_1^S(v_2, q_2)$ as functions of (v_2, q_2) , where the superscript S denotes separate negotiation. Similarly, when addressing (2.8), we obtain $v_2^S(v_1, q_1)$ and $q_2^S(v_1, q_1)$ as functions of (v_1, q_1) . The equilibrium outcome is then characterized by the fixed-point relationships $v_i^S = v_i^S(v_j^S, q_j^S)$ and $q_i^S = q_i^S(v_j^S, q_j^S)$, where $i \in \{1, 2\}$, $j \in \{1, 2\}$, and $i \neq j$.

Assumption 2.3. $(Q_1/G_1) = (Q_2/G_2)$.

Assumption 2.3 imposes a common annualized expected unit value for the project output across the two buyers. Because Q_i/G_i is buyer i 's subjective expected present unit value normalized by his cumulative discounting factor, the assumption does not require identical period-by-period beliefs, identical discount factors, or buyer coordination. It only requires the two buyers' discounted average valuations of one unit of project output during the contract term to coincide.

This condition is especially natural in the collective mechanisms, where buyers communicate before the seller-facing negotiation and can align on project-level valuation inputs. In the separate-negotiation benchmark, however, the assumption is not a coordination requirement, but an equilibrium compatibility condition. In separate negotiation, in the absence of agreement on the annualized unit value, the simultaneous bilateral benchmark would generally lead to incompatible total project sizes. We therefore impose Assumption 2.3 to obtain a well-defined equilibrium outcome and keep the baseline focused on how negotiation architecture affects surplus allocation under apple-to-apple comparisons against the collective mechanisms.

Economically, this common-value benchmark is also plausible in PPA procurement. Even when buyers negotiate separately, they evaluate the same project, the same market node,

and the same REC stream, often using shared public price information, consultant forecasts, or similar internal valuation models. Thus, Assumption 2.3 can be viewed as a controlled benchmark in which buyers may differ in bargaining power and discounting but not in the annualized expected value they assign to the project output. Later extensions relax this restriction and show that the main mechanism insights are not driven by this common-value benchmark.

Lemma 2.4 (Separate-negotiation benchmark). *Under Assumptions 2.2 and 2.3, separate negotiation yields project size $q^S = q^*$, where*

$$C'(q^*) = \frac{G_s}{G_1} Q_1 + Q_s = \frac{G_s}{G_2} Q_2 + Q_s.$$

The equilibrium bargaining gains are

$$\Delta_1^S = \frac{(1 - \alpha_{2s}) \alpha_{1s} G_1}{1 - \alpha_{1s} \alpha_{2s}} \frac{G_1}{G_s} \Delta^{full}, \quad \Delta_2^S = \frac{(1 - \alpha_{1s}) \alpha_{2s} G_2}{1 - \alpha_{1s} \alpha_{2s}} \frac{G_2}{G_s} \Delta^{full}, \quad \Delta_s^S = \frac{1 - \alpha_{1s} - \alpha_{2s} + \alpha_{1s} \alpha_{2s}}{1 - \alpha_{1s} \alpha_{2s}} \Delta^{full},$$

where

$$\Delta^{full} \equiv \frac{G_s}{G_1} Q_1 q^* + Q_s q^* - C(q^*) = \frac{G_s}{G_2} Q_2 q^* + Q_s q^* - C(q^*).$$

Contract-price formulas and the non-unique project split are reported in Appendix A.1.

The Collective Mechanisms

We now turn to the two collective procurement structures, joint negotiation and agent-based negotiation.

Joint Negotiation. Under the joint negotiation mechanism, the bargaining gains for the two buyers and the seller are:

$$\Delta_1 = Q_1 q z_1 - G_1 v q z_1, \quad \Delta_2 = Q_2 q z_2 - G_2 v q z_2, \quad \Delta_s = G_s v q + Q_s q - C(q). \quad (2.9)$$

The stage-one negotiation (between the two buyers) maximizes Equation (2.10), and the stage-two negotiation (between the buyer group and the seller) maximizes the Nash product in Equation (2.11).

$$\text{Stage 1.} \quad \max_{z_1} \quad \Omega_{12}^J = \Delta_1^{\alpha_{12}} \Delta_2^{1-\alpha_{12}} \quad \text{Subject to } \Delta_1 \geq 0, \Delta_2 \geq 0; \quad (2.10)$$

$$\text{Stage 2.} \quad \max_{v,q} \quad \Omega_{bs}^J = (\Delta_1 + \Delta_2)^{\alpha_{bs}} \Delta_s^{1-\alpha_{bs}} \quad \text{Subject to } \Delta_1 + \Delta_2 \geq 0, \Delta_s \geq 0. \quad (2.11)$$

We solve the joint negotiation problem by backward induction. We start with the stage-two negotiation between the buyer group and the seller, given known values of z_1 and z_2 obtained from the stage-one negotiation. Solving (2.11) yields $v^J(z_1, z_2)$ and $q^J(z_1, z_2)$ as a function of (z_1, z_2) . Next we move on to solving stage-one with $v^J(z_1, z_2)$ and $q^J(z_1, z_2)$ plugged in to derive the optimal project allocation z_1^* .

Lemma 2.5 (Joint negotiation outcome). *Under Assumptions 2.2 and 2.3, joint negotiation yields project size $q^J = q^*$, where*

$$C'(q^*) = \frac{G_s}{G_1}Q_1 + Q_s = \frac{G_s}{G_2}Q_2 + Q_s.$$

The equilibrium bargaining gains are

$$\Delta_1^J = \alpha_{bs}\alpha_{12}\frac{G_1}{G_s}\Delta^{full}, \quad \Delta_2^J = \alpha_{bs}(1 - \alpha_{12})\frac{G_2}{G_s}\Delta^{full}, \quad \Delta_s^J = (1 - \alpha_{bs})\Delta^{full}.$$

The project split is $q_1^J = \alpha_{12}q^$ and $q_2^J = (1 - \alpha_{12})q^*$. The contract-price formula is reported in Appendix A.1.*

This outcome closely mirrors the classic Nash bargaining logic. The buyer group first secures total bargaining gain $\alpha_{bs}\Delta^{full}$ in the seller-facing negotiation, and the two buyers then divide that group gain according to their respective bargaining powers α_{12} and $1 - \alpha_{12}$. As will be discussed next, this clean “first aggregate, then divide” logic no longer holds in agent-based negotiation, where the sequential structure lets the follower secure surplus before the seller-facing negotiation, so the final allocation can no longer be interpreted as a simple Nash split of a pre-determined buyer-group gain.

Agent-Based Negotiation. We analyze the case where buyer 1 is the agent, with symmetric results if buyer 2 assumes the role. The bargaining gains under the agent-based negotiation mechanism are given by:

$$\begin{aligned} \Delta_1 &= Q_1q_1 + G_1v_2q_2 - G_1v_1(q_1 + q_2), \quad \Delta_2 = Q_2q_2 - G_2v_2q_2, \\ \Delta_s &= G_sv_1(q_1 + q_2) + Q_s(q_1 + q_2) - C(q_1 + q_2). \end{aligned} \tag{2.12}$$

The agent’s gain Δ_1 varies slightly from Equation (2.3), which consists of the value of his own portion of the project, Q_1q_1 , and the payment collected from the follower, $G_1v_2q_2$,

less the total payment made to the seller, $G_1 v_1(q_1 + q_2)$. The seller's gain Δ_s is the sum of the present value of the total payment received, $G_s v_1(q_1 + q_2)$, and the seller's post-contract valuation, $Q_s(q_1 + q_2)$, minus the project cost, $C(q_1 + q_2)$.

For each stage of the negotiation, we define a Nash bargaining product based on the individual bargaining gains and relative bargaining power of the two parties involved. The corresponding maximization problems are outlined in Equation (2.13) for the first stage between the two buyers and in Equation (2.14) for the second stage between the agent and the seller.

$$\text{Stage 1. } \max_{v_2, q_2} \Omega_{12}^A = \Delta_1^{\alpha_{12}} \Delta_2^{1-\alpha_{12}}, \quad \text{Subject to } \Delta_1 \geq 0, \Delta_2 \geq 0; \quad (2.13)$$

$$\text{Stage 2. } \max_{v_1, q_1} \Omega_{1s}^A = \Delta_1^{\alpha_{as}} \Delta_s^{1-\alpha_{as}}, \quad \text{Subject to } \Delta_1 \geq 0, \Delta_s \geq 0. \quad (2.14)$$

Starting with the second-stage problem, (2.14), where v_2 and q_2 have been determined from the first stage, we obtain the equilibrium value of $v_1^A(v_2, q_2)$ and $q_1^A(v_2, q_2)$ as a function of v_2 and q_2 . Here, the superscript “A” highlights the equilibrium outcomes derived for agent-based negotiation. Then, we substitute $v_1^A(v_2, q_2)$ and $q_1^A(v_2, q_2)$ into the optimization problem for the first stage.

Lemma 2.6 (Agent-based negotiation outcome). *Under Assumptions 2.2 and 2.3, when buyer 1 acts as the agent and buyer 2 acts as the follower, agent-based negotiation yields project size $q^A = q^*$, where*

$$C'(q^*) = \frac{G_s}{G_1} Q_1 + Q_s = \frac{G_s}{G_2} Q_2 + Q_s.$$

The equilibrium bargaining gains are

$$\begin{aligned} \Delta_1^A &= \alpha_{as} \alpha_{12} \frac{G_1}{G_s} \Delta^{full}, \\ \Delta_2^A &= (1 - \alpha_{12}) \frac{G_2}{G_s} \Delta^{full}, \\ \Delta_s^A &= (1 - \alpha_{as}) \alpha_{12} \Delta^{full}. \end{aligned}$$

Contract-price formulas and the non-unique project split are reported in Appendix A.1. The case where buyer 2 acts as the agent is symmetric.

Like alluded to earlier, the allocation in agent-based negotiation departs from the classic simultaneous Nash-bargaining logic. This is because the negotiation sequence ensures that the follower settles with the agent before the agent bargains with the seller. This early agreement lets the follower lock in a total profit of $(1 - \alpha_{12})(G_2/G_s)\Delta^{full}$, independent of the agent's seller-facing bargaining power α_{as} . Hence, the follower's payoff is solely governed by his bargaining position relative to the agent, not by the buyer group's bargaining position relative to the seller. The residual α_{12} -portion of the pie is then carried into the second stage, where the agent and seller divide it according to α_{as} and $1 - \alpha_{as}$. The agent-based mechanism therefore works like a sequential bargaining arrangement with pre-commitment where the first-stage agreement removes part of the surplus from the later negotiation and limits the seller's ability to extract it. This interpretation echoes the bargaining literature on commitment and sequential offers, where the ability to commit early or move first can alter the subsequent bargaining frontier [Schelling, 1956, Rubinstein, 1982, Muthoo, 1999].

2.3.3 Project Size, Viability, and Mechanism Choice

Lemmas 2.4–2.6 yield three main implications for the base model. We first state the common project-size result. We then compare joint negotiation and agent-based negotiation to separate negotiation, and finally compare the two collective mechanisms when both are viable.

Theorem 2.7 (Common equilibrium project size). *Under Assumptions 2.2 and 2.3, and for risk-neutral buyers, the equilibrium project size is unique and identical across the separate, joint, and agent-based mechanisms:*

$$q^S = q^J = q^A \equiv q^*, \quad \text{where } C'(q^*) = \frac{G_s}{G_1}Q_1 + Q_s = \frac{G_s}{G_2}Q_2 + Q_s.$$

Theorem 2.7 shows that the equilibrium project size is determined by the point where the project's marginal cost, $C'(q)$, equals its marginal lifetime value. This marginal value consists of two components: (i) buyers' beliefs about the project's expected net present unit value during the contractual period adjusted by the seller's discount factor, namely $\frac{G_s}{G_1}Q_1$ and $\frac{G_s}{G_2}Q_2$; and (ii) the seller's belief about the project's post-contract expected unit value, Q_s .

We next compare the collective mechanisms with the separate-negotiation benchmark. Since Theorem 2.7 shows that all three mechanisms induce the same risk-neutral project size, the comparison in this subsection is about how the common surplus is allocated across the buyers and the seller. To make the comparison transparent, we impose the following structure on the bargaining-power parameters.

Assumption 2.8. *For $\alpha_{1s}, \alpha_{2s} \in (0, 1)$, the bargaining-power parameters satisfy*

$$\alpha_{12} = \frac{\alpha_{1s}}{\alpha_{1s} + \alpha_{2s}}, \quad \alpha_{bs} = \min\{\alpha_{1s} + \alpha_{2s}, 1\}.$$

When buyer i serves as the agent and buyer j as the follower, where $\{i, j\} = \{1, 2\}$, the acting agent's bargaining power against the seller is

$$\alpha_{as}^{(i)}(\lambda) = \alpha_{is} + \lambda(\alpha_{bs} - \alpha_{is}), \quad \lambda \in [0, 1].$$

The first condition states that the relative bargaining power between the two buyers is proportional to their respective bargaining powers against the seller. Thus, if buyer 1 has stronger seller-facing bargaining power than buyer 2, then buyer 1 also has stronger bargaining power in the internal buyer-buyer negotiation. The second condition states that the buyer group's seller-facing bargaining power is additive, but capped at one. The last condition indicates that if buyer i is designated as the agent and buyer j as the follower, then the agent's bargaining power against the seller is at least as large as buyer i 's standalone bargaining power, but cannot exceed the pooled buyer bargaining power or one. This captures the idea that serving as an agent may strengthen the buyer's position by aggregating demand, while still respecting the cap imposed by the total bargaining strength available to the two buyers.

We acknowledge the possibility of alternative assumptions to the conditions outlined above. First, it can be argued that α_{12} is greater or less than $\alpha_{1s}/(\alpha_{1s} + \alpha_{2s})$ because one of the buyers might have a specific advantage or disadvantage relative to the seller that is not present in bilateral direct negotiations between the two buyers. Second, one might also wonder whether the qualitative insights would change if α_{bs} is less than $\alpha_{1s} + \alpha_{2s}$, implying that forming a buyer group could dilute rather than consolidate bargaining power against the seller. To test the robustness of our results, we have investigated mechanism

comparisons under alternative bargaining-power assumptions. Due to space constraints, however, we provide the details in Appendix A.2.

We next compare the two collective mechanisms relative to the separate-negotiation benchmark. For any mechanism m , let Δ_i^m denote buyer i 's equilibrium bargaining gain. We say that mechanism m is *buyer-wise viable* if it is a Pareto improvement for the two buyers relative to separate negotiation, i.e., if $\Delta_1^m \geq \Delta_1^S$ and $\Delta_2^m \geq \Delta_2^S$. This criterion requires that a collective mechanism must first be acceptable to both buyers individually. To better characterize buyer-wise viability, define

$$\tau_i = \frac{(1 - \alpha_{js})(\alpha_{1s} + \alpha_{2s})}{1 - \alpha_{1s}\alpha_{2s}}, \quad \{i, j\} = \{1, 2\}. \quad (2.15)$$

The quantity τ_i is the minimum seller-facing bargaining power that a collective mechanism must deliver for buyer i to weakly prefer that mechanism to separate negotiation.

For joint negotiation, using Lemmas 2.4 and 2.5, we have

$$\Delta_i^J - \Delta_i^S = \frac{G_i}{G_s} \Delta^{full} \frac{\alpha_{is}}{\alpha_{1s} + \alpha_{2s}} (\alpha_{bs} - \tau_i), \quad \{i, j\} = \{1, 2\}.$$

Thus, joint negotiation is buyer-wise viable if and only if $\alpha_{bs} \geq \tau_i$ for both buyers. When $\alpha_{1s} + \alpha_{2s} \leq 1$, we have $\alpha_{bs} = \alpha_{1s} + \alpha_{2s}$, and both buyer participation constraints are automatically satisfied. When $\alpha_{1s} + \alpha_{2s} > 1$, we have $\alpha_{bs} = 1$, and the two buyer participation constraints become $\alpha_{1s} + \alpha_{2s} \leq 1 + \alpha_{2s}^2$ for buyer 1 and $\alpha_{1s} + \alpha_{2s} \leq 1 + \alpha_{1s}^2$ for buyer 2. Combining these observations gives the following result.

Proposition 2.9 (Joint negotiation relative to separate negotiation). *Suppose $\Delta^{full} > 0$ and $0 < \alpha_{1s}, \alpha_{2s} < 1$. Under Assumptions 2.2–2.8, joint negotiation is buyer-wise viable relative to separate negotiation if and only if*

$$\alpha_{1s} + \alpha_{2s} \leq 1 + \min\{\alpha_{1s}^2, \alpha_{2s}^2\}. \quad (2.16)$$

In particular, joint negotiation is always buyer-wise viable when $\alpha_{1s} + \alpha_{2s} \leq 1$.

Proposition 2.9 shows that joint negotiation is always buyer-wise viable in the uncapped region ($\alpha_{1s} + \alpha_{2s} \leq 1$). In the capped region ($\alpha_{1s} + \alpha_{2s} > 1$), however, viability can fail. Because the coalition's aggregate bargaining power cannot exceed one, the additional surplus

generated by joint negotiation is strictly limited. This capped benefit may be insufficient to attract a buyer who already possesses a strong separate-negotiation benchmark. Consequently, condition (2.16) dictates that joint negotiation remains viable in the capped region only if the buyers have relatively balanced bargaining powers. For example, if the buyers are highly asymmetric – such as when α_{1s} approaches zero and α_{2s} approaches one – the right-hand side of the inequality is bottlenecked by the weaker buyer ($\min\{\alpha_{1s}^2, \alpha_{2s}^2\} \rightarrow 0$). The condition effectively reverts to $\alpha_{1s} + \alpha_{2s} \leq 1$, which is violated. Intuitively, a severely disadvantaged buyer contributes too little additional leverage to the coalition to compensate the stronger buyer for sharing the capped joint surplus.

We next consider agent-based negotiation. Let $A^{(i)}$ denote the agent-based mechanism in which buyer i acts as the agent and buyer j acts as the follower, where $\{i, j\} = \{1, 2\}$. A role assignment $A^{(i)}$ is buyer-wise viable if $\Delta_i^{A^{(i)}} \geq \Delta_i^S$ and $\Delta_j^{A^{(i)}} \geq \Delta_j^S$. Agent-based negotiation is buyer-wise viable if at least one of the two role assignments is buyer-wise viable.

For role assignment $A^{(i)}$, the agent's participation constraint is obtained from

$$\Delta_i^{A^{(i)}} - \Delta_i^S = \frac{G_i}{G_s} \Delta^{full} \frac{\alpha_{is}}{\alpha_{1s} + \alpha_{2s}} \left(\alpha_{as}^{(i)}(\lambda) - \tau_i \right).$$

Thus, the agent participates if and only if $\alpha_{as}^{(i)}(\lambda) \geq \tau_i$. The follower's gain does not depend on the agent's seller-facing bargaining power. Using Lemmas 2.4 and 2.6, the follower's participation constraint reduces to

$$\Delta_j^{A^{(i)}} - \Delta_j^S = \frac{G_j}{G_s} \Delta^{full} \frac{\alpha_{js} (1 - \alpha_{1s} - \alpha_{2s} + \alpha_{is}^2)}{(\alpha_{1s} + \alpha_{2s})(1 - \alpha_{1s}\alpha_{2s})}.$$

Hence, the follower participates if and only if $\alpha_{1s} + \alpha_{2s} \leq 1 + \alpha_{is}^2$.

Proposition 2.10 (Agent-based negotiation relative to separate negotiation). *Suppose $\Delta^{full} > 0$ and $0 < \alpha_{1s}, \alpha_{2s} < 1$. Under Assumptions 2.2–2.8, agent-based negotiation with buyer i as the agent and buyer j as the follower is buyer-wise viable relative to separate negotiation if and only if*

$$\alpha_{as}^{(i)}(\lambda) \geq \tau_i \quad \text{and} \quad \alpha_{1s} + \alpha_{2s} \leq 1 + \alpha_{is}^2, \quad \{i, j\} = \{1, 2\}.$$

Proposition 2.10 highlights the additional role-assignment friction in agent-based negotiation. The agent must receive enough seller-facing bargaining-power lift to prefer serving as the anchor buyer rather than negotiating separately. The follower, by contrast, benefits from the internal buyer-buyer bargain but may still reject the arrangement when his separate-negotiation outcome is too strong. Thus, agent-based negotiation can fail because the agent lacks enough seller-facing bargaining power, because the follower has too strong an outside option, or both.

An immediate implication of Propositions 2.9 and 2.10 is that the conditions for agent-based viability are stronger than those for joint viability under Assumption 2.8. Specifically, if role assignment $A^{(i)}$ is viable, the agent's participation constraint requires $\alpha_{as}^{(i)}(\lambda) \geq \tau_i$. Because the agent's bargaining power is bounded by the group's total collective power ($\alpha_{as}^{(i)}(\lambda) \leq \alpha_{bs}$), joint negotiation automatically satisfies buyer i 's participation constraint. Furthermore, the follower condition in Proposition 2.10 exactly matches the requirement for joint negotiation to satisfy buyer j 's participation constraint in the capped region. As a result, whenever an agent-based arrangement is buyer-wise viable, joint negotiation is also guaranteed to be viable. Importantly, these two propositions collectively demonstrate that both collective mechanisms have the potential to Pareto dominate the separate-negotiation benchmark.

We now compare joint and agent-based negotiation from the buyer-group perspective, conditional on buyer-wise viability. And we use $\Delta_b^m = \Delta_1^m + \Delta_2^m$ to denote the buyer group's total bargaining gain under mechanism m .

To facilitate this comparison, we first determine which buyer should act as the agent from the group's perspective when both role assignments can be buyer-wise viable. Specifically, we assume the group selects the agent that maximizes the buyer group's total bargaining gain. For agent-based negotiation with buyer i as the agent and buyer j as the follower, where $\{i, j\} = \{1, 2\}$,

$$\Delta_b^{A^{(i)}} = \frac{\Delta^{full}}{G_s} \left[G_i \frac{\alpha_{is}}{\alpha_{1s} + \alpha_{2s}} \alpha_{as}^{(i)}(\lambda) + G_j \frac{\alpha_{js}}{\alpha_{1s} + \alpha_{2s}} \right] = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \left[G_1 \alpha_{1s} + G_2 \alpha_{2s} - G_i \alpha_{is} (1 - \alpha_{as}^{(i)}(\lambda)) \right].$$

The first two terms, $G_1 \alpha_{1s} + G_2 \alpha_{2s}$, are common across role assignments. The only role-

dependent term is

$$\frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} G_i \alpha_{is} (1 - \alpha_{as}^{(i)}(\lambda)).$$

This term is the *buyer-group value lost to seller extraction* when buyer i acts as the agent. To see this, note that $\alpha_{is}/(\alpha_{1s} + \alpha_{2s})$ is buyer i 's share of the buyer-side pie in the internal buyer-buyer negotiation, while $1 - \alpha_{as}^{(i)}(\lambda)$ is the seller's residual bargaining share in the second-stage agent-seller negotiation. Thus, the seller extracts $\frac{\alpha_{is}}{\alpha_{1s} + \alpha_{2s}} (1 - \alpha_{as}^{(i)}(\lambda)) \Delta^{full}$ from the portion of surplus carried by the agent into the second-stage negotiation. In the buyer-group objective, this seller-extracted amount is scaled by G_i/G_s , because the lost surplus would otherwise have accrued to buyer i . Hence, the role-choice problem is equivalent to choosing the agent assignment that minimizes $\frac{G_i \alpha_{is}}{\alpha_{1s} + \alpha_{2s}} (1 - \alpha_{as}^{(i)}(\lambda))$.

Lemma 2.11 (Agent-role choice under buyer-wise viability). *Suppose $\Delta^{full} > 0$, $G_1, G_2 > 0$, and $0 < \alpha_{1s}, \alpha_{2s} < 1$. Under Assumptions 2.2–2.8, when both agent-based role assignments are buyer-wise viable, the buyer-group gain difference between the two role assignments is*

$$\Delta_b^{A(1)} - \Delta_b^{A(2)} = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \left[G_2 \alpha_{2s} (1 - \alpha_{as}^{(2)}(\lambda)) - G_1 \alpha_{1s} (1 - \alpha_{as}^{(1)}(\lambda)) \right].$$

Equivalently, buyer 1 should be selected as the agent if and only if $\Delta_b^{A(1)} \geq \Delta_b^{A(2)}$, and buyer 2 should be selected as the agent otherwise.

Adopting Assumption 2.8, the difference expands to

$$\Delta_b^{A(1)} - \Delta_b^{A(2)} = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \begin{cases} G_2 \alpha_{2s} (1 - \alpha_{2s} - \lambda \alpha_{1s}) - G_1 \alpha_{1s} (1 - \alpha_{1s} - \lambda \alpha_{2s}), & \text{if } \alpha_{1s} + \alpha_{2s} \leq 1, \\ (1 - \lambda) [G_2 \alpha_{2s} (1 - \alpha_{2s}) - G_1 \alpha_{1s} (1 - \alpha_{1s})], & \text{if } \alpha_{1s} + \alpha_{2s} > 1. \end{cases}$$

The distinction between the uncapped and capped regions stems from how the coalition's aggregate power limit, $\alpha_{bs} = \min\{\alpha_{1s} + \alpha_{2s}, 1\}$, shapes the agent's seller-facing power, $\alpha_{as}^{(i)}(\lambda)$. In the uncapped region ($\alpha_{1s} + \alpha_{2s} \leq 1$), the buyer group has not exhausted the bargaining-power upper bound. As a result, the seller's extraction, $1 - \alpha_{as}^{(i)}(\lambda) = (1 - \alpha_{1s} - \alpha_{2s}) + (1 - \lambda) \alpha_{js}$, comes from (i) the buyer group's remaining aggregate weakness and (ii) the imperfect transmission of the follower's bargaining power to the agent. Because the follower locks in their share before facing the seller, agent-based negotiation functions

primarily as a protection device. Consequently, the group should assign the lower-power buyer as the agent, effectively shielding the higher-power buyer's larger internal share from seller extraction.

Conversely, in the capped region ($\alpha_{1s} + \alpha_{2s} > 1$), the combined power hits the upper bound ($\alpha_{bs} = 1$). The seller's residual share simplifies to $1 - \alpha_{as}^{(i)}(\lambda) = (1 - \lambda)(1 - \alpha_{is})$. Because aggregate weakness is no longer the bottleneck, the follower's standalone power becomes irrelevant. The sole remaining source of seller extraction is delegation loss, namely the untransmitted gap between the agent's standalone power and the coalition cap. In this region, the priority shifts entirely to carrying as much of the capped coalition power as possible into the final negotiation. Since a stronger standalone buyer is naturally closer to the cap, thereby leaving a smaller residual pie for the seller, the higher-power buyer is the optimal agent.

We now return to the comparison between agent-based and joint negotiation, assuming at least one agent-based role assignment is buyer-wise viable. If only one role assignment is viable, it serves as the unique implementable agent-based mechanism to be compared against joint negotiation. If both are viable, the buyer group selects the optimal assignment per Lemma 2.11, yielding a group gain of $\Delta_b^A = \max\{\Delta_b^{A^{(1)}}, \Delta_b^{A^{(2)}}\}$. Meanwhile, the buyer group's gain under joint negotiation is

$$\Delta_b^J = \frac{\Delta^{full}}{G_s} \alpha_{bs} \left[G_1 \frac{\alpha_{1s}}{\alpha_{1s} + \alpha_{2s}} + G_2 \frac{\alpha_{2s}}{\alpha_{1s} + \alpha_{2s}} \right].$$

Theorem 2.12 (Collective-mechanism choice under buyer-wise viability). *Suppose $G_1, G_2 > 0$, $\Delta^{full} > 0$, and $0 < \alpha_{1s}, \alpha_{2s} < 1$. Under Assumptions 2.2–2.8, for any viable role assignment $A^{(i)}$, the comparison with joint negotiation is*

$$\Delta_b^{A^{(i)}} - \Delta_b^J = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \left[G_j \alpha_{js}(1 - \alpha_{bs}) - G_i \alpha_{is}(\alpha_{bs} - \alpha_{as}^{(i)}(\lambda)) \right].$$

Equivalently,

$$\Delta_b^{A^{(i)}} - \Delta_b^J = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \begin{cases} \alpha_{js} [G_j(1 - \alpha_{1s} - \alpha_{2s}) - G_i \alpha_{is}(1 - \lambda)], & \text{if } \alpha_{1s} + \alpha_{2s} \leq 1, \\ -G_i \alpha_{is}(1 - \lambda)(1 - \alpha_{is}), & \text{if } \alpha_{1s} + \alpha_{2s} > 1. \end{cases}$$

Thus, in the uncapped region, role assignment $A^{(i)}$ gives the buyer group a higher total gain than joint negotiation if and only if

$$G_j(1 - \alpha_{1s} - \alpha_{2s}) > G_i\alpha_{is}(1 - \lambda).$$

In the capped region, joint negotiation weakly dominates any viable agent-based role assignment, with equality only when $\lambda = 1$.

Theorem 2.12 shows that the buyer group's preference between agent-based and joint negotiation again depends on whether their combined seller-facing bargaining power exceeds one.

In the uncapped region, agent-based negotiation protects the follower's share before the seller-facing negotiation; joint negotiation instead exposes the entire buyer group to the seller-facing bargain. Therefore, agent-based negotiation is preferred precisely when the value of protecting the follower's share, $G_j(1 - \alpha_{1s} - \alpha_{2s})$, exceeds the cost of leaving the agent's retained share exposed to seller extraction, $G_i\alpha_{is}(1 - \lambda)$. A larger λ reduces this exposure cost because it gives the acting agent greater access to the buyer group's seller-facing bargaining power.

In the capped region, the result is sharper. Joint negotiation already gives the buyer group the maximum feasible seller-facing bargaining power, $\alpha_{bs} = 1$. Agent-based negotiation cannot improve on this aggregate seller-facing position. It can only match joint negotiation when the acting agent fully accesses the buyer group's seller-facing bargaining power, namely when $\lambda = 1$.

Figure 2.3 illustrates the participation-constrained mechanism-choice map implied by Propositions 2.9–2.10 and Theorem 2.12 under a special case where $G_1 = G_2$. The figure maps the full bargaining-power space $(\alpha_{1s}, \alpha_{2s}) \in (0, 1)^2$. For each λ , gray regions indicate where neither collective mechanism is buyer-wise viable, while light blue indicates viability for joint negotiation alone. When both mechanisms are viable, we compare joint negotiation against the optimal agent-based role assignment, with orange and dark blue regions denoting where agent-based and joint negotiations are respectively preferred by the buyer group. Finally, the dashed diagonal marks the cap boundary ($\alpha_{1s} + \alpha_{2s} = 1$), and the dotted curve shows the joint-viability condition from Proposition 2.9.




Figure 2.3: Mechanism Viability and Buyer-Group Preference Map

As the delegated agent secures greater seller-facing leverage (λ), the viability constraint for agent-based negotiation relaxes, expanding the strategic space where both collective mechanisms are feasible. Within this space, Theorem 2.12 reveals a sharp operational divide where agent-based negotiation emerges as the optimal mechanism when buyers possess low or highly asymmetric standalone bargaining power. In contrast, joint negotiation dominates in balanced or high-power regimes, particularly once the coalition’s aggregate bargaining-power cap binds.

Managerially, the result suggests a two-step mechanism-design rule rather than a universal ranking of collective structures. Joint negotiation should be treated as the robust default which has the broader buyer-wise viability region, and whenever an agent-based arrangement is viable, joint negotiation is viable as well. This makes joint negotiation especially attractive when buyers are balanced or already strong as a coalition, because pooling their bargaining power directly exhausts the available seller-facing leverage. Agent-based negotiation, by contrast, is a more targeted strategic instrument. It is valuable when the buyer group is weak or asymmetric, because the sequential structure lets the follower lock in his share before the seller-facing negotiation and leaves only the agent’s retained share exposed to seller extraction. But this advantage is conditional on implementation. The acting agent must receive enough seller-facing bargaining-power lift, captured by λ , to satisfy his own participation constraint. Moreover, when buyers differ in discounting or contract valuation, role assignment should not be based on bargaining power alone; the group should assign the agent role to minimize the G_i -weighted buyer-group value exposed to seller extraction.

Thus, the practical prescription is not simply to appoint an anchor buyer or to form a consortium. It is to use joint negotiation as the participation-safe baseline, and to switch to agent-based negotiation only when the group can both empower the agent and protect the most valuable follower-side surplus from the seller-facing bargain.

2.4 Risk-Averse Buyers and Price Correlation

In this section, we extend the base model to allow buyers to be risk averse and to account for correlation among energy-related prices. Section 2.4.1 first studies risk aversion while allowing within-period correlation between the project wholesale price and the buyer's local retail price, but without inter-temporal price correlation. Section 2.4.2 then shows that inter-temporal price correlation only changes the definitions of the risk terms without affecting the mathematical structure of the bargaining.

2.4.1 Risk-Averse Buyers without Inter-temporal Price Correlation

We extend the base model by allowing buyers to be risk averse while continuing to treat the seller as risk neutral. This asymmetry reflects the institutional setting of corporate PPAs. Buyers are typically less experienced in wholesale power markets and bear the contract-period market risk associated with the PPA. The seller, by contrast, is a renewable-energy developer with greater market expertise and is partly insulated from contract-period price fluctuations through the fixed PPA payment.

To model the behavior of risk-averse buyers, we adopt a standard mean-variance utility function [Markowitz, 1952, Tobin, 1958, Björk et al., 2014, Markowitz, 2014]:

$$U_i = \mathbb{E}_i[\pi_i] - \rho_i \text{Var}_i(\pi_i), \quad i \in \{1, 2\}, \quad (2.17)$$

where $\rho_i > 0$ is buyer i 's risk-aversion coefficient and π_i denotes the realized financial payoff with or without the contract. All expectations, variances, and covariances, denoted with subscript i , are taken under buyer i 's subjective beliefs. Substituting (2.1) and (2.2) into

(2.17) yields

$$\begin{aligned}
U_i(\pi_i(q_i, v_i)) &= \mathbb{E}_i \left[\sum_{t=1}^T \gamma_i^t \left((p_t - v_i)q_i + r_t(q_i - g_{it}) - c_{it}d_{it} \right) \right] \\
&\quad - \rho_i \text{Var}_i \left[\sum_{t=1}^T \gamma_i^t \left((p_t - v_i)q_i + r_t(q_i - g_{it}) - c_{it}d_{it} \right) \right].
\end{aligned} \tag{2.18}$$

Fully characterizing the variance in (2.18) under arbitrary cross-time and cross-market dependencies yields an analytically intractable covariance matrix. To keep the model manageable while preserving its core economic interactions, we first prioritize the contemporaneous correlation between markets. Consequently, we assume prices are independent across time in the current section and explicitly relax it later in Section 4.2.

Assumption 2.13 (No inter-temporal price correlation). *For each buyer i and any two periods $s \neq t$, all cross-time covariances among p_t , r_t , and c_{it} are zero. Within each period t , the wholesale price and the buyer's local retail price may be correlated: $\text{Cov}_i(p_t, c_{it}) = \sigma_{pc,it}$. REC prices are independent of wholesale prices and local retail prices: $\text{Cov}_i(p_t, r_t) = 0$, $\text{Cov}_i(r_t, c_{it}) = 0$.*

Assumption 2.13 reflects that the wholesale market price and local retail prices may be jointly influenced by common shocks, such as fuel prices, weather conditions, or regional demand fluctuations. In contrast, we assume that the REC price r_t and the wholesale market price p_t are independent, consistent with prior studies showing that REC prices are largely driven by state-specific policies rather than market fundamentals, resulting in high volatility and unpredictability [O'Shaughnessy et al., 2018, Barbose, 2021].

Let the buyer's subjective period means be $\mathbb{E}_i[p_t] = \bar{p}_{it}$, $\mathbb{E}_i[r_t] = \bar{r}_{it}$, $\mathbb{E}_i[c_{it}] = \bar{c}_{it}$, and let the corresponding perceived variances be $\text{Var}_i(p_t) = \sigma_{p,it}^2$, $\text{Var}_i(r_t) = \sigma_{r,it}^2$, $\text{Var}_i(c_{it}) = \sigma_{c,it}^2$. Under Assumption 2.13, (2.18) becomes

$$\begin{aligned}
U_i(\pi_i(q_i, v_i)) &= \sum_{t=1}^T \gamma_i^t \left[(\bar{p}_{it} - v_i)q_i + \bar{r}_{it}(q_i - g_{it}) - \bar{c}_{it}d_{it} \right] \\
&\quad - \rho_i \sum_{t=1}^T \gamma_i^{2t} \left[q_i^2 \sigma_{p,it}^2 + (q_i - g_{it})^2 \sigma_{r,it}^2 + d_{it}^2 \sigma_{c,it}^2 - 2q_i d_{it} \sigma_{pc,it} \right].
\end{aligned} \tag{2.19}$$

Without a PPA contract, buyer i 's utility is

$$U_i(\pi_i(0,0)) = \sum_{t=1}^T \gamma_i^t \left[-\bar{r}_{it}g_{it} - \bar{c}_{it}d_{it} \right] - \rho_i \sum_{t=1}^T \gamma_i^{2t} \left[g_{it}^2 \sigma_{r,it}^2 + d_{it}^2 \sigma_{c,it}^2 \right]. \quad (2.20)$$

Subtracting the no-contract utility from the contract utility yields four aggregate risk terms:

$$V_{pi} := \sum_{t=1}^T \gamma_i^{2t} \sigma_{p,it}^2, \quad V_{ri} := \sum_{t=1}^T \gamma_i^{2t} \sigma_{r,it}^2, \quad V_{gri} := \sum_{t=1}^T \gamma_i^{2t} g_{it} \sigma_{r,it}^2, \quad V_{dpci} := \sum_{t=1}^T \gamma_i^{2t} d_{it} \sigma_{pc,it}. \quad (2.21)$$

The utility gain from entering the PPA is therefore

$$\begin{aligned} \tilde{\Delta}_i &= U_i(\pi_i(q_i, v_i)) - U_i(\pi_i(0,0)) \\ &= \sum_{t=1}^T \gamma_i^t (\bar{p}_{it} + \bar{r}_{it} - v_i) q_i - \rho_i \left[(V_{pi} + V_{ri}) q_i^2 - 2(V_{gri} + V_{dpci}) q_i \right] \\ &= [Q_i - G_i v_i + 2\rho_i (V_{gri} + V_{dpci})] q_i - \rho_i (V_{pi} + V_{ri}) q_i^2, \end{aligned} \quad (2.22)$$

where Q_i and G_i are the discounted expected-price aggregate and cumulative discounting factor defined in Section 2.3. In agent-based negotiation, when buyer i acts as the agent and buyer j acts as the follower, the agent's utility gain becomes

$$\tilde{\Delta}_i^{A(i)} = [Q_i - G_i v_i + 2\rho_i (V_{gri} + V_{dpci})] q_i + G_i (v_j - v_i) q_j - \rho_i (V_{pi} + V_{ri}) q_i^2. \quad (2.23)$$

Here, V_{pi} captures wholesale-price risk and V_{ri} captures REC-price risk. The terms V_{gri} and V_{dpci} enter with the opposite sign because they represent hedging benefits. V_{gri} reflects the reduction in REC-market exposure from receiving project-generated RECs, while V_{dpci} captures the cross-market hedge that arises when project wholesale prices co-move with the buyer's local retail price. The standalone variance of local retail prices cancels out because the buyer pays local utility costs regardless of whether he signs the PPA; local retail price risk therefore affects PPA value only through its covariance with project wholesale prices. In agent-based negotiation, the additional term $G_i(v_j - v_i)q_j$ is the agent's deterministic intermediation margin from contracting with the follower at price v_j and with the seller at price v_i . This term does not introduce additional price risk, because the follower's project share is passed through to the follower rather than retained by the agent. Thus, the

risk-averse gain separates the buyer's expected project value, market-risk exposure, hedging benefits, and, under agent-based negotiation, the deterministic margin from intermediation, which altogether render a sophisticated trade-off.

Substituting (2.22) and (2.23) into the three Nash bargaining problems in Section 2.3 yields the following results.

Proposition 2.14 (Equilibrium project size with risk-averse buyers). *Define*

$$A_i = Q_i + 2\rho_i(V_{gri} + V_{dpci}), \quad B_i = \rho_i(V_{pi} + V_{ri}), \quad i \in \{1, 2\}.$$

The equilibrium project size under separate and agent-based negotiation, denoted by $\tilde{q}^{S,A}$, satisfies

$$C'(\tilde{q}^{S,A}) = Q_s + \frac{G_s(A_1B_2 + A_2B_1 - 2B_1B_2\tilde{q}^{S,A})}{G_2B_1 + G_1B_2}. \quad (2.24)$$

Under joint negotiation, the equilibrium project size $\tilde{q}^J$ is implicitly defined by

$$C'(\tilde{q}^J) = Q_s + \frac{G_s[z_1A_1 + (1 - z_1)A_2 - 2(z_1^2B_1 + (1 - z_1)^2B_2)\tilde{q}^J]}{z_1G_1 + (1 - z_1)G_2}, \quad (2.25)$$

where z_1 is buyer 1's equilibrium project share. A closed-form expression for z_1 is generally unavailable. The complete equilibrium outcomes for separate and agent-based negotiation are summarized in Appendix A.4.

Proposition 2.14 shows that the risk-averse model preserves the bargaining structure of the risk-neutral base model after replacing each buyer's expected unit value Q_i with the risk-adjusted term A_i and adding the marginal risk penalty B_i . Unlike the risk-neutral base model, however, the risk-averse characterization does not require Assumption 2.3. Under risk neutrality, each buyer's marginal benefit from project size is the expected annualized unit value, so stable agreement requires the two buyers to agree on that value, which leads to the equal-ratios condition $Q_1/G_1 = Q_2/G_2$. Under risk aversion, by contrast, marginal utility depends on both expected value and risk exposure. A buyer with a higher expected unit value but stronger risk aversion may value additional project size less than a buyer with a lower expected unit value but weaker risk aversion. This additional risk-adjustment channel allows the bargaining solution to accommodate heterogeneous expected values and

heterogeneous risk preferences. Mathematically, this leads to the relaxation of the assumption $Q_1/G_1 = Q_2/G_2$ of the base model without the risk-adjustment channel when both agents are risk-neutral.

Next, we compare the project size under risk-aversion to that under the risk-neutral setting. It turns out that, in contrast to the conventional wisdom that risk-aversion always leads to smaller equilibrium project size, with PPA contract, risk aversion can either increase or decrease the equilibrium project size relative to the risk-neutral benchmark, depending on whether hedging benefits brought by the PPA contracts dominate the additional exposure to market risk when entering the contracts.

Proposition 2.15 (Risk-averse vs. risk-neutral project sizes). *Let q^* be the unique risk-neutral equilibrium project size defined in Theorem 2.7.*

(a) Separate / Agent-Based Negotiation. *Set*

$$q^+ = \frac{A_1B_2 + A_2B_1 - \frac{Q_1}{G_1}(G_2B_1 + G_1B_2)}{2B_1B_2}.$$

Then the risk-averse equilibrium $\tilde{q}^{S,A}$ satisfies

$$\tilde{q}^{S,A} \begin{cases} > q^*, & \text{if } q^+ > q^*, \\ < q^*, & \text{if } q^+ < q^*. \end{cases}$$

(b) Joint Negotiation. *For any coalition weight $z_1 \in [0, 1]$, define*

$$D_J = z_1G_1 + (1 - z_1)G_2, \quad B_J(z_1) = z_1^2B_1 + (1 - z_1)^2B_2,$$

and

$$q_J^+(z_1) = \frac{z_1A_1 + (1 - z_1)A_2 - \frac{Q_1}{G_1}D_J}{2B_J(z_1)}.$$

The joint risk-averse equilibrium $\tilde{q}^J$ obeys

$$\tilde{q}^J \begin{cases} > q^*, & \text{if } q_J^+(z_1) > q^*, \\ < q^*, & \text{if } q_J^+(z_1) < q^*. \end{cases}$$

The intuition behind Proposition 2.15 is that a PPA creates both hedging benefits and new risk exposure for a risk-averse buyer. On the hedging side, the buyer receives RECs through the contract and therefore reduces exposure to REC market price volatility, as the buyers have dual channels to satisfy her renewable target — the receipt of RECs from the project and the purchase of RECs from the REC market. In addition, when the project wholesale price is positively correlated with the buyer’s local retail price, high local utility bills tend to coincide with high project-market revenues, so the PPA provides a cross-market hedge. These two benefits are captured by V_{gri} and V_{dpci} and enter the risk-adjusted value term $A_i = Q_i + 2\rho_i(V_{gri} + V_{dpci})$.

On the risk-exposure side, a larger project exposes the buyer to more wholesale-price and REC-price volatility. This effect is captured by $B_i = \rho_i(V_{pi} + V_{ri})$. Hence, risk aversion does not monotonically reduce project size. If the hedging benefits are sufficiently large relative to the additional market-risk exposure, risk aversion increases the equilibrium project size. Conversely, if the additional exposure dominates the hedging benefits, risk aversion reduces the equilibrium project size. The cutoff quantities q^+ and $q_J^+(z_1)$ provide a compact characterization of this hedge-exposure tradeoff.

For mechanism comparisons, the key point is that risk aversion changes the equilibrium project size and total risk-adjusted surplus, but it does not change the surplus-allocation structure under separate and agent-based negotiation. As shown in Appendix A.4, the equilibrium gains under these two mechanisms retain the same coefficient of surplus allocation as in the risk-neutral model after replacing total surplus with its risk-adjusted counterpart. Therefore, the buyer-wise viability comparison between agent-based negotiation and separate negotiation follows the same bargaining-power logic as in Section 2.3.

Joint negotiation is less transparent because the equilibrium split z_1 is not available in closed form. Since joint negotiation does not admit a closed-form allocation in this extension, we report the full mechanism-comparison map numerically in Appendix A.7, where Figure A.3(a) shows that the mechanism-choice logic from the base model is largely preserved under risk aversion. Joint negotiation remains the broadly more viable mechanism that especially performs better when the buyer coalition can credibly bargain with the seller as a strong unified party. On the other hand, agent-based negotiation performs better when

the buyers are weak or asymmetric, because the internal buyer-buyer agreement gives the follower its share before the seller-facing negotiation, thereby limiting how much surplus the seller can extract from the buyer side. Risk aversion primarily shifts the boundaries of these regions through its effect on risk-adjusted project size and total surplus.

2.4.2 inter-temporal Price Correlation

The preceding subsection rules out inter-temporal price correlation to obtain transparent expressions for the four risk terms. This restriction is mainly expositional. With inter-temporal correlation, buyer i 's mean-variance gain retains the same quadratic form in project size; only the four risk aggregates need to be replaced by discounted covariance sums.

Continue to assume that REC prices are independent of wholesale prices and local retail prices, but now allow inter-temporal dependence within the wholesale-price process, the REC-price process, and the local-utility-cost process. Define

$$\begin{aligned} V_{pi}^{\text{ser}} &:= \sum_{s=1}^T \sum_{t=1}^T \gamma_i^{s+t} \text{Cov}_i(p_s, p_t), & V_{ri}^{\text{ser}} &:= \sum_{s=1}^T \sum_{t=1}^T \gamma_i^{s+t} \text{Cov}_i(r_s, r_t), \\ V_{gri}^{\text{ser}} &:= \sum_{s=1}^T \sum_{t=1}^T \gamma_i^{s+t} g_{it} \text{Cov}_i(r_s, r_t), & V_{dpci}^{\text{ser}} &:= \sum_{s=1}^T \sum_{t=1}^T \gamma_i^{s+t} d_{it} \text{Cov}_i(p_s, c_{it}). \end{aligned} \quad (2.26)$$

When cross-time covariances are zero, these terms reduce to the no-inter-temporal definitions in (2.21). Under inter-temporal correlation, the utility gain becomes

$$\tilde{\Delta}_i^{\text{ser}} = [Q_i - G_i v_i + 2\rho_i(V_{gri}^{\text{ser}} + V_{dpci}^{\text{ser}})]q_i - \rho_i(V_{pi}^{\text{ser}} + V_{ri}^{\text{ser}})q_i^2. \quad (2.27)$$

Thus, all bargaining characterizations derived under (2.22) continue to apply after replacing

$$(V_{pi}, V_{ri}, V_{gri}, V_{dpci}) \quad \text{with} \quad (V_{pi}^{\text{ser}}, V_{ri}^{\text{ser}}, V_{gri}^{\text{ser}}, V_{dpci}^{\text{ser}}).$$

Inter-temporal correlation changes the numerical values of the risk aggregates, but not the quadratic structure of the buyer's gain or the resulting bargaining logic.

A convenient parameterization is a common-factor representation, which is widely used in the distribution and inventory systems literature to capture intertemporal dependence

and cross-location correlation in demand processes [Erkip et al., 1990]. Under buyer i 's subjective beliefs, suppose

$$p_t = \bar{p}_{it} + \beta_p L_t + \varepsilon_{p,t}, \quad c_{it} = \bar{c}_{it} + \beta_{c,i} L_t + \varepsilon_{c,it}, \quad (2.28)$$

where L_t is a zero-mean stationary AR(1) factor which represents a common stochastic process that drives the evolution of both the wholesale price and the local retail price,

$$L_t = \phi L_{t-1} + \eta_t, \quad \mathbb{E}_i[\eta_t] = 0, \quad \text{Var}_i(\eta_t) = (1 - \phi^2) \sigma_{L,i}^2, \quad (2.29)$$

with $|\phi| < 1$. The idiosyncratic shocks $\varepsilon_{p,t}$ and $\varepsilon_{c,it}$ are mean zero, serially independent, mutually independent, and independent of the factor process, with

$$\text{Var}_i(\varepsilon_{p,t}) = \sigma_{p,it}^2, \quad \text{Var}_i(\varepsilon_{c,it}) = \sigma_{c,it}^2.$$

The normalization in (2.29) implies

$$\text{Cov}_i(L_s, L_t) = \sigma_{L,i}^2 \phi^{|t-s|}. \quad (2.30)$$

Suppose further that REC prices are serially independent and independent of the wholesale-price and retail-price processes, with $\text{Var}_i(r_t) = \sigma_{r,it}^2$. Then the four risk aggregators become

$$\begin{aligned} V_{pi}^{\text{ser}} &= \beta_p^2 \sigma_{L,i}^2 \sum_{s=1}^T \sum_{t=1}^T \gamma_i^{s+t} \phi^{|t-s|} + \sum_{t=1}^T \gamma_i^{2t} \sigma_{p,it}^2, & V_{ri}^{\text{ser}} &= \sum_{t=1}^T \gamma_i^{2t} \sigma_{r,it}^2, \\ V_{gri}^{\text{ser}} &= \sum_{t=1}^T \gamma_i^{2t} g_{it} \sigma_{r,it}^2, & V_{dpci}^{\text{ser}} &= \beta_p \beta_{c,i} \sigma_{L,i}^2 \sum_{s=1}^T \sum_{t=1}^T \gamma_i^{s+t} d_{it} \phi^{|t-s|}. \end{aligned} \quad (2.31)$$

The first line of (2.31) contains the risk-exposure terms, which reduce buyer i 's utility gain through $(V_{pi}^{\text{ser}} + V_{ri}^{\text{ser}}) q_i^2$. The second line contains the hedging terms, which increase the utility gain through $2(V_{gri}^{\text{ser}} + V_{dpci}^{\text{ser}}) q_i$. Thus, the factor model preserves the same mean-variance structure as before, while incorporating both contemporaneous and intertemporal co-movement between project wholesale prices and buyer i 's local retail prices.

The expressions also show how sensitivity to the common-factor affects the buyer's risk-hedging tradeoff. Suppose $\beta_p, \beta_{c,i} > 0$, so project wholesale prices and buyer i 's local retail prices move in the same direction in response to the common factor. An increase in $\beta_{c,i}$

affects only V_{dpci}^{ser} . Thus, when buyer i 's local retail price is more sensitive to the common factor, the PPA provides a stronger cross-market hedge since periods with high local retail prices are more likely to coincide with high project wholesale-price revenues.

By contrast, an increase in β_p affects both sides of the tradeoff. It increases the common-factor component of wholesale-price risk in V_{pi}^{ser} , but it also increases the cross-market hedge in V_{dpci}^{ser} . The former effect is quadratic in β_p , whereas the latter is linear in β_p . Therefore, holding other parameters fixed, a stronger sensitivity of the project wholesale price to the common factor can make wholesale-price risk grow faster than the cross-market hedge. Similarly, higher REC-price volatility increases both V_{ri}^{ser} and V_{gri}^{ser} , while a larger renewable target g_{it} increases only V_{gri}^{ser} , making the REC hedge more valuable.

These comparative statics help interpret the mechanism-comparison maps under risk aversion. The bargaining-power parameters determine the baseline comparison between joint and agent-based negotiation. The risk terms then shift this comparison by changing the equilibrium project size and, therefore, the total surplus available to the buyer coalition. When the exposure terms V_{pi}^{ser} and V_{ri}^{ser} are large relative to the hedging terms V_{gri}^{ser} and V_{dpci}^{ser} , risk aversion lowers the value of expanding the project. This reduction is more pronounced under joint negotiation because the seller-facing negotiation reflects both buyers' risk-adjusted valuations. As a result, the region in which agent-based negotiation is preferred tends to expand. Conversely, when REC hedging or cross-market hedging is strong, for example, when g_{it} , d_{it} , or $\beta_{c,i}$ is large, risk aversion is less likely to reduce project size, making joint negotiation relatively more attractive.

Finally, intertemporal correlation does not affect the risk-neutral base model because expected gains depend only on period means. It matters only under risk aversion, where cross-time covariances enter through the variance term and are summarized by the redefined risk aggregators in (2.31).

In summary, this section shows that allowing buyers to be risk averse and allowing prices to be correlated changes the risk-adjusted value of a PPA, but not the main mechanism-comparison logic from the base model. Contemporaneous correlation between project wholesale prices and local retail prices creates a cross-market hedge, while intertemporal correlation changes the magnitudes of the risk aggregators. In both cases, the buyer's utility

gain retains the same quadratic structure. Wholesale-price and REC-price risks reduce the value of larger projects, whereas REC hedging and cross-market hedging increase it. As a result, risk aversion may increase or decrease equilibrium project size depending on which force dominates. For mechanism choice, collective negotiation continues to improve the buyer side relative to separate negotiation. Between the two collective mechanisms, joint negotiation remains more attractive when the buyers can bargain with the seller as a strong unified group, whereas agent-based negotiation remains more attractive when buyer bargaining powers are low or asymmetric because the follower secures its terms before the seller-facing negotiation. Thus, price correlation mainly shifts the boundaries of the mechanism-comparison regions through its effect on risk-adjusted surplus, rather than changing the underlying comparison among mechanisms.

2.5 Extension: REC Banking

In this section, we consider an alternative extension concerning buyers' operational decisions in the banking of RECs. Following the establishment of a PPA, the buyer receives two distinct components: the energy commodity and the associated RECs. The energy can be used to offset the buyer's own electricity consumption or sold directly to the grid to generate immediate revenue. The RECs, by contrast, do not need to be claimed or monetized immediately. Buyers have the option to "bank" RECs, that is, to defer their use or sale, allowing them to strategically time REC transactions based on market conditions or regulatory developments. This intertemporal flexibility introduces an additional layer of decision-making that can significantly influence the buyer's valuation of the contract under uncertainty. This is the focus of the current subsection. We begin with the risk-neutral case, which yields a clean contract-level capacity result, and then turn to the risk-averse banking problem.

2.5.1 Risk-Neutral Buyers with REC Banking

Following the notation introduced at the beginning of Section 2.3, we assume that the buyer may use RECs either generated under the PPA or purchased from the market to satisfy this

obligation. Any RECs generated in excess of the obligation can be banked at the time of generation and either sold or used in a future period.

Optimal Banking Policy

We formulate a dynamic program (DP) to derive the optimal REC-banking policy for a representative buyer; the policy structure is identical across buyers, so we suppress the buyer index. At the beginning of period t , the buyer holds an inventory of $I_t \geq 0$ banked RECs. After receiving new RECs q and facing the renewable target g_t , the period- t REC surplus is $S_t = [I_t + q - g_t]^+$. The decision variable is an unconstrained *banking target* s_t .

Operationally, the REC inventory carried over to the next period is $I_{t+1} = \min\{s_t^*, S_t\}$ where s_t^* is the optimal banking target. This rule mirrors an order-up-to policy, or more precisely a bank-down-to policy in the present setting [Porteus, 2002]. In particular, when $s_t^* = 0$, the buyer will sell all surplus RECs to the market, which is the assumption of the base model analyzed in the previous sections. When $s_t^* < S_t$, the buyer sells the surplus $S_t - s_t^*$. When $s_t^* > S_t$, the available inventory falls short of the target, so the buyer retains all available RECs and sells none. Importantly, We assume that the buyer does not purchase RECs to go beyond the target. This reflects the institutional role of the buyer. He is not a speculative REC trader, but as a compliance-motivated participants whose REC holdings are tied to renewable targets. While we focus on analyzing our model under this assumption, relaxing it would lead to a trivial solution in which the optimal target level s_t^* becomes unbounded, as discussed below.

The one-period profit can be written as

$$R_t = (p_t - v)q - c_t d_t + r_t(I_t + q - g_t) - r_t I_{t+1} = m_t + r_t(I_t - I_{t+1}),$$

where $m_t = (p_t - v)q - c_t d_t + r_t(q - g_t)$ is the buyer's period- t profit when REC banking is not allowed, that is, when all surplus RECs beyond the renewable target are sold immediately as in Equation (2.1). The additional term $r_t(I_t - I_{t+1})$ therefore captures the incremental gain or loss generated by the option to bank RECs. The operational objective over the

contract horizon is

$$\max_{s_t} \mathbb{E} \left[\sum_{t=1}^T \gamma^t [m_t + r_t(I_t - I_{t+1})] \right] = \mathbb{E} \left[\sum_{t=1}^T \gamma^t m_t \right] + \max_{s_t} \mathbb{E} \left[\sum_{t=1}^T \gamma^t r_t(I_t - \min\{s_t, S_t\}) \right],$$

which is equivalent to maximizing the additional expected profit generated by REC banking,

$$\max_{s_t} \mathbb{E} \left[\sum_{t=1}^T \gamma^t r_t(I_t - \min\{s_t, S_t\}) \right] = \max_{0 \leq s_t \leq S_t} \mathbb{E} \left[\sum_{t=1}^T \gamma^t r_t(I_t - s_t) \right].$$

Let $J_t(I_t)$ denote the optimal expected discounted profit from period t to period T , given initial REC inventory I_t at period t . The problem can then be solved by dynamic programming through the Bellman equation

$$J_t(I_t) = \max_{0 \leq s_t \leq S_t} \{ r_t(I_t - s_t) + \gamma \mathbb{E}_{r_{t+1}} [J_{t+1}(I_{t+1})] \}. \quad (2.32)$$

In this formulation, we assume that the REC price process $\{r_t\}$ is independent across time. As discussed in Section 2.4.2, allowing inter-temporally correlated REC prices changes the continuation value but not the basic bank-or-sell logic of the policy.

Proposition 2.16 (Banking Under Risk Neutrality). *For each period t , the value function $J_t(I_t)$ is convex and piecewise-linear in I_t . Consequently, the optimal banking rule is*

$$s_t^* = \begin{cases} 0, & \text{when } r_t \geq \gamma \frac{d\mathbb{E}[J_{t+1}(I_{t+1})]}{dI_{t+1}}, \text{ sell all available RECs,} \\ [I_t + q - g_t]^+, & \text{when } r_t < \gamma \frac{d\mathbb{E}[J_{t+1}(I_{t+1})]}{dI_{t+1}}, \text{ bank the entire feasible amount.} \end{cases}$$

Moreover, holding all other parameters fixed, J_t is (i) non-decreasing in the project size q , and (ii) non-increasing in the renewable target g_t . This policy applies regardless of cross-time price correlation.

Proposition 2.16 shows that the risk-neutral buyer's trade-off is whether to sell RECs immediately at the current market price r_t or to bank them for future use or sale. The decision depends on comparing the immediate payoff from selling against the expected future value of inventory. Formally, the buyer weighs r_t against the discounted expected marginal continuation value

$$\gamma \frac{d\mathbb{E}[J_{t+1}(I_{t+1})]}{dI_{t+1}}.$$

If this continuation value exceeds the current price, the buyer banks the maximum feasible quantity of RECs. Otherwise, the buyer sells all available RECs.

Impact on Contract Negotiation

To connect the operational problem back to contract negotiation, denote

$$\hat{J}(q, \gamma, \{g_t\}) = \max_{0 \leq s_t \leq S_t} \mathbb{E} \left[\sum_{t=1}^T \gamma^t r_t (I_t - s_t) \right]$$

as the objective value of the DP solved above. Thus, $\hat{J}(q, \gamma, \{g_t\})$ is the buyer's expected incremental profit from optimally banking RECs, given project size q , discount factor γ , and an exogenous renewable-target path $\{g_t\}$. We incorporate $\hat{J}(q, \gamma, \{g_t\})$ into the Nash bargaining equations for each negotiation mechanism, namely Equations (2.7)-(2.8), (2.13)-(2.14), and (2.10)-(2.11), by adding it to the buyer's bargaining payoff. The bargaining problems therefore retain the same structure as in the base model, with $\hat{J}$ acting as an additional buyer-side value term. This yields the following first-order characterization of equilibrium project size.

Corollary 2.17 (Project size with REC banking). *When REC banking is permitted, the equilibrium project size is weakly greater than in the baseline case without REC banking. It is strictly greater whenever at least one buyer has a strictly positive marginal banking value at the no-banking equilibrium.*

Specifically, under separate or agent-based negotiation, the equilibrium project size satisfies

$$C'(q^*) = \frac{G_s}{G_1} (Q_1 + \hat{J}'_1(q_1^*)) + Q_s = \frac{G_s}{G_2} (Q_2 + \hat{J}'_2(q_2^*)) + Q_s,$$

where $\hat{J}'_i(q_i^*)$ denotes a subgradient of buyer i 's banking value at q_i^* . Under joint negotiation,

$$C'(q^*) = G_s \frac{(Q_1 + \hat{J}'_1(q_1^*)) z_1^* + (Q_2 + \hat{J}'_2(q_2^*)) z_2^*}{G_1 z_1^* + G_2 z_2^*} + Q_s.$$

Therefore, REC banking increases the marginal value of project size through the buyer-side terms $\hat{J}'_i(q_i^*)$. The increase in equilibrium project size is strict if the relevant marginal banking value is strictly positive.

Corollary 2.17 isolates the main analytical implication of REC banking; that is, banking adds a nonnegative marginal value term to buyer payoffs and therefore weakly increases

equilibrium project size relative to the no-banking benchmark. The increase is strict whenever banking has positive marginal value at the relevant equilibrium project size. A full comparison of bargaining mechanisms is less transparent because $\hat{J}$ is itself defined by a dynamic operating problem. We therefore treat the capacity effect as the main analytical result here and report the mechanism-comparison map numerically in Appendix A.7. As shown in Figure A.3(b), the preference map under REC banking largely preserves the base-model structure. Joint negotiation remains the more robust participation-safe mechanism, while agent-based negotiation is preferred when the protected follower surplus is sufficiently valuable relative to the delegation loss. REC banking mainly shifts the viability and preference boundaries by adding buyer-specific option value to project output.

2.5.2 Risk-Averse Buyers with REC Banking

Finally, we combine the two preceding extensions by allowing buyers to be risk averse while also banking RECs after entering a contract. As in Section 2.5.1, we begin with the buyer's post-contract operational problem. The buyer chooses a banking policy to maximize the mean-variance utility of operating profits:

$$\max_{0 \leq s_t \leq S_t} \mathbb{E} \left[\sum_{t=1}^T \gamma^t R_t \right] - \rho \text{Var} \left[\sum_{t=1}^T \gamma^t R_t \right],$$

where $R_t = m_t + r_t(I_t - I_{t+1})$ is the single-period profit and $I_{t+1} = \min\{s_t^*, S_t\}$ as defined earlier. Because the REC price r_t is observed at the beginning of period t , immediate period- t profit is known when the banking decision is made. The continuation problem from period t onward is therefore

$$\max_{0 \leq s_t \leq S_t} \gamma^t R_t + \mathbb{E} \left[\sum_{\tau=t+1}^T \gamma^\tau R_\tau \right] - \rho \text{Var} \left[\sum_{\tau=t+1}^T \gamma^\tau R_\tau \right].$$

Assuming that prices are independent across time, the variance of continuation profits simplifies to

$$\text{Var} \left[\sum_{\tau=t+1}^T \gamma^\tau R_\tau \right] = \sum_{\tau=t+1}^T \gamma^{2\tau} \text{Var}[R_\tau],$$

where $\text{Var}[R_t] = \text{Var}[r_t] [(I_t - s_t)^2 - 2(g_t - q)(I_t - s_t) + (g_t - q)^2]$.

A closed-form solution to this dynamic mean-variance problem is generally not obtainable. Our analysis therefore proceeds in two steps. First, Appendix A.5 studies two tractable heuristic policies. The first is a two-period myopic policy, which isolates the basic tradeoff between selling RECs today and carrying them forward to the next period. Under this policy, the buyer sells all RECs in the terminal period; in the preceding period, the banking target equals the next-period REC shortfall, adjusted upward or downward according to the expected REC price change and scaled by risk aversion and REC price variance. The second is a generic unconstrained interior policy, which extends the same logic to a multi-period setting by treating the next-period target as given. This policy sets the current banking target equal to the next-period REC requirement, again adjusted for the expected intertemporal REC price difference and the buyer’s risk exposure. These two policies make transparent how future compliance needs, expected REC price movements, and risk aversion jointly shape the buyer’s banking decision. Second, because embedding the full banking problem into the bargaining model does not yield closed-form mechanism comparisons, we examine the robustness of the base-model insights numerically in Appendix A.7.

Because the truncation of the banking target generates a state-contingent piecewise recursion, a full analytical solution is unavailable for the general problem. Embedding this operational problem into Nash bargaining would therefore create a three-stage optimization that must track both the mean and the variance of banking profits. We therefore restrict the main text to the formulation and leave the detailed policy characterizations to the appendix.

The analytical results above focus on how risk aversion and REC banking affect project size. We next ask whether these operational extensions also change the mechanism-choice implications of the base model. Because joint negotiation under risk aversion and REC banking does not yield the same closed-form surplus allocation as in the risk-neutral benchmark, we examine this question numerically in Appendix A.7. Figure A.3 shows that the buyer-wise viability and buyer-group preference maps are largely preserved across the RA No-Banking and RN Banking settings. In particular, joint negotiation remains the more robust participation-safe collective mechanism, while agent-based negotiation is preferred when the value of shielding the follower’s surplus exceeds the delegation loss from sending only the agent into the seller-facing negotiation. Thus, the extensions primarily shift the

boundaries of the preference regions rather than changing the underlying economic tradeoff.

2.6 Concluding Remarks

This paper studies how collective PPA negotiation mechanisms affect renewable project size, surplus allocation, and buyer participation. We compare separate negotiation, joint negotiation, and agent-based negotiation in a unified generalized Nash bargaining framework. Under risk neutrality and convex project costs, the three mechanisms yield the same equilibrium project size. Thus, in the base model, the negotiation architecture does not change renewable capacity addition; instead, it changes how the value created by the project is allocated among the buyers and the developer. This separation between capacity and surplus clarifies the central challenge in collective PPA design. The main barrier is not necessarily whether collective procurement creates value in aggregate, but whether the chosen mechanism makes each buyer individually better off than negotiating separately.

Our analysis shows that this participation requirement materially shapes the choice between joint and agent-based negotiation. Joint negotiation pools buyer bargaining power before the seller-facing negotiation and therefore provides the broader participation-safe collective structure. Agent-based negotiation has a tighter implementability requirement because the acting agent must carry the residual buyer-side surplus into the seller-facing negotiation and must therefore receive enough seller-facing leverage to participate. Conditional on buyer-wise viability, however, agent-based negotiation can outperform joint negotiation. Its sequential structure allows the follower to lock in a share of surplus before the seller-facing negotiation, shielding that share from seller extraction. The cost of this protection is delegation loss since only the agent, rather than the full buyer group, bargains with the seller. Hence, joint negotiation is preferred when pooled bargaining power is strong, especially for relatively balanced buyer groups, whereas agent-based negotiation is attractive when the shielding benefit for the follower exceeds the delegation loss. This tradeoff also determines the optimal agent role assignment. The group should choose the agent that minimizes the buyer-side surplus exposed to seller extraction, rather than simply appointing the largest or most powerful buyer by default.

The extensions show that operational features can change project size while largely

preserving this mechanism-choice logic. Risk aversion introduces a hedge-versus-exposure tradeoff. A PPA can hedge REC-price risk and, when project-market prices co-move with local retail prices, can also hedge electricity-cost risk; at the same time, larger contracted quantities expose buyers to additional wholesale-price and REC-price volatility. As a result, risk aversion can either increase or decrease equilibrium project size. Serial price correlation changes the relevant covariance aggregates but not the underlying bargaining structure. REC banking, by contrast, adds a nonnegative marginal value to project output and therefore weakly increases equilibrium project size, with a strict increase when the marginal banking value is positive. Numerical results calibrated to market data show that the viability and buyer-group preference maps are largely preserved under the risk-aversion and REC-banking extensions, although the boundaries of the regions shift.

Managerially, the results suggest that collective PPA design should proceed in three steps. First, buyers and intermediaries should verify buyer-wise participation rather than evaluating a collective deal only by aggregate surplus or contracted capacity. Second, they should choose the governance structure based on bargaining-power profiles. Joint negotiation is the robust default when the buyer group can credibly pool leverage, whereas agent-based negotiation is useful when protecting a follower's surplus is valuable and the agent can be sufficiently empowered. Third, operational contract features should be incorporated before the bargaining structure is finalized. Matching buyers with compatible risk profiles, quantifying covariance between project-market prices and local retail prices, and enabling REC banking when feasible can improve both project sizing and implementation. These insights help explain why collective PPAs can be difficult to organize despite their aggregate benefits, and how buyers, developers, and procurement intermediaries can design collective arrangements that are both value-enhancing and individually acceptable.

Chapter 3

BATTERY DISPATCH IN ELECTRICITY MARKETS WITH NEGATIVE PRICES: FORECAST-EMBEDDED DYNAMIC PROGRAMMING

This chapter is based on joint work with Michael R. Wagner.

3.1 Introduction

Electricity markets with growing renewable penetration increasingly rely on utility-scale batteries to absorb surplus energy and release it during scarcity. For storage operators, dispatch quality is now an economically first-order operational problem where profits depend on responding well to sharp intraday price swings, while the system value of storage depends on how effectively electricity is shifted across time [Augustine and Blair, 2021, Fitzgerald et al., 2015, Newbery, 2012]. These issues are especially salient in markets where renewable output and congestion create recurrent low and sometimes negative prices [Seel et al., 2021]. In such settings, battery dispatch is not just a generic control problem. It is an industry-specific operating problem shaped by market design, storage duration, and the quality of short-horizon forecasts.

A common operating practice is to generate a point forecast and repeatedly reoptimize dispatch over a rolling horizon. This approach is transparent and implementable, but it does not value recourse explicitly. Stochastic dynamic programming (DP) is the natural alternative, yet it is often viewed as impractical for realistic storage settings because the problem combines a continuous state of charge (SoC), charging and discharging rate limits, and an exogenous price process with rich forecasting structure. The negative-price case is especially important because it removes the classical positive-price threshold shortcut and makes realistic rate-constrained storage analytically harder [Zhou et al., 2016, 2019]. More broadly, prior work on storage arbitrage under stochastic prices highlights both the appeal of stochastic control and the computational difficulties created by realistic market frictions

[Jiang and Powell, 2015, Zhou et al., 2019, Salas and Powell, 2018].

This paper studies a narrower and more operationally relevant question than the usual “DP versus heuristics” comparison. Namely, when does stochastic look-ahead add enough value relative to receding-horizon deterministic reoptimization to justify its additional modeling effort in electricity-market dispatch? That question matters for industry practice because rolling deterministic reoptimization is already a strong baseline [Secomandi, 2025], and it matters for public-policy discussions because the realized flexibility value of storage depends on the operating tools used to dispatch it in markets with recurrent negative-price episodes.

We study a price-taking storage operator in a single electricity market and develop a finite-horizon DP that is exact in the SoC dimension conditional on a finite exogenous-state graph. The key observation is not merely that the value function is piecewise affine in SoC, but that for a fixed exogenous state the Bellman update can change only when a continuation breakpoint enters or leaves the feasible charging or discharging window. This yields an endogenous event grid and a linear-time Bellman update for a fixed exogenous state. Approximation is therefore confined to the exogenous process, which we represent on a reduced graph through shock discretization, aggregation, and projection of forecasting states. The resulting architecture applies to forecasting models that admit a Markov latent-state representation; in our implementation, a 48-period problem solves in seconds.

Empirically, we benchmark the proposed DP against two deterministic controls that are operationally relevant in practice: a certainty-equivalent open-loop linear program (LP-open) and a receding-horizon deterministic reoptimization controller (LP-MPC). We also report a perfect-foresight benchmark as an unattainable upper bound. The results paint a more qualified picture than the standard claim that stochastic control dominates deterministic heuristics. In matched synthetic settings, DP delivers substantial gains over LP-open. Against LP-MPC, however, the comparison is much tighter. In real data, the performance of the stochastic policy depends critically on how forecast errors are modeled. A baseline DP that ignores dependence in forecast errors can underperform deterministic reoptimization, whereas a dependence-aware specification materially improves performance, especially for longer-duration batteries. Across days, the relative advantage of DP is strongest when

deterministic benchmark value is low and forecast uncertainty is high, and these same signals support a simple switching rule that improves over always using either controller. The paper’s message is therefore not that stochastic control should replace deterministic re-optimization by default, but that the value of stochastic look-ahead depends on market conditions, forecast-error structure, and storage duration.

Contributions. We contribute along three dimensions. First, once the exogenous price process is represented on a finite graph, we show that the storage recursion can be solved exactly in the SoC dimension on an endogenous event grid; for a fixed exogenous state, the Bellman update is linear in the number of continuation segments. Second, we provide a forecasting-to-control architecture that embeds time-series models through a Markov latent-state representation while making explicit that approximation enters only through shock discretization and projection in the exogenous dimension. Third, using out-of-sample tests in California’s real-time electricity market, we contribute an industry-study result rather than a blanket dominance claim. Specifically, stochastic look-ahead is most useful in high-uncertainty, weak-deterministic-signal regimes, becomes more attractive for longer-duration batteries, and can be deployed selectively relative to deterministic reoptimization through a simple switching rule.

The paper is therefore not offered as a generic new DP algorithm with an arbitrary application. Its purpose is to understand an electricity-market operating problem in context, in terms of how negative prices, forecast-error dependence, and storage duration jointly determine when more sophisticated stochastic look-ahead changes economically relevant dispatch decisions.

The remainder of the paper is organized as follows. Section 3.2 reviews related work on stochastic storage control. Section 3.3 presents our analytic DP formulation and efficient implementation. Section 3.4 details our numerical experiments and results, and Section 4.6 concludes with practical implications.

3.2 Literature Review

We focus on storage operation in a single market under the price-taker assumption. In this setting, the basic problem is to shift energy intertemporally in response to uncertain prices while respecting storage capacity and charging and discharging limits. A foundational stream models storage as a stochastic dynamic program or Markov decision process under exogenous price uncertainty. For example, Secomandi [2010] studies commodity storage with capacity and injection/withdrawal limits under an exogenous Markov spot-price process, and Kim and Powell [2011] analyze advance energy commitments with storage and intermittent supply under stochastic prices and wind generation. In the electricity-storage setting, Wu et al. [2012] study the price-adjusted rolling intrinsic policy for seasonal energy storage with limited flexibility, while Zhou et al. [2016] and Zhou et al. [2019] derive structural results for storage problems with negative prices and with wind, storage, and transmission constraints, respectively. Beyond operating a storage system against a single price stream, related work studies participation across sequential electricity markets such as day-ahead and intraday settings; see, for example, Löhndorf et al. [2013], Finnah et al. [2022], and Löhndorf and Wozabal [2023]. In these settings, the objective is to allocate storage flexibility across markets with different trading timelines and settlement rules.

Across these settings, a central computational challenge is the curse of dimensionality created by combining storage dynamics with rich stochastic price information and, in some applications, additional operational coupling such as wind, networks, or multiple assets. Much of the literature therefore uses approximate dynamic programming or related approximations. Lai et al. [2010] develop a tractable ADP approach for natural-gas storage valuation under uncertain price dynamics and use it to benchmark practice-based heuristics. Jiang and Powell [2015] develop a convergent ADP algorithm for battery bidding in real-time electricity markets and also study a distribution-free variant that can be trained from historical price data. Salas and Powell [2018] benchmark a scalable ADP approach for stochastic control of grid-level storage under stochastic wind, demand, and price processes. Nadarajah and Secomandi [2018] formulate merchant energy trading in a network of storage and transport assets as an MDP and develop simulation-based heuristics and bounds.

More recent work incorporates machine-learning tools for storage control and prediction-to-decision pipelines: Yi et al. [2024] study decision-focused learning for strategic energy storage, and Wang and Zhang [2018] develop a reinforcement-learning approach to real-time storage arbitrage. Relative to this literature, we exploit the geometry of the price-taking storage problem so that the endogenous SoC dimension can be solved exactly, while approximation is confined to the exogenous price-information process through discretization, aggregation, and projection.

The papers most closely related to ours are Zhou et al. [2016] and Zhou et al. [2019]. Zhou et al. [2016] show that negative prices fundamentally change the structure of optimal storage policies, but their analytical characterization relies on a fast-storage setting in which the device can fully charge or discharge within one period. Zhou et al. [2019] incorporate storage and transmission capacity in a wind-storage problem and derive structural results in the positive-price case. When negative prices are allowed, they turn to heuristic policies. These papers make clear that the combination of negative prices and rate limits is both operationally important and analytically difficult. Our paper builds on this line of work, but its goal is different. Rather than pursuing a closed-form threshold policy, we develop a finite-horizon recursion that remains exact in the storage dimension once the exogenous price process has been represented on a finite graph.

Beyond the stochastic-control literature, a related stream studies the operational role of storage in electricity systems with growing renewable penetration, congestion, and negative-price episodes. This work helps explain why storage duration, deployment context, negative-price exposure, and market design matter for storage value and system flexibility [Augustine and Blair, 2021, Fitzgerald et al., 2015, Newbery, 2012, Seel et al., 2021].

Our paper complements that literature by taking the market setting as given and asking a narrower operational question: when does richer stochastic look-ahead change battery-dispatch decisions relative to the rolling deterministic reoptimization that operators can implement in practice? Our empirical question also differs from much of the literature in the benchmark it emphasizes. In electricity-market operations, the practically relevant alternative to stochastic control is often not a static heuristic, but deterministic rolling-horizon reoptimization based on repeatedly updated point forecasts. We therefore treat

LP-MPC as the main operational comparator and use the numerical study not only to evaluate a new DP architecture, but also to identify the market conditions under which stochastic look-ahead materially changes decisions in this industry setting.

3.3 Model

This section develops the DP formulation for optimal storage dispatch under uncertain electricity prices. We first show that once the exogenous price process is represented on a finite graph, the recursion in the state-of-charge (SoC) dimension can be carried out without imposing an a priori SoC grid. We then explain how approximation enters through the exogenous price process via shock discretization, aggregation, and projection. Section 3.3.1 presents the generic finite-horizon DP model. Section 3.3.2 studies the recursion in the SoC dimension under a finite exogenous-state graph. Section 3.3.3 constructs the reduced exogenous graph, states the resulting backward-pass complexity, and clarifies the approximation introduced by shock discretization, aggregation, and projection.

3.3.1 The generic DP model

We consider a price-taking energy storage system operating over periods $k = 0, \dots, N - 1$. Following prior work on storage arbitrage under stochastic prices [Zhou et al., 2016, 2019, Jiang and Powell, 2015, Salas and Powell, 2018], we formulate the problem as a finite-horizon dynamic program. At the beginning of period k , the controller observes the state $(b_k, \mathbf{x}_k)$, where $b_k \in [0, M^e]$ is the storage state of charge (SoC) and $\mathbf{x}_k$ is the exogenous price-information state. The vector $\mathbf{x}_k$ is a Markov sufficient statistic for forecasting. Conditional on $\mathbf{x}_k$, the next realized price admits the decomposition

$$p_k = \hat{p}_k(\mathbf{x}_k) + \varepsilon_k, \quad \hat{p}_k(\mathbf{x}_k) := \mathbb{E}[p_k \mid \mathbf{x}_k],$$

and the exogenous state evolves according to

$$\mathbf{x}_{k+1} = h(\mathbf{x}_k, \varepsilon_k). \tag{3.1}$$

Here ε_k is the zero-mean period- k forecast error, which is also the one-step shock arriving between periods k and $k+1$. We allow its distribution to depend on the period (for example,

through a seasonal index), but conditional on the Markov state $\mathbf{x}_k$ the shocks are stagewise independent. Serial dependence in forecast errors can be accommodated by augmenting the state so that the remaining random shock is independent; Section 3.4.3 implements exactly this idea by augmenting $\mathbf{x}_k$ with the lagged forecast error.

The control variable u_k denotes the net change in stored energy during period k , with $u_k > 0$ corresponding to charging and $u_k < 0$ corresponding to discharging. The feasible set, denoted as $u_k \in U(b_k) := [-M^d, M^c] \cap [-b_k, M^e - b_k]$, is the intersection of two intervals: (i) the interval defined by the storage's maximum charge and discharge rate limits, $[-M^d, M^c]$; and (ii) the state-dependent energy balance constraint, $u_k \in [-b_k, M^e - b_k]$, which ensures that the next state of charge remains within the storage capacity M^e . The resulting state transition in the inventory dimension is thus

$$b_{k+1} = b_k + u_k. \quad (3.2)$$

We model the battery's period- k decision as a scalar net charge/discharge choice u_k , so the control is a net hourly position rather than separate gross charge and gross discharge quantities. This representation is appropriate for the single-market, single-settlement-price setting studied here, whose purpose is to model intertemporal arbitrage across periods rather than within-period cycling against the same price. In particular, the Bellman step compares net charging, net discharging, and no-action decisions only. If one wished to model settings with distinct intra-period prices, simultaneous participation in multiple services, or explicit within-period cycling, then a richer action representation with additional complementarity or throughput restrictions would be required.

The one-period revenue is

$$R_k(u_k, \mathbf{x}_k, \varepsilon_k) = -\max\{\eta_d u_k, u_k/\eta_c\} (\hat{p}_k(\mathbf{x}_k) + \varepsilon_k), \quad (3.3)$$

where $\eta_c, \eta_d \in (0, 1)$ denote charging and discharging efficiencies. Thus charging by $u_k > 0$ withdraws u_k/η_c units from the grid, whereas discharging by $u_k < 0$ injects $-u_k\eta_d$ units into the grid. Finally, the terminal value $J_N(b_N) := 0$. All structural results below continue to hold with any terminal value that is continuous and piecewise affine in SoC; we set $J_N \equiv 0$ for expositional simplicity.

The finite-horizon problem is

$$\max_{u_0, \dots, u_{N-1}} \mathbb{E} \left[\sum_{k=0}^{N-1} R_k(u_k, \mathbf{x}_k, \varepsilon_k) + J_N(b_N) \right].$$

The Bellman recursion is

$$J_k(b_k, \mathbf{x}_k) = \max_{u_k \in U(b_k)} \mathbb{E}_{\varepsilon_k} \left[R_k(u_k, \mathbf{x}_k, \varepsilon_k) + J_{k+1}(b_{k+1}, \mathbf{x}_{k+1}) \right].$$

Using (3.1), (3.2), (3.3) and $\mathbb{E}[\varepsilon_k | \mathbf{x}_k] = 0$, this becomes

$$J_k(b_k, \mathbf{x}_k) = \max_{u_k \in U(b_k)} \left\{ -\max\{\eta_d u_k, u_k/\eta_c\} \hat{p}_k(\mathbf{x}_k) + \mathbb{E}_{\varepsilon_k} [J_{k+1}(b_k + u_k, h(\mathbf{x}_k, \varepsilon_k))] \right\}. \quad (3.4)$$

3.3.2 Structural properties in the SoC dimension

We first use the continuous Bellman recursion to explain why negative prices eliminate the classical threshold shortcut. We then rewrite the problem on a finite exogenous-state graph and show that the resulting recursion in SoC is closed and can be computed by sweeping an endogenous event grid.

Continuous Bellman geometry

Fix period k and current exogenous state $\mathbf{x}_k$. Define the expected continuation function

$$\bar{J}_{k+1}(b_k + u_k; \mathbf{x}_k) := \mathbb{E}_{\varepsilon_k} [J_{k+1}(b_k + u_k, h(\mathbf{x}_k, \varepsilon_k)) | \mathbf{x}_k].$$

Then the Bellman recursion can be written as

$$J_k(b_k, \mathbf{x}_k) = \max_{u_k \in [-M^d, M^c] \cap [-b_k, M^e - b_k]} \left\{ -\max\{\eta_d u_k, u_k/\eta_c\} \hat{p}_k(\mathbf{x}_k) + \bar{J}_{k+1}(b_k + u_k; \mathbf{x}_k) \right\}.$$

For a given exogenous state $\mathbf{x}_k$, the problem is one-dimensional in SoC. It is convenient to separate charging and discharging and write

$$J_k(b_k, \mathbf{x}_k) = \max\{V_k^c(b_k, \mathbf{x}_k), V_k^d(b_k, \mathbf{x}_k)\}, \quad (3.5)$$

where

$$V_k^c(b_k, \mathbf{x}_k) = \max_{u_k \in [0, \min\{M^e - b_k, M^c\}]} \left\{ -\frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c} u_k + \bar{J}_{k+1}(b_k + u_k; \mathbf{x}_k) \right\}, \quad (3.6)$$

$$V_k^d(b_k, \mathbf{x}_k) = \max_{u_k \in [\max\{-b_k, -M^d\}, 0]} \left\{ -\eta_d \hat{p}_k(\mathbf{x}_k) u_k + \bar{J}_{k+1}(b_k + u_k; \mathbf{x}_k) \right\}. \quad (3.7)$$

Hence, for fixed $\mathbf{x}_k$, the Bellman step reduces to two one-dimensional optimization problems, one on the feasible charging interval and one on the feasible discharging interval.

This decomposition helps explain how the policy structure changes when forecast prices can be negative. Let $w \in \partial \bar{J}_{k+1}(\cdot; \mathbf{x}_k)$, where $\partial \bar{J}_{k+1}(\cdot; \mathbf{x}_k)$ denotes the subdifferential of the continuation function with respect to the post-decision SoC. In the charging problem, the local slope is $w - \frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c}$, whereas in the discharging problem, the local slope is $w - \eta_d \hat{p}_k(\mathbf{x}_k)$. When $\hat{p}_k(\mathbf{x}_k) \geq 0$, the two thresholds satisfy $\eta_d \hat{p}_k(\mathbf{x}_k) \leq \frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c}$, so there is a no-action region $\eta_d \hat{p}_k(\mathbf{x}_k) \leq w \leq \frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c}$, in which both mode-specific problems prefer to remain at b_k . This yields a monotone threshold structure from the positive-price setting; see, for example, Zhou et al. [2019]. When $\hat{p}_k(\mathbf{x}_k) < 0$, however, the ordering reverses to $\frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c} < \eta_d \hat{p}_k(\mathbf{x}_k)$. The no-action region then disappears. In this case, when $\frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c} \leq w \leq \eta_d \hat{p}_k(\mathbf{x}_k)$, the charging problem prefers higher post-decision inventory, while the discharging problem prefers lower post-decision inventory, so charging and discharging boundary actions may compete directly. Accordingly, the monotone threshold structure in the positive-price case no longer applies. Since negative prices therefore preclude the classical positive-price threshold shortcut, the Bellman step must compare the charging and discharging envelopes directly.

Finite-state reformulation

To obtain a computable recursion, we next replace the continuous exogenous transition induced by $h(\mathbf{x}_k, \varepsilon_k)$ with a discrete finite-state approximation. Once the exogenous process is represented on a finite graph, the remaining recursion is only in SoC.

Suppose that, at each period k , the exogenous price information is represented by a finite state set $\mathcal{X}_k$, and the one-step transition probabilities are $P_k(\mathbf{x}_k, \mathbf{x}')$ for $\mathbf{x}_k \in \mathcal{X}_k$ and $\mathbf{x}' \in \mathcal{X}_{k+1}$. The Bellman recursion can then be written as

$$\tilde{J}_k(b_k, \mathbf{x}_k) = \max_{u_k \in U(b_k)} \left\{ -\max\{\eta_d u_k, u_k/\eta_c\} \hat{p}_k(\mathbf{x}_k) + \sum_{\mathbf{x}' \in \mathcal{X}_{k+1}} P_k(\mathbf{x}_k, \mathbf{x}') \tilde{J}_{k+1}(b_k + u_k, \mathbf{x}') \right\}, \quad (3.8)$$

with terminal condition $\tilde{J}_N(b_N) \equiv 0$ for all $\mathbf{x} \in \mathcal{X}_N$. For the structural results below, it is convenient to rewrite the finite-state Bellman recursion in the form of post-decision SoC,

$b' = b_k + u_k$. Define the finite-state continuation function

$$\tilde{J}_{k+1}(b'; \mathbf{x}_k) := \sum_{\mathbf{x}' \in \mathcal{X}_{k+1}} P_k(\mathbf{x}_k, \mathbf{x}') \tilde{J}_{k+1}(b', \mathbf{x}'), \quad b' \in [0, M^e].$$

Then (3.8) can be written as

$$\tilde{J}_k(b_k, \mathbf{x}_k) = \max\{\tilde{V}_k^c(b_k, \mathbf{x}_k), \tilde{V}_k^d(b_k, \mathbf{x}_k)\}, \quad (3.9)$$

where $\tilde{V}_k^c(b_k, \mathbf{x}_k)$ and $\tilde{V}_k^d(b_k, \mathbf{x}_k)$ are defined similarly as in (3.6) – (3.7).

$$\tilde{V}_k^c(b_k, \mathbf{x}_k) = \max_{b' \in [b_k, \min\{M^e, b_k + M^e\}]} \left\{ -\frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c} (b' - b_k) + \tilde{J}_{k+1}(b'; \mathbf{x}_k) \right\}, \quad (3.10)$$

$$\tilde{V}_k^d(b_k, \mathbf{x}_k) = \max_{b' \in [\max\{0, b_k - M^d\}, b_k]} \left\{ -\eta_d \hat{p}_k(\mathbf{x}_k) (b' - b_k) + \tilde{J}_{k+1}(b'; \mathbf{x}_k) \right\}. \quad (3.11)$$

Thus, for fixed $\mathbf{x}_k$, the finite-state Bellman step again reduces to two one-dimensional optimization problems over post-decision SoC, one over the feasible charging window and one over the feasible discharging window.

Structural and computational results

The finite-state reformulation reduces the problem, for each fixed exogenous state $\mathbf{x}_k \in \mathcal{X}_k$, to a one-dimensional recursion in SoC. The next three results explain why this recursion is computationally manageable. Proposition 3.1 shows that backward induction preserves a continuous piecewise-affine value function in SoC. Proposition 3.2 then identifies the only SoC values at which the update can change form. These values define an endogenous event grid determined by the continuation breakpoints and the charging and discharging rate limits. Corollary 3.3 combines the two results and shows that, for a fixed exogenous state, the recursion can be computed by a linear-time sweep over this event grid.

The first question is whether the Bellman recursion preserves a manageable functional form in the SoC dimension. If the recursion destroyed piecewise affinity after one backward step, then avoiding an a priori SoC grid would buy little. The next result shows that, once the exogenous process is represented on a finite graph, backward induction preserves a continuous piecewise-affine value function and a piecewise-affine optimal policy in SoC.

Proposition 3.1 (Piecewise-affine closure in SoC). *Fix a finite exogenous-state graph $\{\mathcal{X}_k, P_k\}_{k=0}^{N-1}$, then for each $k = 0, \dots, N-1$ and each $\mathbf{x}_k \in \mathcal{X}_k$, the value function $\tilde{J}_k(b_k, \mathbf{x}_k)$ is continuous and piecewise affine on $b_k \in [0, M^e]$. Moreover, there exists an optimal policy $u_k^*(b_k, \mathbf{x}_k)$ that is piecewise affine on $b_k \in [0, M^e]$.*

Proposition 3.1 establishes closure, but closure alone is not enough for tractability. A piecewise-affine continuation function could still generate many candidate switching points in the Bellman step. What we need next is a sharp description of where the update can actually change as current SoC, b_k varies. The next result shows that these changes can occur only on an endogenous event grid induced by continuation breakpoints and the charge/discharge rate limits. Specifically, in the post-decision formulation (3.10)–(3.11), the feasible charging and discharging windows move with b_k . Fixing $\mathbf{x}_k$, a qualitative change to $\tilde{V}_k^c$ and $\tilde{V}_k^d$ can occur only when an edge of one of these windows hits a breakpoint of $\tilde{J}_{k+1}(\cdot; \mathbf{x}_k)$, which is also piecewise affine following Proposition 3.1. This observation leads to the event grid below.

Proposition 3.2 (Event grid for the SoC update). *Fix a period k and an exogenous state $\mathbf{x}_k \in \mathcal{X}_k$. Suppose the expected continuation function $\tilde{J}_{k+1}(b'; \mathbf{x}_k)$ is continuous and piecewise affine on $b' \in [0, M^e]$ with breakpoints $0 = s_0 < s_1 < \dots < s_n = M^e$. Define the event grid*

$$E_k(\mathbf{x}_k) := \left(\{s_j\}_{j=0}^n \cup \{s_j - M^c\}_{j=0}^n \cup \{s_j + M^d\}_{j=0}^n \right) \cap [0, M^e].$$

Then, on any interval of b_k between two consecutive points of $E_k(\mathbf{x}_k)$, $\tilde{J}_k(b_k, \mathbf{x}_k)$ is the upper envelope of at most five affine functions corresponding to: (i) staying at the current SoC, (ii) charging to the upper boundary of the feasible charging window, (iii) charging to the best reachable interior breakpoint of $\tilde{J}_{k+1}(\cdot; \mathbf{x}_k)$, (iv) discharging to the lower boundary of the feasible discharging window, and (v) discharging to the best reachable interior breakpoint of $\tilde{J}_{k+1}(\cdot; \mathbf{x}_k)$.

Consequently, the Bellman update can change form only at event-grid points or at intersections of these affine candidates within an event interval.

Proposition 3.2 shows that a continuation breakpoint becomes relevant at period k only when it enters or leaves a feasible charging or discharging window. These entry and exit

locations are exactly $s_j - M^c$, s_j , and $s_j + M^d$. Between two consecutive event-grid points, the same set of continuation breakpoints remains reachable, so the Bellman step reduces to comparing a fixed small set of affine candidates. The remaining question is whether they can be scanned without repeatedly solving local optimization problems from scratch. The next corollary shows that the Bellman update for a fixed exogenous state reduces to a linear-time sweep over the event grid.

Corollary 3.3 (Linear-time Bellman update for a fixed exogenous state). *If $\tilde{J}_{k+1}(\cdot; \mathbf{x}_k)$ has n affine segments, then the event grid $E_k(\mathbf{x}_k)$ has at most $3n + 3$ points, and $\tilde{J}_k(\cdot, \mathbf{x}_k)$ together with an optimal policy $u_k^*(\cdot, \mathbf{x}_k)$ can be computed in $O(n)$ time and $O(n)$ memory. In particular, the updated value function $\tilde{J}_k(\cdot, \mathbf{x}_k)$ has $O(n)$ affine segments.*

Piecewise-affine closure alone does not make the Bellman step computationally useful. In principle, a piecewise-affine continuation function could still require searching over many candidate breakpoints. Proposition 3.2 and Corollary 3.3 show why that does not happen here. In the rate-constrained storage problem with possibly negative prices, the Bellman update can change only when a continuation breakpoint enters or leaves the feasible charging or discharging window, which yields an endogenous event grid of size at most $3n + 3$ when the continuation function has n affine segments. Between consecutive event points, the same fixed small set of affine candidates remains relevant, so the update reduces to a one-dimensional linear-time sweep. Thus, the contribution is not merely that $\tilde{J}_k(\cdot, \mathbf{x}_k)$ is piecewise affine, but that the relevant breakpoints and candidate actions are explicitly controlled, yielding an exact recursion in the SoC dimension conditional on the finite exogenous-state graph.

3.3.3 Deterministic reduction of the exogenous state

The structural results in Section 3.3.2 resolve the SoC dimension once the exogenous process is represented on a finite state graph. What remains is the exogenous dimension. Under realistic forecasting models, the state $\mathbf{x}_k$ is continuous and potentially high-dimensional. The goal of this subsection is therefore to construct, for each period k , a finite set of representative exogenous states that preserves the price information most relevant for storage

decisions.

A natural baseline is to simulate many sample paths of the forecasting model and cluster the resulting states, for example by applying k -means directly to the raw state vectors. We do not pursue that approach for two reasons. First, the Bellman expectation would then be evaluated by Monte Carlo, introducing simulation noise into the reduced state graph and making the approximation sensitive to the particular sample. Second, closeness in the raw state vector need not translate into closeness in dispatch value, especially when the state includes lagged prices, residuals, and seasonal components whose effect on operations is indirect. Numerical evidence on this point is reported in Section 3.4.5. We therefore adopt a deterministic reduced-state graph through aggregation.

The construction proceeds in three steps. First, because all uncertainty enters through the scalar shock ε_k , we replace its distribution with a small deterministic support. Second, to prevent the number of candidate states from growing exponentially over time, we aggregate the propagated states to satisfy a fixed computational budget. Third, we project each propagated child state onto a representative state. The resulting representative states, together with a projected transition kernel, define a reduced deterministic state graph that can be solved efficiently by dynamic programming.

Deterministic shock discretization

Proposition 3.1 is stated for an abstract finite-state transition kernel $P_k(\mathbf{x}_k, \mathbf{x}')$. We first describe the shock discretization used to build that kernel. The projection step that maps propagated child states to representative states is defined in Section 3.3.3.

For each period k , let F_k denote the distribution of the continuous one-step shock ε_k . We approximate F_k by a discrete distribution $\tilde{F}_k$ with discretization points $\{\epsilon_{k1}, \dots, \epsilon_{kQ_k}\}$ and weights $\{\pi_{k1}, \dots, \pi_{kQ_k}\}$, where

$$\mathbb{P}_{\tilde{F}_k}(\varepsilon_k = \epsilon_{kq}) = \pi_{kq}, \quad q = 1, \dots, Q_k, \quad \pi_{kq} \geq 0, \quad \sum_{q=1}^{Q_k} \pi_{kq} = 1.$$

In the implementation, these discretization points and weights are obtained by period-specific quantile discretization of the fitted residual distribution.

Given a current exogenous state $\mathbf{x}_k$ and a discretized shock node ϵ_{kq} , the forecasting model generates the child state $h(\mathbf{x}_k, \epsilon_{kq})$. Given the discretized shock nodes, the child states follow directly from the forecasting model. In our setting, this step is numerically simple because all uncertainty enters through a single scalar shock.

State space reduction and aggregation

With the shock distribution discretized, each representative state at period k generates at most Q_k candidate children at period $k + 1$. Because the exogenous process evolves independently of the storage action, this forward construction depends only on the forecasting model and the discretized shock support. The total number of candidates can nevertheless grow rapidly with k , so at each period we aggregate them to a reduced graph according to a fixed state budget.

Let $\mathcal{X}_k$ denote the representative exogenous states retained at period k , and let $\mu_k(\mathbf{x}_k)$ denote the probability mass carried by $\mathbf{x}_k \in \mathcal{X}_k$. We initialize at the observed initial state

$$\mathcal{X}_0 = \{\mathbf{x}_0\}, \quad \mu_0(\mathbf{x}_0) = 1.$$

Given $\mathcal{X}_k$, we generate the raw candidate states for period $k + 1$ by propagating each representative state through the discretized shock support

$$\mathcal{C}_{k+1} = \{h(\mathbf{x}_k, \epsilon_{kq}) : \mathbf{x}_k \in \mathcal{X}_k, q = 1, \dots, Q_k\}.$$

Each candidate $\mathbf{y} \in \mathcal{C}_{k+1}$ receives probability mass

$$\tilde{\mu}_{k+1}(\mathbf{y}) = \sum_{\mathbf{x}_k \in \mathcal{X}_k} \sum_{q=1}^{Q_k} \mu_k(\mathbf{x}_k) \pi_{kq} \mathbf{1}\{h(\mathbf{x}_k, \epsilon_{kq}) = \mathbf{y}\}.$$

At this stage, the only approximation is the shock discretization $F_k \mapsto \tilde{F}_k$. Given the discretized shock nodes, the candidate children follow directly from the forecasting model.

We next decide which candidate states in period $k + 1$ should be retained as representatives. In time-series models, two latent states can be close in Euclidean distance yet imply very different current and near-term price forecasts. The reverse can also occur, where two states can be far apart in the raw coordinates while implying nearly identical dispatch

decisions. For storage arbitrage, the economically relevant object is therefore not the raw state vector itself, but the forecast path that it implies. We therefore aggregate candidate states in a dispatch-relevant forecast space.

For each candidate state $\mathbf{y} \in \mathcal{C}_{k+1}$, let $g_{k+1}(\mathbf{y})$ denote a model-based summary of the near-term forecast path implied by $\mathbf{y}$. In the experiments, $g_{k+1}(\mathbf{y})$ is a short vector of conditional point forecasts over a fixed look-ahead window after period $k + 1$. We then define the aggregation feature vector

$$\phi_{k+1}(\mathbf{y}) = (\lambda_p \hat{p}_{k+1}(\mathbf{y}), \lambda_f g_{k+1}(\mathbf{y})), \quad (3.12)$$

where $\lambda_p, \lambda_f > 0$ are scaling weights. The first term captures the immediate forecast of period $k + 1$, while the second term captures the short-horizon forecast trajectory, which is informative about continuation value. Note that the feature vector is only used for selecting representative states and the Bellman operation still happens on the original state vector.

To compare feature components on similar scales, let D_{k+1} be the diagonal matrix of empirical componentwise standard deviations of the feature vectors $\{\phi_{k+1}(\mathbf{y}) : \mathbf{y} \in \mathcal{C}_{k+1}\}$, with any zero diagonal entry replaced by one. We define the distance between candidate states by

$$d_{k+1}(\mathbf{y}, \mathbf{y}') := \|D_{k+1}^{-1}(\phi_{k+1}(\mathbf{y}) - \phi_{k+1}(\mathbf{y}'))\|_2. \quad (3.13)$$

Given a prescribed state budget M_{k+1} , we choose the representative states by a weighted k -medoids heuristic that approximately solves

$$\min_{\mathcal{S} \subseteq \mathcal{C}_{k+1}, |\mathcal{S}|=M_{k+1}} \sum_{\mathbf{y} \in \mathcal{C}_{k+1}} \tilde{\mu}_{k+1}(\mathbf{y}) \min_{\mathbf{s} \in \mathcal{S}} d_{k+1}(\mathbf{y}, \mathbf{s}). \quad (3.14)$$

We use medoids rather than centroids so that every representative remains an admissible model state that can be propagated exactly through $h(\cdot, \cdot)$ [Kaufman and Rousseeuw, 1990]. This is important because the reduced graph is constructed by repeated forward propagation of representative states.

Let $\mathcal{X}_{k+1} \subseteq \mathcal{C}_{k+1}$ denote the selected medoid set, and let

$$A_{k+1}(\mathbf{y}) \in \arg \min_{\mathbf{s} \in \mathcal{X}_{k+1}} d_{k+1}(\mathbf{y}, \mathbf{s})$$

be the nearest-representative projection. The representative-state masses are then obtained by aggregating candidate masses, which will be used to weight the aggregation step in the next period.

$$\mu_{k+1}(\mathbf{x}') = \sum_{\mathbf{y} \in \mathcal{C}_{k+1}: A_{k+1}(\mathbf{y}) = \mathbf{x}'} \tilde{\mu}_{k+1}(\mathbf{y}), \quad \mathbf{x}' \in \mathcal{X}_{k+1}.$$

The same map induces the projection used in the Bellman recursion. For a current representative state $\mathbf{x}_k \in \mathcal{X}_k$ and shock node ϵ_{kq} , define

$$\Pi_k(\mathbf{x}_k, \epsilon_{kq}) = A_{k+1}(h(\mathbf{x}_k, \epsilon_{kq})) \in \mathcal{X}_{k+1}.$$

Equivalently, the reduced one-step transition kernel is

$$P_k(\mathbf{x}_k, \mathbf{x}') = \sum_{q=1}^{Q_k} \pi_{kq} \mathbf{1}\{A_{k+1}(h(\mathbf{x}_k, \epsilon_{kq})) = \mathbf{x}'\}, \quad \mathbf{x}_k \in \mathcal{X}_k, \mathbf{x}' \in \mathcal{X}_{k+1}.$$

Thus, $P_k(\mathbf{x}_k, \mathbf{x}')$ records how the shock mass associated with a current representative state $\mathbf{x}_k$ is redistributed across next-period representatives after projection.

With this notation, the implemented Bellman recursion is

$$\tilde{J}_k(b_k, \mathbf{x}_k) = \max_{u_k \in U(b_k)} \left\{ -\max\{\eta_d u_k, u_k/\eta_c\} \hat{p}_k(\mathbf{x}_k) + \sum_{q=1}^{Q_k} \pi_{kq} \tilde{J}_{k+1}(b_k + u_k, \Pi_k(\mathbf{x}_k, \epsilon_{kq})) \right\}, \quad \mathbf{x}_k \in \mathcal{X}_k. \quad (3.15)$$

The only approximation in this recursion comes from shock discretization and projection onto representative states.

In words, the forward pass first generates candidate children under the discretized shock support and assigns them probability masses $\tilde{\mu}_{k+1}$. It then aggregates those candidates to the prescribed state budget using a forecast-based distance and projects each child state onto a representative. This yields both the next representative set $\mathcal{X}_{k+1}$ and the reduced transition kernel P_k . The backward pass is then a standard finite-state dynamic program on the resulting graph, with the SoC recursion solved using the results of Section 3.3.2.

Algorithm 1 summarizes the forward construction of the reduced exogenous-state graph. Starting from the observed initial state, the procedure propagates representative states through the discretized shock support, aggregates the resulting candidate children in the forecast-feature space, and records the induced reduced transition kernel.

Algorithm 1 Forward construction of the reduced exogenous-state graph

- 1: For each period k , estimate the one-step shock distribution F_k and construct a deterministic support

$$\{(\epsilon_{kq}, \pi_{kq})\}_{q=1}^{Q_k}$$

by quantile discretization.

- 2: Initialize

$$\mathcal{X}_0 = \{\mathbf{x}_0\}, \quad \mu_0(\mathbf{x}_0) = 1.$$

- 3: **for** $k = 0, \dots, N - 2$ **do**

- 4: Generate the candidate set

$$\mathcal{C}_{k+1} = \{h(\mathbf{x}_k, \epsilon_{kq}) : \mathbf{x}_k \in \mathcal{X}_k, q = 1, \dots, Q_k\}.$$

- 5: Compute candidate masses $\tilde{\mu}_{k+1}$ by aggregating inherited mass over duplicate children.
6: Compute the aggregation features ϕ_{k+1} , the scaling matrix D_{k+1} , and the distance d_{k+1} from (3.12)–(3.13).
7: Run weighted k -medoids with budget M_{k+1} to select $\mathcal{X}_{k+1} \subseteq \mathcal{C}_{k+1}$.
8: Assign each candidate $\mathbf{y} \in \mathcal{C}_{k+1}$ to

$$A_{k+1}(\mathbf{y}) \in \arg \min_{\mathbf{s} \in \mathcal{X}_{k+1}} d_{k+1}(\mathbf{y}, \mathbf{s}).$$

- 9: Aggregate candidate masses to obtain the representative-state masses μ_{k+1} .
10: Define the projection map

$$\Pi_k(\mathbf{x}_k, \epsilon_{kq}) = A_{k+1}(h(\mathbf{x}_k, \epsilon_{kq}))$$

and the reduced transition kernel

$$P_k(\mathbf{x}_k, \mathbf{x}') = \sum_{q=1}^{Q_k} \pi_{kq} \mathbf{1}\{\Pi_k(\mathbf{x}_k, \epsilon_{kq}) = \mathbf{x}'\}.$$

- 11: **end for**
-

Given the reduced state graph $\{\mathcal{X}_k, P_k\}$, the backward pass solves the resulting finite-state DP by standard backward induction. For each representative exogenous state $\mathbf{x}_k \in \mathcal{X}_k$, define

$$\tilde{J}_{k+1}(\cdot; \mathbf{x}_k) = \sum_{\mathbf{x}' \in \mathcal{X}_{k+1}} P_k(\mathbf{x}_k, \mathbf{x}') \tilde{J}_{k+1}(\cdot, \mathbf{x}').$$

Applying the one-dimensional SoC recursion from Section 3.3.2 then yields $\tilde{J}_k(\cdot, \mathbf{x}_k)$ and an associated optimal policy $u_k^*(\cdot, \mathbf{x}_k)$ for every representative state.

For $\mathbf{x}_k \in \mathcal{X}_k$, let $n_k(\mathbf{x}_k)$ denote the number of affine segments of $\tilde{J}_k(\cdot, \mathbf{x}_k)$, and define

$$n_k^{\max} := \max_{\mathbf{x}_k \in \mathcal{X}_k} n_k(\mathbf{x}_k), \quad m_k := |\mathcal{X}_k|.$$

Once the SoC recursion is under control, the main remaining computational burden comes from assembling continuation functions across representative exogenous states and shock nodes. The next result makes that burden explicit and states the backward-pass complexity on the pre-constructed reduced graph.

Proposition 3.4 (Backward-pass complexity). *Fix the discretized shock supports $\{(\epsilon_{kq}, \pi_{kq})\}_{q=1}^{Q_k}$ and suppose that the reduced exogenous graph, including the projection maps Π_k , has already been constructed. Treat Q_k as a fixed discretization parameter. Then the continuation function for a fixed representative state $x_k \in X_k$ can be written as*

$$\tilde{\tilde{J}}_{k+1}(\cdot; x_k) = \sum_{q=1}^{Q_k} \pi_{kq} \tilde{J}_{k+1}(\cdot, \Pi_k(x_k, \epsilon_{kq})).$$

Hence it is assembled from at most Q_k child value functions. The backward pass over the reduced graph can therefore be executed in

$$O\left(\sum_{k=0}^{N-1} m_k Q_k n_{k+1}^{\max}\right)$$

time and

$$O\left(\sum_{k=0}^N m_k n_k^{\max}\right)$$

memory, excluding storage of the precomputed reduced exogenous graph.

Proposition 3.4 makes clear what the reduced graph buys computationally. Once the graph is built, the remaining workload is governed by the number of representative states, discretized shock nodes, and continuation segments.

Approximation from the reduced exogenous graph

The preceding subsection makes clear where approximation does and does not enter the model. Once the exogenous process is represented on a finite graph, the recursion in the SoC dimension introduces no further approximation. The only approximation in our solution procedure comes from the reduced exogenous graph. In particular, we approximate the continuous shock distribution by a finite support, and then replace each propagated child state by a nearby representative state. This subsection makes those two sources of approximation explicit.

Let $\mathbb{X}_k$ denote the full exogenous state space at period k , and let $\mathcal{X}_k \subseteq \mathbb{X}_k$ denote the representative states retained in the reduced graph. Using the same feature map ϕ_{k+1} and distance d_{k+1} introduced in (3.12)–(3.13), extend the nearest-representative rule to the full space by defining

$$A_{k+1}(\mathbf{y}) \in \arg \min_{\mathbf{s} \in \mathcal{X}_{k+1}} d_{k+1}(\mathbf{y}, \mathbf{s}), \quad \mathbf{y} \in \mathbb{X}_{k+1}.$$

For representative current states $\mathbf{x}_k \in \mathcal{X}_k$, the implemented map in (3.15) is simply

$$\Pi_k(\mathbf{x}_k, \epsilon_{kq}) = A_{k+1}(h(\mathbf{x}_k, \epsilon_{kq})).$$

Now fix any bounded continuation function

$$J : [0, M^e] \times \mathbb{X}_{k+1} \rightarrow \mathbb{R}.$$

The exact one-step Bellman update is

$$(T_k J)(b, \mathbf{x}) = \max_{u \in U(b)} \left\{ -\max\{\eta_d u, u/\eta_c\} \hat{p}_k(\mathbf{x}) + \int J(b + u, h(\mathbf{x}, \epsilon)) dF_k(\epsilon) \right\},$$

whereas the update induced by the reduced graph is

$$(\tilde{T}_k J)(b, \mathbf{x}) = \max_{u \in U(b)} \left\{ -\max\{\eta_d u, u/\eta_c\} \hat{p}_k(\mathbf{x}) + \sum_{q=1}^{Q_k} \pi_{kq} J(b + u, A_{k+1}(h(\mathbf{x}, \epsilon_{kq}))) \right\}.$$

To separate the two sources of approximation, define

$$\Delta_k^{\text{disc}}(J) := \sup_{b \in [0, M^e], \mathbf{x} \in \mathbb{X}_k} \sup_{u \in U(b)} \left| \int J(b + u, h(\mathbf{x}, \epsilon)) dF_k(\epsilon) - \sum_{q=1}^{Q_k} \pi_{kq} J(b + u, h(\mathbf{x}, \epsilon_{kq})) \right|,$$

and

$$\Delta_k^{\text{proj}}(J) := \sup_{b \in [0, M^e], \mathbf{x} \in \mathbb{X}_k} \sup_{u \in U(b)} \sum_{q=1}^{Q_k} \pi_{kq} \left| J(b + u, h(\mathbf{x}, \epsilon_{kq})) - J(b + u, A_{k+1}(h(\mathbf{x}, \epsilon_{kq}))) \right|.$$

A direct add-and-subtract argument gives the following decomposition.

Proposition 3.5 (Bellman error decomposition). *For any bounded continuation function J ,*

$$\|T_k J - \tilde{T}_k J\|_\infty \leq \Delta_k^{\text{disc}}(J) + \Delta_k^{\text{proj}}(J), \quad (3.16)$$

where $\|\cdot\|_\infty$ denotes the supremum norm over $[0, M^e] \times \mathbb{X}_k$.

Proposition 3.5 shows that the reduced graph affects the Bellman step only through shock discretization and projection onto representative exogenous states. Conditional on the reduced graph, the recursion in the SoC dimension is exact. This result also helps motivate our aggregation rule. Because the quality of the projection matters directly for the Bellman step, representative states should be selected using a notion of similarity that preserves the price information relevant for storage decisions. Two raw model states can be close in their latent coordinates yet imply quite different current and near-term forecasts. The feature map in (3.12) is therefore constructed to group states by dispatch-relevant forecast information rather than by raw statistical proximity alone.

3.4 Electricity-Market Setting and Numerical Experiments

Using data from California’s real-time electricity market, this section answers four questions in sequence. First, how does the proposed DP compare with operationally relevant deterministic benchmarks? Second, if gains observed in matched synthetic data fail to transfer to realized prices, is the gap driven by misspecified error dynamics rather than by stochastic control per se? Third, once a dependence-aware specification is introduced, when is stochastic look-ahead most likely to add value relative to deterministic benchmarks? Fourth, which ingredients of the proposed method are necessary to capture that value? The subsections below follow this logic from introducing the market setting and experimental setup (Section 3.4.1–3.4.2), average performance and diagnosis (Section 3.4.3), to regime dependence (Section 3.4.4), to ablation (Section 3.4.5).

3.4.1 Electricity-market setting and forecasting model

We use hourly prices from the CAISO real-time market (RTM), specifically the NP-15 node, from January 1, 2023 through December 31, 2024. The raw market data are available at 15-minute frequency and are aggregated to hourly averages. We use 2023 for model estimation and 2024 for out-of-sample evaluation. In the 2024 evaluation sample, negative prices are operationally meaningful rather than anecdotal: 4.5% of hourly prices are below zero, the 5th percentile is \$2.1MWh, and the 95th percentile is \$73.7/MWh. These sample features make the NP-15 real-time market a useful setting for studying when stochastic look-ahead changes dispatch decisions.

This setting is useful not merely as a source of data, but as an industry environment in which the paper’s question is operationally meaningful: negative prices recur, short-horizon forecast updates matter, and battery dispatch is shaped by renewable-driven volatility and congestion. For a price-taking storage operator, these features make deterministic reoptimization a credible practical baseline while also creating scope for stochastic look-ahead to add value when uncertainty is persistent. The California real-time market therefore provides a natural test bed for asking not only whether stochastic control can work, but when it is worth the additional modeling effort relative to simpler rolling-horizon dispatch.

For illustration, we use an additive HW model with seasonal period $m = 24$. The smoothing parameters (α, β, γ) are selected by blocked time-series cross-validation with a 24-hour forecast horizon. Forecast accuracy is measured by relative mean absolute error (rMAE), defined as

$$\text{rMAE} = \frac{\sum_{t=1}^N |y_t - \hat{y}_t|}{\sum_{t=1}^N |y_t - y_{t-7m}|},$$

following Lago et al. [2021]. In our CAISO sample, the resulting HW model attains an out-of-sample rMAE of 0.73.

The framework in Section 3.3 only requires a Markov forecasting state, a conditional mean forecast, and a one-step transition law. To align the numerical notation with Section 3.3, we index the HW state as

$$\mathbf{x}_k^{\text{HW}} := (l_{k-1}, t_{k-1}, \mathbf{s}_{k-1}),$$

where $\mathbf{s}_{k-1} \in \mathbb{R}^{24}$ stores the seasonal factors. The conditional mean forecast is then

$$\hat{p}_k(\mathbf{x}_k^{\text{HW}}) = l_{k-1} + t_{k-1} + s_{k-m}.$$

We allow the one-step shock variance to vary by hour of day, estimated from training residuals. This delivers the Markov state representation required by the DP. Because the numerical section focuses on policy performance rather than forecast-model development, details of the HW state-transition embedding are given in Appendix B.2.

3.4.2 Experimental design and benchmark policies

We study batteries with durations $H \in \{2, 4, 8\}$ hours. With energy capacity normalized to M^e , the corresponding power limits are $M^c = M^d = M^e/H$. Charging and discharging efficiencies are set to $\eta_c = \eta_d = 0.9$. Each policy is evaluated over a 48-hour horizon.

Environments. We report results in three environments. For the *matched synthetic environment*, we randomly sample 24 starting dates from the 2024 test period (6 for each quarter), and generate 50 synthetic 48-hour price paths from the fitted HW model. This is the model-consistent environment in which the same forecast model defines both the controller and the data-generating process. The *bridge-to-real environment* is next designed to study why matched synthetic gains may not transfer to realized prices. We keep the same initial forecast states but replace i.i.d. error sampling with residual block bootstrap. Specifically, we resample realized residuals from model fitting in contiguous blocks of length $L \in \{1, 4, 8, 12, 24\}$. This preserves the fitted mean model while injecting serial dependence into the forecast errors. Finally, in the *real-data environment*, we evaluate policies directly on realized 2024 CAISO prices.

Policies. We compare four policies. The *DP* policy is the proposed controller that solves the reduced-state stochastic DP at the start of each 48-hour window. During execution, the realized price updates the continuous forecasting state. We denote the resulting period- k state by $\mathbf{x}_k$, and its nearest-representative projection onto the reduced state graph by

$$\hat{\mathbf{x}}_k := A_k(\mathbf{x}_k).$$

The stored policy is then applied in closed loop at $\hat{\mathbf{x}}_k$. The *LP-open* policy solves a certainty-equivalent deterministic storage problem using the conditional mean forecast path over the full 48-hour horizon and then executes the resulting plan in open loop, without reoptimization. Finally, the *LP-MPC* policy also solves a certainty-equivalent deterministic storage problem, but reoptimizes every hour. After each realized price, the HW state is updated, a fresh deterministic forecast is formed, and the problem is re-solved over the remaining horizon. LP-MPC is therefore the strongest deterministic benchmark and the main operational comparator. The distinction between DP and LP-MPC is important. Both use the same observed price stream, but LP-MPC reoptimizes from the exact updated forecast state every hour, whereas DP is solved once and acts on a reduced stochastic graph. The comparison is therefore not merely stochastic versus deterministic optimization; it is also reduced-state stochastic control versus exact-state receding-horizon reoptimization. As an upper bound, we further solve the same deterministic LP problem using the realized future price path and refer to this setting as *LP-perfect*.

Inference. All significance markers in Tables 3.1 and 3.2 are based on two-sided paired *t*-tests of day-level profit differences relative to the named benchmark. In the matched synthetic environment, we first average each policy’s profit across the 50 simulated paths within a starting date and then test across the 24 starting dates; in real data, the unit of inference is the day.

3.4.3 Overall performance

Table 3.1 reports average relative profit differences between the proposed DP policy and the two LP benchmarks, with positive values favoring DP. The matched synthetic results establish the first benchmark fact. Relative to LP-open, DP delivers large and statistically significant profit gains across all battery durations. This confirms that stochastic recourse creates value relative to a forecast-then-optimize policy that commits to a single mean path. At the same time, LP-MPC is difficult to beat in the matched synthetic setting. Across all three battery durations, LP-MPC outperforms DP. This is because LP-MPC refreshes the forecast from the exact updated state each hour, while DP relies on a compressed

representative-state graph constructed at the start of the horizon. In a model-consistent environment, that approximation cost can outweigh the gain from stochastic optimization. This benchmark is informative: our contribution is not that a reduced-state stochastic controller must mechanically dominate an exact-state deterministic controller, but that a tractable stochastic controller becomes more attractive when uncertainty is persistent and the error process is modeled well. However, when evaluated on real market prices, the proposed baseline DP calibrated under stagewise-independent one-step shocks loses much of its synthetic advantage and can even underperform LP-open, as shown in the right panel of Table 3.1.

Table 3.1: Baseline DP under independent one-step shocks versus LP benchmarks.

	Matched Synthetic			Real		
	2h	4h	8h	2h	4h	8h
Proposed DP vs. LP-Open	+63.2%***	+40.6%***	+26.6%***	−9.3%***	−4.2%**	−5.7%**
Proposed DP vs. LP-MPC	−34.7%***	−27.3%***	−26.5%***	−10.7%***	−5.6%**	−8.9%***
n (days)	24	24	24	182	182	182

Note: Relative profit differences are computed as $\overline{(\text{DP} - \text{LP})} / |\overline{\text{LP}}|$, where the overbar denotes the sample average over days and LP denotes the relevant benchmark in that row. Significance stars are based on two-sided paired t -tests of daily profit differences.

Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

A natural explanation is that in realized CAISO prices, forecast errors may exhibit serial dependence that this baseline specification does not capture. Under that misspecification, stochastic optimization around the wrong error law can underperform both LP-open and LP-MPC, even though the same DP adds value in the matched synthetic environment. To test whether the synthetic-to-real gap is driven by misspecified error dependence rather than by stochastic optimization per se, we construct the bridge-to-real environment (using 4-hour storage for illustration), in which we keep the same initial forecast states and fitted mean

model but replace independent error sampling with residual block bootstrap. Specifically, we resample realized residuals in contiguous blocks of length $L \in \{1, 4, 8, 12, 24\}$. This isolates the role of dependence in forecast errors while holding fixed the rest of the modeling pipeline.




Figure 3.1: Bridge-to-real via residual block bootstrap

Figure 3.1 supports this interpretation. Starting from the same initial forecast states as in the matched synthetic experiment, we gradually inject serial dependence through longer residual blocks. As the block length grows, the DP - LP-open profit gap shrinks rapidly. We therefore interpret the baseline real-data underperformance as a diagnostic failure mode of the independence specification, not as evidence that stochastic control is intrinsically dominated by deterministic reoptimization. This motivates the dependence-aware DP below, which is the main real-data specification in the remainder of this subsection.

Specifically, in the main real-data comparison, we use a dependence-aware version of the DP. Let

$$\mathbf{x}_k^{\text{HW}} := (l_{k-1}, t_{k-1}, \mathbf{s}_{k-1})$$

denote the baseline HW forecasting state and let

$$e_{k-1} := p_{k-1} - \hat{p}_{k-1}(\mathbf{x}_{k-1}^{\text{HW}})$$

denote the previous forecast error. We augment the forecasting state to

$$\mathbf{x}_k^{\text{dep}} := (\mathbf{x}_k^{\text{HW}}, e_{k-1}).$$

As a first step, we model the residual dependence with a simple AR(1) structure, so the transition law is updated accordingly but the problem remains Markov under $\mathbf{x}_k^{\text{dep}}$. Under this augmented state, the next forecast error is stagewise independent, so the formulation in Section 3.3 continues to apply. Richer dependence models are certainly possible, for example by tracking multi-period error covariance or longer error histories. We do not pursue those extensions here because the goal of this section is not to optimize the residual model, but to test whether even a parsimonious dependence correction can materially change the control results. In that sense, the AR(1) specification serves as a transparent and economically interpretable stress test of the independence assumption.

Table 3.2: Main real-data comparison: dependence-aware DP versus LP benchmarks.

	2h	4h	8h
Dependence-aware DP vs. LP-Open	−1.8%	+1.7%	+8.7%***
Dependence-aware DP vs. LP-MPC	−3.3%	+0.2%	+5.0%
n (days)	182	182	182

Note: Relative profit differences are computed as $\overline{(\text{DP} - \text{LP})} / |\overline{\text{LP}}|$, where the overbar denotes the sample average over days and LP denotes the relevant benchmark in that row. Significance stars are based on two-sided paired t -tests of daily profit differences.

Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

Table 3.2 reports the main real-data comparison. Relative to the baseline DP in Table 3.1, the dependence-aware specification closes the gap to LP-MPC substantially, and its performance improves with storage duration. For 2-hour storage, the dependence-aware DP still slightly trails the deterministic benchmarks. For 4-hour storage, it slightly exceeds LP-open and is essentially tied with LP-MPC. For 8-hour storage, it outperforms LP-open and numerically exceeds LP-MPC, with a larger and statistically significant advantage over the open-loop benchmark. The real-data message is therefore not that stochastic control uni-

versally dominates deterministic reoptimization, but that DP becomes competitive once the exogenous-state model captures salient dependence in forecast errors, especially for longer-duration storage. In other words, the contribution of the real-data analysis is a boundary condition on when stochastic control is worth deploying, not a claim that it should displace deterministic reoptimization in every market or on every day.

The same real-data experiments also show that the dependence-aware DP remains computationally manageable. Table ?? reports solve times and active piecewise-affine segment counts for the real data runs with the shock discretization fixed at $Q_k = 50$ and the maximum representative-state budget fixed at 50. We report two segment diagnostics: the day-level average number of active affine segments across periods, and the day-level maximum across periods. For each metric, the table reports the mean across sampled runs. The purpose is to show that enriching the error process does not create an excessive computational burden in the real-data study.

Table 3.3: Computational footprint of the dependence-aware DP on CAISO real data.

Metric	2h	4h	8h
Solve time (s)	3.30	5.00	6.87
Active segments* (period avg.)	3.50	10.68	18.46
Active segments* (period max.)	11.25	21.54	27.17

Note: For each day, we compute the average and maximum number of active segments across periods; the table reports the mean of those daily statistics.

With both Q_k and the representative-state budget fixed at 50, solve times remain modest and segment counts stay well below any sign of combinatorial growth.

Tables 3.2 and ?? establish that dependence-aware DP can be both computationally practical and competitive in real data, especially for longer-duration assets, but they also show that LP-MPC remains a strong comparator and that the daily ranking varies sub-

stantially. That heterogeneity shifts the operational question from average dominance to regime selection. In particular, on what kinds of days should an operator expect stochastic look-ahead to earn its additional modeling effort? We turn to that *ex ante* question next.

3.4.4 When is stochastic look-ahead most attractive?

Because LP-MPC remains a strong practical benchmark, the more useful question is not whether DP wins on average, but when stochastic look-ahead is most likely to add value. To study this, we regress the daily profit difference,

$$\Delta_d := \text{Dependence-aware DP}_d - \text{LP-MPC}_d,$$

on *ex ante* variables observable at the start of the horizon. Specifically, for each battery duration we estimate linear models of the form

$$\Delta_d \sim \alpha + \beta^\top z_d,$$

where z_d is a vector of candidate *ex ante* features. We screened a broader set of *ex ante* variables and found that two were the most consistently informative: the deterministic benchmark value V_d^{open} and a model-implied uncertainty level $\bar{\sigma}_d$. The deterministic benchmark value V_d^{open} is the optimal objective value of LP-open solved on the point forecast available at the start of the day. It measures how much arbitrage value is already visible in the mean forecast path. The model-implied uncertainty level $\bar{\sigma}_d$ is the average, across hours of day, of the estimated residual standard deviations from the HW model using data available up to the focal day. Because it is estimated from an expanding sample, this variable moves slowly over time and can be interpreted as a regime-level proxy for forecast uncertainty. Figure 3.2 plots the relationship between Δ_d and these two variables.

Three patterns emerge. First, DP tends to perform better relative to LP-MPC when the deterministic benchmark value is lower. When the mean forecast already reveals strong arbitrage opportunities, frequent deterministic reoptimization is difficult to beat. When that deterministic signal is weaker, the value of stochastic look-ahead becomes more pronounced. Second, DP tends to perform better when the model-implied uncertainty level is higher.




Figure 3.2: *Ex ante* regime analysis with recomputed duration-specific deterministic benchmark values

This is consistent with the role of stochastic optimization. When the market is in a higher-uncertainty regime, explicitly accounting for future uncertainty has more room to improve decisions. Finally, the patterns are also stronger for longer-duration batteries, which is consistent with the idea that stochastic look-ahead matters more when the asset is more exposed to medium-term forecast risk.

To turn these descriptive patterns into an operational prescription, we also estimate a shallow CART switching rule using only the two *ex ante* variables above. For each battery duration, we fit a separate rolling regression tree on an expanding sample and re-estimate it each day using only prior observations for that same duration. The tree predicts the next-day profit difference Δ_d ; we choose dependence-aware DP whenever the predicted gap is positive and LP-MPC otherwise. To align this exercise with the full-year real-data comparison, we evaluate the rule on a 181-day rolling sample, using day 1 only to initialize the history and scoring days 2–182 out of sample. Among the simple candidate trees we considered, a fixed depth-2 specification performed best, with minimum leaf size 30 for the 2-hour and 4-hour batteries and 20 for the 8-hour battery. Table 3.4 reports the results.

The switching rule improves on both always-DP and always-LP-MPC in all three durations, while remaining about 6–8% below the infeasible oracle that chooses the better of DP and LP-MPC ex post.

Table 3.4: Rolling CART switching rule on real data.

	2h	4h	8h
Gain vs. always dependence-aware DP	+5.2%	+2.7%	+2.4%
Gain vs. always LP-MPC	+1.7%	+3.1%	+7.5%
Oracle gap	7.8%	7.2%	6.3%
Share of days choosing DP	40.3%	56.9%	66.3%
n (evaluation days)	181	181	181


The regime analysis identifies when DP is most attractive, but it does not yet explain which parts of the algorithm make those gains possible. In principle, the advantage in Table 3.2 could come from stochastic modeling alone or from the specific approximation architecture in Section 3.3. The next subsection isolates the impact of the design choices intended to preserve dispatch-relevant value.

3.4.5 *Where do the gains come from?*


In this subsection, we connect the regime findings back to the algorithmic design. Section 3.4.4 shows that DP is most useful when uncertainty matters and deterministic signal is weak. In exactly those regimes, coarse approximations are most likely to destroy the value of flexibility. We therefore test whether the gains of the proposed DP depend on preserving the right structure in each dimension of the problem, namely exact piecewise-affine treatment of the SoC dimension and forecast-feature-based reduction of the exogenous state.

The first ablation examines the treatment of the SoC dimension by comparing the pro-

posed exact piecewise-affine recursion with a naive uniform SoC grid. Figure 3.3a presents the profit of using a uniform SoC discretization relative to the proposed exact piecewise affine Bellman solve (exact-PWL). Exact-PWL attains the highest mean profit for every battery duration and every tested grid size. Coarse grids leave noticeable value on the table, with the profit gap most prominent for longer duration batteries. Although a fine enough grid can eventually approximate the benchmark well, it requires careful tuning and testing. In comparison, the exact-PWL construction removes that tuning and delivers the best observed profit directly.



(a) Uniform SoC grid vs. exact PWL



(b) Raw-state clustering vs. forecast-feature clustering

Figure 3.3: Profit impact of two approximation choices relative to the proposed DP.

The second ablation concerns exogenous-state aggregation. We compare the proposed forecast-feature clustering with a naive alternative that clusters directly on the raw HW state vector. Figure 3.3b shows that the proposed method outperforms the raw-state benchmark at every tested state budget, with a substantial profit advantage for more generous state budget. This is because distances in the raw HW state space do not reliably reflect distances in control value. Specifically, two states that are far apart in Euclidean coordinates can imply very similar near-term forecast paths and therefore very similar dispatch incentives. Under a limited state budget, separating such states into different clusters uses capacity inefficiently. This is also consistent with the regime analysis in Section 3.4.4. When uncertainty is

high, the relevant notion of state similarity is whether two states imply similar current and near-term forecast paths, not whether they are nearby in the raw latent coordinates. By clustering on forecast-feature representations, the proposed method groups states according to the information that matters most for control and preserves more of the economically relevant structure.

The ablation patterns show that the benefit of exact piecewise-affine treatment is largest for longer-duration batteries, which is also where dependence-aware DP becomes most competitive in Table 3.2. Meanwhile, the benefit of forecast-feature clustering is largest when representative-state budgets are tight, precisely the setting in which preserving the economically relevant part of uncertainty matters most. Taken together, these results suggest that the gains of DP in favorable regimes come not from stochastic modeling alone, but from preserving the parts of the problem that matter most for dispatch value.

3.5 Discussion, Managerial Implications, and Conclusion

We develop a finite-horizon stochastic control framework for price-taking storage that separates the problem into two parts. We embed the exogenous price process through the latent state of a forecasting model and represent that exogenous evolution on a finite graph. Conditional on that graph, we solve the SoC recursion exactly through a continuous piecewise-affine Bellman recursion on an endogenous event grid. This gives a plug-and-play architecture for any forecasting models with a Markov latent-state representation, while making clear that approximation enters only through the exogenous reduction.

We then show that these modeling choices matter economically, not just computationally. The ablation results show that solving the SoC dimension in exact piecewise-affine form improves on uniform SoC grids, especially for longer-duration assets, and that reducing the exogenous state in forecast-feature space improves on clustering directly on the raw latent state. The value of the framework therefore comes from preserving the parts of the problem that are most relevant for dispatch, rather than from stochastic modeling alone.

The central operational question is then when such a stochastic framework is worth using. Our results show that the answer depends on both the benchmark and the error model. Relative to an open-loop deterministic policy, stochastic recourse creates clear

value in matched environments. Relative to LP-MPC, however, the comparison is more demanding. The baseline DP loses much of its advantage in real data when forecast errors are treated as independent. Once dependence in forecast errors is modeled, the stochastic policy becomes much more competitive, particularly for longer-duration batteries, where medium-term uncertainty matters more. The regime analysis tells the same story. DP is most attractive when the deterministic forecast signal is weaker and uncertainty is higher, and a simple switching rule can improve on always using either DP or LP-MPC.

Our empirical claims should also be interpreted with the scope of the study in mind. The out-of-sample evidence comes from one market, one node, and one forecasting family with a parsimonious dependence correction. We therefore do not claim universal superiority of stochastic control across electricity markets. Rather, the evidence identifies a set of market conditions under which stochastic look-ahead is more or less likely to matter, and it shows that those conditions can be studied in a computationally tractable way within an industry-relevant dispatch model.

For operators, the results support selective adoption rather than universal replacement. Receding-horizon deterministic control remains a strong default when point forecasts already reveal clear arbitrage opportunities or when short-horizon uncertainty is limited. By contrast, stochastic look-ahead is most attractive for longer-duration batteries and on days when the deterministic value signal is weak and forecast uncertainty is high. For market designers and regulators, the results suggest that the realized value of storage depends not only on installed capacity but also on the information environment and dispatch tools through which that capacity is operated. In markets with recurring low or negative prices, forecast quality and the ability to model forecast-error dependence can materially affect how much flexibility storage actually delivers. The evidence points to a practical implication: better information and longer-duration assets can improve the operational absorption of renewable variability only if dispatch methods are able to use that information effectively.

Taken together, these results suggest a more disciplined role for stochastic control in storage operations. We do not find that DP should replace deterministic reoptimization as a default. Rather, we show where its value comes from and when it is most likely to pay off: solve the storage dimension exactly, represent the exogenous process in a control-relevant

way, and use stochastic look-ahead in regimes where uncertainty materially changes dispatch opportunities.

Chapter 4

**EXPLAIN WORKER SCHEDULE QUALITY DISPARITIES:
EVIDENCE FROM AN AMERICAN RESTAURANT CHAIN**

This chapter is based on joint work with Qiuping Yu.

4.1 Introduction

Advances in large-scale data collection and AI-enabled demand forecasting have made just-in-time (JIT) scheduling increasingly prevalent in the service industry. This approach involves managers initially scheduling a minimal number of employees and then dynamically adjusting staffing levels based on real-time demand and attendance fluctuations [Sampson and Froehle, 2006, Lambert, 2008]. While JIT scheduling is effective in reducing labor cost for the companies as they can schedule just enough (not too many) employees, it results in unpredictable and often insufficient hours for employees. National estimates indicate that roughly 17% of U.S. workers experience such unpredictable schedules [Golden, 2015]. According to the national survey data from the Shift Project [Schneider and Harknett, 2019b], such precarious scheduling practice is particularly prevalent in the food services and retail sectors: 90% of the workers have unstable schedules, 75% express a desire for more predictable schedules, and 33% prefer to be scheduled for more hours.

Precarious scheduling has significant economic and social consequences for employees. Studies have demonstrated that such scheduling practices can severely impact mental health by increasing work-related stress, exacerbating work-life conflict, and diminishing a sense of control over one's life [National Health Interview Survey, 2023]. It also often leads to household hardships, including economic insecurity due to unpredictable income fluctuations, insufficient scheduled hours, and the barriers it creates when workers attempt to secure a second source of income [Schneider and Harknett, 2021]. Perhaps even more profoundly, these precarious schedules pose significant challenges in arranging childcare and can lead to developmental disadvantages for the children of impacted workers [Harknett et al., 2022,

Schneider and Harknett, 2019a].

Given the prevalence of the precarious scheduling practice and its crucial impact on worker well-being, we investigate whether it affects all workers equally. Previous research on workplace fairness primarily focused on hiring, wages, benefits, and promotions (see e.g., Cancio et al. 1996, Blau and Kahn 2000, 2017, Goldin 2014, Mandel and Semyonov 2016, Jones and Schmitt 2016, Snipp and Cheung 2016, Kristal et al. 2018). However, disparities in scheduling quality have received little attention, despite their clear significance — particularly for hourly workers who make up 55.6% of the U.S. workforce and whose livelihoods and well-being are directly shaped by when and how much they are scheduled to work. The limited research that explores this issue (see e.g., Finnigan and Hunter 2018, Lambert et al. 2014, Storer et al. 2020) relies on survey data, which may be affected by sample selection and self-reporting biases. They also frequently rely on relatively small subgroup sample sizes, which may limit the robustness of their findings, and lack the data granularity needed to account for key confounding factors or to clearly identify the mechanisms underlying observed disparities. To address this gap, our study investigates racial and gender disparities in scheduling practices within the service sector using *real-world* shift data. Our dataset includes *all* workers at the covered restaurant locations, eliminating sample selection bias. The granularity of our data and our carefully crafted identification strategy allow for a more nuanced analysis of the mechanisms behind racial and gender disparities than has been possible in previous work. Specifically, with our dataset, we explore two key questions: (1) Whether there are, and if so, what are the gender and racial disparities in scheduling quality across various dimensions? (2) What are the mechanisms that drive these disparities?

We conduct our research in the restaurant industry, a sector well-known for its prevalent use of precarious scheduling practices that significantly impact a large and diverse workforce [Schneider and Harknett, 2019b]. Representing approximately 10% of the U.S. workforce, the industry has a higher proportion of female and minority workers compared to the national average [National Restaurant Association, 2024], making it an ideal setting to investigate scheduling disparities. Our dataset comes from a national full-service casual dining restaurant chain in the U.S., covering three years (2014–2016) and spanning 25 loca-

tions in the Pacific Northwest. It includes nearly half a million initially offered and actually worked shifts for over 2,400 part-time customer-facing workers, as well as more than 1.7 million transactions and nearly 300,000 customer reviews. For our main analysis, we focus on a detailed 9-month subset of this data, leveraging its granularity to account for worker skills and preferences, while also replicating our analysis on the full 3-year dataset to ensure robustness.

At our partner restaurant chain, managers offer workers shifts a week before the service day based on demand forecasts and workers’ past preferences. Shifts with this seven-day advance notice are called *regular* shifts. Shortly prior to the day of service (mostly two days in our data), managers may add shifts based on worker attendance information, resulting in what we term *short-notice additions (SNA)*.¹ On the day of service, managers can observe actual demand and worker attendance, and make further adjustments in real-time by adding hours for workers who are already onsite—referred to as *real-time additions (RTA)*. We refer to these short-notice and real-time additions collectively as *last-minute* additions. Workers are compensated at the same rate for the hours they work regardless of the advance notice given. Guided by these practices, our investigation focuses on two key dimensions of scheduling quality at the worker-week level: (1) sufficiency, which assesses the number of hours they receive; and (2) predictability, which examines the extent of advance notice to workers about their shifts, particularly in terms of the frequencies of their last-minute additions.

To compare scheduling quality across worker groups, we begin by accounting for a set of factors commonly identified in the workplace inequality literature as drivers of disparities in job outcomes. Central to our framework is the concept of human capital—the economic value workers bring through their experience, skills, and abilities [Blau and Kahn, 2006b, Cook et al., 2021]. We capture multiple dimensions of human capital, including tenure at the restaurant, role type (e.g., server vs. bartender), absenteeism, and performance in sales, service speed, and customer satisfaction. Our ability and absenteeism measures are made possible by highly detailed field data, enabling a level of granularity rarely attainable

¹Here we follow the terminologies in Kamalahmadi et al. [2021].

through surveys or less comprehensive sources. Specifically, we assess absenteeism based on the long-term frequency of self-initiated, last-minute shift cancellations—a direct proxy for reliability in fulfilling scheduled commitments, which can shape managerial scheduling decisions. Likewise, our performance metrics offer a multidimensional view of each worker’s on-the-job effectiveness in areas most valued by employers in our context, providing a more accurate and objective picture of worker ability than what is available in prior research. In addition to human capital, we account for operational factors such as work location and schedule seasonality. When analyzing disparities in scheduling predictability, we further control for the number of hours initially offered in the focal week, as well as each worker’s long-term average weekly hours. These controls help rule out the alternative explanation that workers received more last-minute shift additions simply because they were offered fewer hours in advance. Our analysis compares scheduling outcomes across four demographic groups—White males (reference group), White females, minority males, and minority females—to assess both gender and racial disparities.

Our findings reveal significant differences in scheduling outcomes across gender and racial groups. Human capital factors collectively explain a substantial portion of these disparities. Specifically, accounting for tenure, role, absenteeism, and skill metrics reduces gaps in scheduling sufficiency by 13–45% and in scheduling predictability by up to 15%. However, even after adjusting for these factors, the remaining gaps continue to be sizable. In particular, we show that after accounting for differences in workers’ human capital factors, women and minority men were offered fewer hours in advance but were more likely to receive last-minute additions. Compared to White men, White women are offered 10.7% fewer hours but experience a 7.2% higher short-notice additions (SNA) ratio; minority women receive 6.9% fewer hours, with a 16.9% higher SNA ratio and a 7.0% increase in real-time additions (RTA) ratio; and minority men are scheduled 3.3% fewer hours alongside a 7.4% higher SNA ratio.

Given the sizable disparities that cannot be attributed to human capital factors, we further investigate potential mechanisms related to workers’ preferences and availability that might drive these remaining differences. We first address the possibility that women and minority workers might prefer last-minute additions over regular shifts due to the flexibility

they may bring. Using granular transaction-level data, we show that, when comparing *the same shifts* with different levels of advance notice, there is a uniform preference (measured by workers' sales performance) among workers for regular shifts over last-minute ones regardless of workers' race or gender. Further, last-minute additions generally occur during hours with lower earning potential compared to regular shifts—approximately 40% lower for real-time additions and 2-4% lower for short-notice additions, making it even less likely that workers would prefer last-minute additions. In addition to examining workers' preference for regular shifts, we investigate whether lower availability contributes to the lower offered sufficiency observed among women and minority workers. By comparing the hours initially offered with the hours actually worked, we find that reduced availability explains part of the disparity for White women. However, it does not account for the gaps observed among minority workers, both male and female.

In summary, significant gender and racial disparities in scheduling persist even after accounting for human capital. Women and minority men receive fewer hours and more last-minute shifts than White men. These gaps are not explained by the classic explanation of worker preferences or availability, as workers universally prefer regular shifts and only White women's lower availability may explain part of the shortfall in our data.

To the best of our knowledge, we are the first to study gender and racial disparities in an often-overlooked aspect of job quality—workers' schedules—using real-world shift data. Our analysis not only documents these disparities but also uncovers the mechanisms behind them through a carefully designed identification strategy, advancing the literature in three key ways. First, we leverage granular field data to objectively capture multiple dimensions of human capital—traits that are rarely measured with precision or examined in survey-based research. We assess worker performance across sales ability, service speed, and customer satisfaction, and consider absenteeism as a proxy for reliability. This level of detail allows us to analyze dimensions of worker quality that are typically self-reported or entirely overlooked. Second, we examine mechanisms that are traditionally difficult to quantify, including worker preferences. By comparing sales performance and earning potential across shift types, we infer worker preferences between last-minute and regular shifts. Third, we provide a comprehensive assessment of disparities in both schedule sufficiency and predictability. This

enables us to estimate a lower bound on worker availability and to distinguish between managerial under-scheduling and actual constraints on worker engagement—offering insights that would be missed if each dimension of schedule quality were studied in isolation.

4.2 *Literature Review*

Our research is related to literature on workplace gender and racial disparities and workforce scheduling.

Workplace Gender and Racial Disparities. Literature on workplace disparities consistently documents significant gender and racial gaps in hiring, wages, benefits, and promotions (see e.g., Blau and Kahn 2000, Goldin 2014, Blau and Kahn 2017, Snipp and Cheung 2016, Mandel and Semyonov 2016, Kristal et al. 2018). These disparities stem from several factors, including human capital differences [Altonji and Blank, 1999, Blau and Kahn, 2006a], variations in worker preferences and availability [Marianne, 2011], occupational segregation [Goldin and Katz, 2002, Jones and Schmitt, 2016], unequal access to networks and mentorship [Pedulla and Pager, 2019], as well as differences in social dynamics, negotiation powers [Bowles et al., 2007, Babcock and Laschever, 2009, Hernandez et al., 2019] and biases [Goldin and Rouse, 2000, Bertrand and Mullainathan, 2004].

Given that our study focuses on workers within the same restaurant chain performing the same role, the most relevant factors are human capital, preferences, availability, negotiation powers and biases. Human capital factors typically include education, experience, and intrinsic skills. Increases in human capital among women and minority workers—driven by higher educational attainment and greater labor market participation—have played a key role in reducing both gender and racial pay gaps [Altonji and Blank, 1999, Blau and Kahn, 2006a].

Regarding preferences and availability, a major factor is the fewer hours worked by women and penalties associated with preference for job flexibility. Women often face wage reductions when they seek flexible work arrangements, as jobs requiring long and inflexible hours typically offer higher pay [Goldin, 2014]. Even in gig economies, where workers can choose their hours, women consistently work fewer hours than men [Cook et al., 2021]. Another critical factor is the motherhood penalty, where time off for childbirth and childcare

and the resulted weak continuity in labor force participation leads to long-term wage losses [Cuddy et al., 2004, Blau and Kahn, 2017]. Compared to gender, there is limited evidence suggesting that racial disparities in the workplace stem from differences in availability or preferences.

On the front of social dynamics and negotiation powers, studies consistently find that women tend to be more risk-averse, less competitive, exhibit stronger altruism, and are less likely to initiate negotiations—all of which can exacerbate the gender gap [Marianne, 2011, Eagly and Wood, 1991, Eagly, 2013, Babcock and Laschever, 2009, Babcock et al., 2017]. Meanwhile, racial differences in social dynamics and negotiation can be influenced by cultural norms. For example, Komarraju et al. [2008] highlights that individuals from collectivist cultures—often overlapping with minority groups in Western contexts—tend to prioritize group harmony over assertive self-advocacy, making them less inclined to negotiate.

Additionally, conscious and unconscious managerial biases negatively affect the opportunities available to female and minority workers, both by limiting initial career advancement [Goldin and Rouse, 2000, Bertrand and Mullainathan, 2004] and by increasing resistance to their self-advocacy efforts [Bowles et al., 2007, Bowles and McGinn, 2008, Hernandez et al., 2019].

Scheduling Disparities. While significant strides have been made in understanding gender and racial disparities in hiring, wages, benefits, and promotions, there remains a critical gap in research on the temporal quality of workers’ schedules—an area particularly relevant for hourly workers, who make up 55.6% of the U.S. workforce [US Bureau of Labor Statistics, 2023].

The limited existing research on workforce scheduling disparities has identified some significant cross-group disparities. For example, Lambert et al. [2014] analyzed data from the National Longitudinal Survey of Youth 1997 and found that short-notice schedule changes are significantly more common among part-time workers compared to full-time workers, as well as among working parents. Similarly, based on the 2004-2008 Survey of Income and Program Participation, Finnigan and Hunter [2018] found that the White-Black disparity in work hour variability increased during the late 2000s recession in White-dominated oc-

occupations but decreased in minority-dominated occupations. While these studies provide valuable insights, they also have notable limitations. Both studies face challenges with relatively small subgroup sample sizes after controlling for factors like location, industry, and occupation. Additionally, neither directly investigates the mechanisms behind the observed disparities or adequately adjusts for potential confounders. Furthermore, each study focuses on a single dimension of schedule quality: Finnigan and Hunter [2018] examines variability in work hours, while Lambert et al. [2014] primarily addresses short-notice scheduling changes.

The study most closely related to ours is Storer et al. [2020], which makes important strides by examining a wide range of precarious scheduling practices—including on-call shifts, last-minute cancellations, involuntary part-time work, difficulties obtaining time off, and "clopening" shifts (where workers close one shift and open the next). The authors show that non-White workers, especially non-White women, are disproportionately affected, largely due to firm-level segregation and racial discordance between workers and managers. While this study offers a more comprehensive view of precarious scheduling than prior research, it relies on survey data collected via Facebook and Instagram with a response rate of just 1.1%, and does not directly account for factors such as workers' intrinsic abilities, availability, preferences, or absenteeism.

Our study addresses these gaps by leveraging actual shift data rather than survey responses, thereby addressing the problems of sample selection and subjective reporting biases. We include all workers at the covered restaurant locations, capturing a complete and objective picture of scheduling practices. Beyond standard human capital factors like tenure and role, we leverage granular transaction data to capture intrinsic worker abilities—such as sales performance, service speed, and customer ratings—that are critical to the company but typically unobservable in survey or less detailed field data. Furthermore, we introduce an innovative approach to disentangle the impacts of worker availability, preferences, and absenteeism—factors typically challenging to analyze due to their unobservable nature. This comprehensive approach allows us to offer deeper and complementary insights into the mechanisms driving the disparities.

Impact of Precarious Schedules. A growing body of empirical and sociological re-

search highlights the prevalence of precarious scheduling in real-world operations and its harmful effects on workers and their families [Golden, 2015, Schneider and Harknett, 2019b]. For example, Henly and Lambert [2014] and Golden [2015] find that unpredictable and irregular shifts intensify work-family conflict and heighten stress. Lambert et al. [2019] and Schneider and Harknett [2021] further link precarious scheduling to economic insecurity, financial hardship, and institutional distrust. Harknett et al. [2022] examines the effects on childcare arrangements, showing that volatile schedules force reliance on multiple informal care providers, undermining child development and parental well-being. In addition to harming workers, just-in-time scheduling may also have unintended negative consequences for firms. A growing number of studies find that such practices reduce employee productivity and ultimately erode business profitability, contradicting the cost-saving rationale that motivates them [Kamalahmadi et al., 2021, Kesavan et al., 2022]. Moreover, schedule volatility contributes to higher voluntary turnover among both professional workers, such as nurses [Bergman et al., 2023], and workers in the food and retail sectors [Choper et al., 2022], resulting in added costs related to hiring, training, and operational disruption.

Operational Flexibility in Services. Our study is broadly related to the literature on operational flexibility in service settings. Prior research in this area has primarily explored how firms achieve flexibility through cross-training or adaptable work schedules. When demand is highly heterogeneous, cross-training enhances throughput by enabling workers to perform multiple tasks [Gans et al., 2003, Iravani et al., 2005]. In contrast, when demand is more uniform but volatile over time, firms often rely on flexible scheduling to reallocate capacity across periods. Much of this literature uses analytical staffing models that incorporate dynamic labor adjustments through mechanisms such as temporary staffing, overtime, or on-call shifts [Hur et al., 2004, Bhandari et al., 2008, Mehrotra et al., 2010, Easton, 2014, Gans et al., 2015, Slauch et al., 2018]. Kesavan et al. [2014] are the first to empirically study how volume flexibility created through part-time and temporary workers impacts the financial performance of big box retailers. Our study focuses on a restaurant context in which all employees are part-time and scheduling flexibility is implemented via just-in-time scheduling practices. Rather than optimizing the scheduling policy itself, we investigate whether this approach to flexibility is applied equitably across workers in practice.

4.3 Data

We describe our study context and dataset, and define the the key variables in our main model in this section.

4.3.1 Study Context

We set our study within the restaurant industry, a sector that generates approximately 1 trillion dollars in annual sales [US Census Bureau, 2024]. The industry employs 10% of the US workforce, with notable demographic characteristics: about 45% of restaurant servers are minorities, and 70% are female—both figures exceeding the national averages [National Restaurant Association, 2024]. These demographics underscore the relevance of our discussion on racial and gender scheduling disparity.

Our dataset comes from a large American casual-dining restaurant chain with over 500 locations nationwide, focusing specifically on customer-facing roles—servers and bartenders. On average, each location requires approximately ten customer-facing employees working simultaneously during operating hours, with specific staffing levels varying by the hour of day and the day of week. An algorithm guides managers on demand forecasts, specifying *how many* labor hours are needed for each hour of any given day, while managers retain discretion over *who* to schedule. The scheduling process begins a week before the service day, when managers offer workers shifts based on demand forecasts and workers’ past preferences. These shifts, given with seven days’ notice, are called *regular* shifts.

As the service day approaches, managers may add shifts based on updated demand and worker availability. These additional shifts, known as last-minute additions, occur in two phases. *Short-notice* additions are assigned to servers shortly before the day of service (mostly two days in advance). These shifts typically arise when previously assigned workers request leave, requiring managers to fill gaps. *Real-time* additions happen on the service day itself, when managers extend onsite workers’ shifts due to unexpected worker absences or surges in demand. Workers receive the same pay rate regardless of when their shifts are assigned. Although they can decline additional hours, they may feel pressured to accept them—even when these extra hours conflict with personal plans or commitments—due

to pressures such as fear of retaliation, limited alternative employment opportunities, or a workplace culture in which unpredictable scheduling is normalized [Henly et al., 2006, Golden, 2015, Schneider and Harknett, 2017]. Such pressure can be particularly strong for real-time additions, as workers are asked to extend their hours while already on the job. At our restaurants, managers do not proactively seek to cancel worker shifts. However, as alluded to earlier, workers might request absences with a short advance notice, ask to leave early while already on the job, or fail to show up entirely.

4.3.2 *Data Description*

The main dataset includes data from 25 locations of the national restaurant chain in the Pacific Northwest. Each location in the dataset provides full dining, bar, and takeout services. Operating hours vary by location, with openings between 10 and 11 a.m. and closings between 10 p.m. and midnight, resulting in typical operating windows of 12 to 14 hours. The dataset contains detailed information across the following dimensions.

Worker Data. The dataset includes information on race, gender, roles, employment dates, and employment status (part-time or full-time) for 2,697 customer-facing workers. We focus on part-time workers, who constitute approximately 98% of the workforce. Additionally, we exclude workers with missing tenure length information, resulting in a final sample of 2,485 unique workers for analysis.

Shift Data. Our shift data includes start and end times for all shifts of customer-facing workers from January 2014 to September 2016. For each worker, both offered and worked shifts are recorded. Offered shifts are those assigned by managers during the planning phase a week before the service day, while worked shifts reflect the hours workers ultimately worked, capturing all last-minute adjustments. The raw dataset comprises 499,953 unique offered shifts, totaling 2.18 million worker hours, and 486,877 unique worked shifts, totaling 2.37 million worker hours. The interquartile range (IQR) for the length of offered shifts is 3.25 to 5.25 hours, with a median of 4 hours. For worked shifts, the IQR ranges from 3.5 to 5.9 hours, with a median of 4.7 hours. We use the difference between workers' offered shifts and worked shifts to infer short-notice and real-time changes, as will be detailed in Section

4.3.3.

Point-Of-Sales (POS) Data. The POS data encompasses detailed information on 1,717,191 transactions spanning a 9-month period from January to September 2016, including sales amount, party size, meal duration, timing, location, and the handling worker, among other variables. To mitigate the impact of outliers, we exclude the top and bottom 1% of transactions based on sales amount and meal duration and filter for transactions handled by workers of interest. The final POS dataset includes 1,529,770 transactions.

Customer Survey Data. The survey data includes feedback collected from customers at the end of their meals, also spanning the 9-month period from January to September 2016. The survey covers several aspects, including food quality, service speed, restaurant cleanliness, and the handling worker’s attentiveness and enthusiasm. To assess the handling worker’s performance, we use responses to the question, “How likely are you to recommend this service to others?” rated on a scale from 0 to 10. The dataset comprises 288,048 completed surveys, reflecting a response rate of 17%.

4.3.3 Variable Definitions

We next define the dependent, independent variables of interest, and control variables in our regression analysis.

Dependent Variables. We measure workers’ schedule quality in the dimension of sufficiency and predictability at worker-store-week level. We first define sufficiency as $OfferedWeeklyHrs_{ijt}$, the offered number of hours for worker i at store j during week t .

We next define schedule predictability. As previously noted, all additional hours are assigned by managers, whereas cancellations are initiated by workers. Therefore, in measuring predictability, we focus on last-minute additions. A worked shift is classified as a short-notice addition (SNA) if it does not appear in the originally offered shifts. A real-time addition (RTA) occurs when a worker’s worked shift extends beyond the originally offered duration. For instance, if a worker’s offered shift was from 12 to 4 p.m. but worked until 5 p.m., we consider the period between 4 p.m. and 5 p.m. a real-time addition. We define

$RTARatio_{ijt}$ as the fraction of offered shifts that had real-time addition for worker i at store j in week t , while $SNARatio_{ijt}$ denotes the fraction of worked shifts that were short-notice addition for worker i at store j in week t .²

Independent Variables. Our primary independent variables are workers’ race and gender. We classify the 2,485 unique workers in our dataset into four groups: White females, White males, minority (non-White) females, minority (non-White) males. We use three dummy variables to indicate the classification of each worker: $WhiteFemale_i$, $MinorityFemale_i$, and $MinorityMale_i$, which equals 1 if worker i falls within the corresponding group, and 0 otherwise. If all variables equal 0, the worker is a White male worker, who are treated as the baseline group.

The determination of such demographic classification is as follows. The raw dataset contains each employee’s first and last name, which enables the inference of their demographic background [Bertrand and Mullainathan, 2004, Cui et al., 2020]. Two algorithms are commonly used to infer race and gender from names in previous studies (see, e.g., Barrios et al. 2020, Munz et al. 2020): NamSor³ and NamePrism⁴. NamSor’s machine-learning algorithm is trained on large name datasets from across the world, such as the names of Olympic athletes, and is shown to predict an individual’s gender with over 95% accuracy and race with 75% accuracy [Munz et al., 2020]. NamePrism’s tool is trained on contact lists from an email company, which contains 74 million labeled names covering 118 countries. It assigns each name with a probability for different ethnic groups [Ye et al., 2017].

In our main analysis, we use the gender assignment based on NamSor, while the race assignment was based on both algorithms. When the race predictions are inconsistent, the names were manually searched in Google Images, after which the most likely race was inferred based on human judgment. More detailed explanation of this process can be found in Appendix C.1.2. Using this approach, the workforce comprises 50% White females, 30%

²The number of worked shifts is used as the denominator in calculating $SNARatio_{ijt}$ to prevent division by zero and to ensure the ratio remains between 0 and 1. In addition, qualitatively and quantitatively consistent results are observed when using the ratio of the duration of last-minute changes as the outcome variable.

³<https://namsor.app/>

⁴<https://www.name-prism.com/>

White males, 12% minority (non-White) females, and 8% minority males. Overall, females constitute 62% of the workforce. In Appendix C.1.2, we replicate the main analyses using race classifications derived exclusively from NamSor and NamPrism and demonstrate that the primary conclusions remain consistent across these alternative methods.

Control Variables. We control for several variables that could influence workers’ scheduling outcomes, drawing on key factors identified in the literature as contributors to workplace gender and racial disparities. Specifically, we first account for workers’ role and tenure [Lambert et al., 2014]. $Role_i$ is a categorical variable capturing the focal worker’s role mix. Workers can serve in three roles: server, cocktail server, and bartender, and it is possible for workers to be assigned different roles for different shifts. The variable $Role_i$ represents the combination of roles each worker performs, including “server only”, “server or cocktail server”, “server or bartender”, “cocktail server or bartender”, and “all three roles”. In our dataset, the majority of workers (65%) work exclusively as servers. Approximately 19% of workers serve as either server or cocktail server, about 2% serve as server/cocktail server and bartender, and 13% serve in all three roles.⁵ Experience level is measured as the number of days from the worker’s initial hire date at the store to the worker’s first shift in the focal week. We classify tenure days into three categories - short, mid, long - using the bottom and top 25% thresholds for sufficient variation. Two dummy variables, $MidTenure_{ijt}$ and $LongTenure_{ijt}$, are defined, with short tenure as the baseline category. In our data, the thresholds are 235 days (short to mid) and 1,739 days (mid to long).

Beyond tenure and worker role composition, we also control for workers’ absenteeism. Absenteeism is a significant problem among frontline workers [Muchinsky, 1977, Johansson and Palme, 2002, Junor et al., 2009]. Thus, managers may prioritize more reliable workers when assigning shifts. We use the frequency of last-minute cancellation as a measure of worker absenteeism. Last-minute cancellations are classified into two categories: full cancellations, where an offered shift is absent from the worked shifts, and partial cancellations, where a shift ends earlier than the initially offered time. We define two absen-

⁵To rule out confounding from role-specific shift-assignment rules, we also run our analysis on the subsample of workers serving exclusively as servers. The results—presented in Appendix C.1.1 are consistent with our main findings.

teism metrics, $PartialAbsenteesim_{ijt}$ and $FullAbsenteesim_{ijt}$, representing cumulative averages of weekly partial and full cancellation ratios, respectively, for worker i at store j . Specifically, $PartialAbsenteesim_{ijt}$ is calculated by first determining, for each worker-store-week combination, the fraction of initially offered shifts partially canceled, and then averaging these fractions cumulatively across all weeks prior to the focal week t . Similarly, $FullAbsenteesim_{ijt}$ is computed using the fraction of fully canceled shifts.

Furthermore, we control for time fixed effects using dummy variables $Year_t$, $Week_t$ and store fixed effects using dummy variables $Store_j$. The former accounts for any trends over time, while the latter captures unobserved time-invariant heterogeneity across stores.

When analyzing predictability measures as the outcome variable, we further control for the workers' offered hours for the week, $OfferedWeeklyHrs_{ijt}$, as the baseline schedule affects workers' availability and thus the likelihood of additions. Additionally, we control for the workers' long-term average weekly hours at the store, $OfferedLTWeeklyHours_{ij}$, since workers with consistently longer weekly hours may differ from those with fewer hours, potentially leading to different managerial treatment.

Finally, workers' intrinsic ability levels may influence how managers allocate shifts [Blau and Kahn, 2006b]. Leveraging the granularity of our transaction data, we assess three most important dimensions of worker abilities that restaurant managers should value and matter the most to companies. The three dimensions are: sales (a worker's selling effectiveness), service speed (the rate at which a worker turns tables), and customer ratings (customer-perceived service quality), denoted as $SalesAbility_{ijt}$, $SpeedAbility_{ijt}$, and $ScoreAbility_{ijt}$, respectively. We estimate workers' intrinsic abilities using worker fixed-effect models, following the approach in Tan and Netessine [2019], Kamalahmadi et al. [2021], and Kamalahmadi et al. [2023]. To account for workers' learning over time, we divide all transactions by month and conduct separate estimations for each month and each ability dimension. Specifically, sales ability is quantified as the average sales per customer generated by a worker in a given month, adjusted for store and time fixed effects (hour of day, day of week, week of year). Service speed is measured as the inverse of the average meal duration of transactions a worker handled, adjusted for party size, number of items sold, check size, number of other transactions the worker handled at the same time, along with store and time fixed

effects. Lastly, service quality is measured as the average customer recommendation score a worker received, also adjusted for store and time-fixed effects. It is important to note that customers are assigned to servers in a round robin fashion. As such, customers served by different servers should be statistically similar and thus customer characteristics will not bias the three measurements above for workers' ability. Detailed explanations are provided in Appendix C.2. Higher values of $SalesAbility_{ijt}$, $SpeedAbility_{ijt}$, and $ScoreAbility_{ijt}$ indicate higher ability levels: generating greater sales, turning tables faster for the same sales amount, and achieving higher customer ratings, respectively. Since the transaction data is only available for the 9-month period in 2016, we can only estimate worker intrinsic ability levels for this 9-month period.

Summary Statistics. We focus our main analysis on the 9-month sample, as it allows us to control for the key human capital factors that characterizes servers' abilities in sales, speed, and customer ratings. This is essential to rule out ability-based disparities. Table 4.1 presents the definition and the mean of the main variables mentioned above for the full sample and the 9-month samples. The 9-month sample closely mirrors the full sample across all variables. In Appendix C.1.3, we replicate the main analyses using the full 3-year sample without accounting for intrinsic abilities, which yield consistent directional conclusions.

4.3.4 Model

Our main model to estimate the inter-group disparities in schedule quality captures the classic human capital factors using the control variables presented in Section 4.3:

$$Y_{ijt} = \beta_0 + \beta_1 WhiteFemale_i + \beta_2 MinorityMale_i + \beta_3 MinorityFemale_i + \beta_4 \mathbf{X}_{ijt} + \varepsilon_{ijt}. \quad (4.1)$$

For sufficiency, the dependent variable $Y_{ijt} = \log(OfferedWeeklyHrs_{ijt} + 1)$. For predictability, the dependent variable is the two predictability measures, $RTARatio_{ijt}$ and $SNARatio_{ijt}$. The vector of control variables, $\mathbf{X}_{ijt}$, is given as follows when estimating

Table 4.1: Summary Statistics of the Variables at Worker-Store-Week Level.

Variable	Definition	3-Year Sample	9-Month Sample
<i>WhiteFemale_i</i>	Dummy variable: 1 if worker is a white female, 0 otherwise	0.50	0.49
<i>MinorityMale_i</i>	Dummy variable: 1 if worker is a minority male, 0 otherwise	0.08	0.07
<i>MinorityFemale_i</i>	Dummy variable: 1 if worker is a minority female, 0 otherwise	0.12	0.13
<i>OfferedWeeklyHrs_{ijt}</i>	Offered weekly hours in the focal week	17.77	16.84
<i>WorkedWeeklyHrs_{ijt}</i>	Worked weekly hours in the focal week	19.29	18.32
<i>RTARatio_{ijt}</i>	Fraction of offered shifts with real-time additions (RTA)	0.39	0.38
<i>SNARatio_{ijt}</i>	Fraction of worked shifts resulting from short-notice additions (SNA)	0.10	0.12
<i>MidTenure_{ijt}</i>	Dummy variable: 1 if worker tenure is of short category in week t at store j , 0 otherwise	0.50	0.50
<i>LongTenure_{ijt}</i>	Dummy variable: 1 if worker tenure is long category in week t at store j , 0 otherwise	0.25	0.25
<i>OfferedLTWeeklyHours_i</i>	Long-term average weekly hours worker i receives at store j	18.11	17.19
<i>PartialAbsenteesim_{ijt}</i>	Rolling average of worker i 's weekly RTC ratio at store j by week t	0.12	0.13
<i>FullAbsenteesim_{ijt}</i>	Rolling average of worker i 's weekly SNC ratio at store j by week t	0.11	0.13
<i>SalesAbility_{ijt}</i>	Intrinsic sales ability	-	0.00
<i>SpeedAbility_{ijt}</i>	Intrinsic speed ability	-	-0.00
<i>ScoreAbility_{ijt}</i>	Intrinsic score ability	-	-0.01

inter-group disparities in sufficiency:

$$\mathbf{X}_{ijt} = (\textit{Role}_i, \textit{MidTenure}_{ijt}, \textit{LongTenure}_{ijt}, \textit{SalesAbility}_{ijt}, \textit{SpeedAbility}_{ijt}, \\ \textit{ScoreAbility}_{ijt}, \textit{PartialAbsenteesim}_{ijt}, \textit{FullAbsenteesim}_{ijt}, \textit{Store}_j, \textit{Week}_t). \quad (4.2)$$

When estimating the predictability disparities, as mentioned in Section 4.3.3, we extended the vector with two additional control variables: *OfferedWeeklyHrs_{ijt}* and *OfferedLTWeeklyHours_{ij}*.

4.4 Schedule Disparities - The Role of Human Capital

In this section we present the main results of gender and racial disparities with and without the human capital controls for each dimension of the scheduling quality using the 9-month sample because it allows us to control for workers' intrinsic ability levels. Table 4.2 presents regression results assessing the disparities, comparing the baseline model to the fully controlled model. We also ran the regressions without worker intrinsic abilities using the 3-year sample, which yield the same directional results as the 9-month sample as presented in Section 4.6.3.

Sufficiency Column 1 reveals baseline sufficiency gaps of 12.4% for White females, 6.2% for minority males, and 12.6% for minority females relative to White male workers, controlling only for store and week-of-year fixed effects. In Column 2, after incorporating controls for workers' tenure, roles, intrinsic abilities, and perceived long-term reliability, the sufficiency gaps for all of minority and female workers are narrowed, resulting in a final unexplained sufficiency gap of 10.7% for White females and 6.9% for minority females. Minority males also exhibit an unexplained sufficiency gap of 3.3%. Collectively, human capital factors explain 13.4% of the baseline sufficiency gap for White females, 46.4% for minority males, and 45.2% for minority females.

Last-minute Additions. Columns 3 through 6 present regression results assessing the predictability measures in terms of last-minute additions. Minority workers tend to receive more frequent RTA shifts compared to White workers (column 3). Minority females are found to be 2.6 percentage points (p.p) or 7.0% more likely to receive RTA shifts compared to all other groups, even after accounting for all human capital factors in tenure, role mix,

abilities, and absenteeism in the fully controlled model (column 4). Columns 5 and 6 report the corresponding results for the SNA ratio. After controlling for workers' human capital factors in column 6, White female and minority female workers are 0.8 p.p (or 7.2%) and 1.9 p.p (or 16.9%) more likely to receive SNA than White male workers, respectively. Minority male workers experience SNA 0.8 percentage points (or 7.4%) more frequently than their White male counterparts. Although this gap is only marginally significant in the 9-month sample ($p < 0.10$), it persists at a similar magnitude and is highly significant ($p < 0.01$) in the three-year dataset presented in Section 4.6.3.

Overall, the combination of human capital factors explain 0% to 15.3% of the baseline disparities in scheduling predictability depending on metric and worker group, but we continue to observe sizable gaps in the fully controlled models, especially for female workers.

Table 4.2: Sufficiency and Last-minute Addition Disparities.

Dependent Variables:	$\log(OfferedWeeklyHrs_{ijt} + 1)$		$RTARatio_{ijt}$		$SNARatio_{ijt}$	
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
<i>WhiteFemale_i</i>	-0.1237*** (0.0092)	-0.1073*** (0.0089)	-0.0013 (0.0041)	0.0034 (0.0039)	0.0080*** (0.0028)	0.0080*** (0.0028)
<i>MinorityMale_i</i>	-0.0617*** (0.0172)	-0.0331** (0.0165)	0.0135* (0.0070)	0.0001 (0.0068)	0.0071 (0.0049)	0.0082* (0.0049)
<i>MinorityFemale_i</i>	-0.1257*** (0.0132)	-0.0689*** (0.0129)	0.0285*** (0.0058)	0.0256*** (0.0056)	0.0222*** (0.0041)	0.0188*** (0.0041)
<i>MidTenure_{ijt}</i>		0.0791*** (0.0105)		-0.0073* (0.0042)		-0.0260*** (0.0032)
<i>LongTenure_{ijt}</i>		-0.0121 (0.0121)		-0.0180*** (0.0051)		-0.0478*** (0.0036)
<i>SalesAbility_{ijt}</i>		-0.0254 (0.0488)		0.0364* (0.0209)		0.0112 (0.0152)
<i>SpeedAbility_{ijt}</i>		0.2654*** (0.0483)		-0.2740*** (0.0213)		0.0219 (0.0154)
<i>ScoreAbility_{ijt}</i>		0.2856*** (0.0282)		0.0124 (0.0128)		0.0034 (0.0087)
<i>FullAbsenteesim_{ijt}</i>		-0.7782*** (0.0326)		-0.3767*** (0.0136)		0.1101*** (0.0143)
<i>PartialAbsenteesim_{ijt}</i>		-0.1290*** (0.0385)		-0.4463*** (0.0152)		-0.0232* (0.0121)
<i>OfferedWeeklyHrs_{ijt}</i>			-0.0053*** (0.0003)	-0.0056*** (0.0003)	-0.0149*** (0.0003)	-0.0149*** (0.0003)
<i>OfferedLTWeeklyHours_{ij}</i>			0.0045*** (0.0004)	0.0035*** (0.0005)	0.0107*** (0.0004)	0.0110*** (0.0004)
<i>Fixed-effects</i>						
<i>Store_j</i>	Yes	Yes	Yes	Yes	Yes	Yes
<i>Week_t</i>	Yes	Yes	Yes	Yes	Yes	Yes
<i>Role_i</i>		Yes		Yes		Yes
<i>Fit statistics</i>						
Observations	29,767	29,767	29,014	29,014	28,922	28,922
R ²	0.04351	0.12177	0.06298	0.12582	0.19666	0.20630

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

4.5 *Schedule Disparities - The Role of Worker Preference and Availability*

In the section above, we control for key human capital factors, yet sizable gaps remain. Here, we explore potential drivers related to worker preference and availability.

Our main observation indicates that female and minority workers are initially offered fewer hours but experience more frequent subsequent additions. This pattern is particularly pronounced among female workers, for whom the disparities are larger in magnitude and statistically more significant compared to minority males. Therefore, our examination of mechanisms centers on gender differences, followed by a shorter discussion of racial disparities in Section 4.6.

To better understand this phenomenon, we break down our analysis of worker preference and availability into several potential mechanisms that may explain this outcome.

Mechanism 1. Women prefer to work **more hours** but are initially **under-offered**.

Mechanism 2. Women prefer to work **fewer hours** but are later **imposed** with additional shifts.

Mechanism 3. Women prefer the **flexibility** to start with fewer offered hours, with the option to add hours closer to actual operation.

Mechanism 1 aligns with findings in the literature that hourly workers especially in the food-service sector often desire more hours than they receive [Schneider and Harknett, 2019b], suggesting that this phenomenon may be even more pronounced among female workers. Mechanism 2 relates to a well-documented driver of workplace gender gaps, where women have lower availability due to heavier domestic responsibilities [Blau and Kahn, 2017, Cook et al., 2021]. However, if these workers indeed preferred fewer hours, the imposition of more frequent additions would contrast their preference. Mechanism 3 is consistent with a phenomenon suggested in the literature on gender pay gaps, where women may value temporal flexibility or greater control over their working hours [Goldin, 2014]. For instance, female workers may experience greater uncertainty in their personal lives due to heavier domestic responsibilities, making it challenging to commit to a set schedule a week in advance. Consequently, they may prefer fewer hours offered initially, with additional hours added closer to the workday, when they have greater certainty about their availability.

4.5.1 Ruling Out Mechanism 3 - Preference for Last-Minute Additions

Mechanism 3 suggests that, for an identical shift, women prefer it to be given as a last-minute addition over a regular shift with a week’s advance notice due to the flexibility it provides. To test if this is indeed the case, we use worker’s sales performance as an indicator of preference. Our theory is that if female workers do prefer last-minute hours, their sales performance should be better during such times than during regular shifts as they might exert extra effort as an act of goodwill and reciprocation [Spector and Fox, 2002, Avgoustaki and Bessa, 2019], or to secure more such opportunities in the future [Fehr et al., 2009].

To compare workers’ sales performance on shifts with different levels of advance notice, we follow the approach proposed in Kamalahmadi et al. [2021]. The analysis is conducted at the transaction level using the POS data described in Section 4.3. The check size for transaction k , denoted as $Sale_k$, is the outcome variable representing the handling server’s sales performance. Servers can influence the check size through their cross-selling and up-selling efforts. The primary explanatory variables are SNA_k and RTA_k , indicators of whether the transaction was opened during a handling worker’s SNA or RTA shift, respectively. To estimate the impact of advance notice on workers’ sales performance, we adopt the following model:

$$\log(Sale_k) = \alpha_0 + \alpha_1 SNA_k + \alpha_2 RTA_k + \alpha_3 \mathbf{W}_k + \eta_k. \quad (4.3)$$

where

$$\begin{aligned} \mathbf{W}_k = & (GuestCount_k, \log(Tenure_k), StoreLoad_k, WorkerLoad_k, \\ & WorkerLoad_k^2, \log(Fatigue_k), EndOfShift_k, Hour_k, \\ & DayWeek_k, WeekYear_k, Month_k, Store_k, Worker_k). \end{aligned} \quad (4.4)$$

To rule out potential confounding, we include a number of variables that can impact the check size. Specifically, at the store level, we control for the store fixed effects denoted as $Store_k$, which account for the store-specific characteristics that are time-invariant such as store location and culture characteristics. We also control for restaurant load, denoted as $StoreLoad_k$. It is defined as the average number of customers present during the focal transaction, which proxies for kitchen congestion and dining atmosphere. At the server

level, we control for server fixed effects denoted as $Worker_k$, capturing time-invariant server specific characteristics such as their intrinsic abilities and personalities. We also control for server workload and experience, denoted by $WorkerLoad_k$ and $Tenure_k$, respectively. $WorkerLoad_k$ is the average number of transactions the server is handling during the focal transaction, and we include $WorkerLoad_k^2$ to account for non-linear effects, as suggested by prior work [Tan and Netessine, 2014]. $Tenure_k$ represents the number of days since the server was hired.

Regarding the timing of the transaction, we include time dummies, i.e., $Hour_k$, $DayWeek_k$, $WeekYear_k$, and $Month_k$, which capture temporal heterogeneity and seasonal trends in sales. We also account for two additional schedule-related effects: $Fatigue_k$, the number of transactions the server completed earlier that day (capturing the fatigue effect), and $EndOfShift_k$, a dummy equal to 1 if the transaction began during the last hour of the server’s shift (capturing the end-of-shift effect). Lastly, while customers characteristics could also impact check size, they will not bias our estimates on the impact of last-minute shifts. This is because customers were assigned to servers in a round-robin fashion at our restaurants. This random assignment ensures that the mix of customers across the regular and last-minute shifts shall be statistically the same (conditional on other controls in our model). Nevertheless, we include the number of customers associated with the check, denoted as $GuestCount_k$, as a control, which helps improve our estimation efficiency.

Even with the random assignment of customers and the extensive controls in our model, there may still be unobservable confounding factors, such as managers’ private information about the demand, which may be correlated with both the use of last-minute shifts (SNA_k and RTA_k) and check size ($Sale_k$). It will in turn lead to biased estimates. To address this endogeneity concern, we use the instrumental variable approach. Specifically, following Kamalahmadi et al. [2021], we use the percentage of servers on SNA or RTA shifts in the same hour from the previous week ($Lagged\%SNA_k$ and $Lagged\%RTA_k$) as instruments for SNA_k and RTA_k . They are valid instruments for the following reasons: (1) They satisfy the relevance condition. Managers at each store make staffing and scheduling decisions according to company guidelines, with some additional discretion. Given the persistence inherent in these decision-making processes, we expect $Lagged\%SNA_k$ and $Lagged\%RTA_k$

to correlate with SNA_k and RTA_k , respectively. Our results confirm this expectation, as detailed in Table 4.3; (2) They satisfy the exclusion condition. Specifically, staffing and scheduling decisions made in the previous week should not directly influence the sales of focal transaction in the current week. However, a potential concern with our lagged instrument is that sales might exhibit temporal correlation due to unobservable common shocks. As a result, our instruments could be correlated with the error term η_k through its serial correlation with error terms from transactions occurring during the same hour in the previous week. To examine this possibility, we conduct the serial correlation test described in Chapter 15.7 of Wooldridge [2016]. Our results reject the hypothesis of serial correlation in the error terms across weeks, thus supporting the exclusion condition.

Table 4.3: First-Stage Regression Results.

Dependent Variables:	SNA_k	RTA_k
Model:	(1)	(2)
<i>Variables</i>		
$Lagged\%SNA_k$	0.0934*** (0.0114)	0.0317*** (0.0050)
$Lagged\%RTA_k$	-0.0009 (0.0067)	0.1777*** (0.0089)
Controls and Fixed-effects	Yes	Yes
Observations	1,487,647	1,487,647
R ²	0.11021	0.18684

Fixed effects on Hour_k, DayWeek_k, WeekYear_k, Month_k, Store_k, Worker_k.

Clustered (EmpID) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

We then estimate our model in Equation (4.3) using the Probit-2SLS method described in chapter 18.4 of Woodridge [2010]. Probit-2SLS is commonly used when the endogenous variables are binary, which tends to be more efficient compared to the classic 2SLS estimators. We present our estimation results in Table 4.4, and refer the details of the estimation procedure to Appendix C.4. By incorporating a comprehensive set of controls and leveraging our instrumental variables, our estimates of the effects of SNA and RTA should reflect differences in sales performance driven specifically by variations in advance notice of the shifts, rather than by other confounding factors.

Consistent with prior findings in Kamalahmadi et al. [2021], our analysis shows that workers exhibit approximately 3.28% lower sales performance during RTA shifts compared to regular shifts. This suggests that workers prefer regular shifts to real-time shifts. Instead of offering flexibilities, real-time shifts creates inconvenience for workers due to the unpredictable nature of these shifts. Additionally, we find no significant difference in performance between SNA shifts and regular shifts, indicating workers are indifferent to short-notice additions. As suggested in Kamalahmadi et al. [2021], this may be because the positive *income opportunity effect* cancels out the *unpredictable effect* of the SNA. Chronic underscheduling among food-service and hourly workers nationwide [Golden, 2016, Schneider and Harknett, 2019b, Golden and Kim, 2020] may lead employees to view last-minute shifts as valuable opportunities—or “gifts”—for extra income [Akerlof, 1982]. This perception can drive greater effort, either from feelings of goodwill [Spector and Fox, 2002, Avgoustaki and Bessa, 2019] or as a strategic move to secure better future scheduling [Fehr et al., 2009]. We refer to this phenomenon as the income-opportunity effect, following Kamalahmadi et al. [2021]. If last-minute shifts truly offer desirable flexibility as Mechanism 3 suggests, women’s performance should be higher during these shifts than during regularly scheduled ones. However, our results suggest it is not the case: last-minute shifts create inconvenience rather than flexibility for workers, which we refer to as the unpredictability effect.

This main model above allows us to infer all workers’ average preference between regular and last-minute notice. To rule out the possibility where women and men have different preferences, as suggested in Mechanism 3, we also ran an extension model including interaction terms between the SNA and RTA indicator and worker groups. The results for this

Table 4.4: Productivity for Different Shift Types.

Dependent Variable: Model:	$\log(Sale_k)$	
	(1)	(2)
<i>Variables</i>		
SNA_k	0.0037 (0.0187)	0.0217 (0.0265)
RTA_k	-0.0328*** (0.0069)	-0.0401*** (0.0091)
$GuestCount_k$	0.3259*** (0.0011)	0.3259*** (0.0011)
$StoreLoad_k$	0.0003*** (2.63×10^{-5})	0.0003*** (2.63×10^{-5})
$\log(Tenure_k)$	0.0025 (0.0018)	0.0025 (0.0018)
$WorkerLoad_k^2$	-0.0044*** (0.0001)	-0.0044*** (0.0001)
$WorkerLoad_k$	-0.0193*** (0.0005)	-0.0193*** (0.0005)
$\log(Fatigue_k)$	-0.0016** (0.0007)	-0.0016** (0.0007)
$EndOfShift_k$	-0.0980*** (0.0019)	-0.0980*** (0.0019)
$SNA_k \times MinorityMale_k$		-0.0228 (0.0300)
$SNA_k \times MinorityFemale_k$		-0.0374 (0.0302)
$SNA_k \times WhiteFemale_k$		-0.0173 (0.0239)
$RTA_k \times MinorityMale_k$		0.0174 (0.0131)
$RTA_k \times MinorityFemale_k$		0.0090 (0.0133)
$RTA_k \times WhiteFemale_k$		0.0085 (0.0096)
<i>Fit statistics</i>		
Observations	1,330,485	1,330,485
R ²	0.64432	0.64426

Fixed effects on Hour_k, DayWeek_k, WeekYear_k, Month_k, Store_k, Worker_k.

Clustered (Worker) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

extended models are presented in the second column of Table 4.4. Our results show that RTA and SNA affect all workers equally, regardless of their race or gender.

The above analysis demonstrate that, for an identical shift, all workers prefer them to be scheduled in advance, with no significant differences in preference by race or gender. Next, we show that, in practice, RTA and SNA shifts are typically scheduled during less lucrative hours compared to regular shifts. This scheduling pattern further decreases the likelihood that women prefer last-minute shifts to regular shifts.

To compare the earning potential of regular and last-minute shifts, we first estimate the hourly earning potential for each hour-of-week at the 25 operating locations. With a constant hourly minimum wage across all hours, hourly earning potential is defined by the expected tip income per hour, which drives income variation. For each unique hour in the data, we calculate hourly earnings per worker by dividing total customer tips⁶ from transactions occurring in that hour by the number of workers present. We then calculate the average hourly earning potential per worker for each unique location, day-of-week, hour-of-day combination.

To estimate differences in earning potential across shift types, we break down workers' shifts into individual hours and map each hour to its corresponding hourly earning potential. We then calculate the average hourly earning potential for all regular and last-minute shifts of the focal worker within the focal week, weighting each hour by the duration worked in that hour. This approach enables us to construct the expected hourly earning potential at the worker-location-week level, separately for regular shifts, RTA shifts, and SNA shifts.

We then run four separate regressions—one each for White males, White females, minority males, and minority females. In each case, we model

$$\log(HrlyEP_{ijts}) = \beta_0 + \beta_1 \mathbf{1}\{s = \text{SNA}\} + \beta_2 \mathbf{1}\{s = \text{RTA}\} + \beta_3 Y_{ijt} + u_{ijts}, \quad (4.5)$$

where $HrlyEP_{ijts}$ denotes the expected hourly earning potential of worker i under schedule type s at location j in week t , $Y_{ijt} = (\text{Store}_j, \text{Week}_t, \text{Worker}_i)$ captures store, week, and

⁶We estimate tip amounts as 15% of total transaction revenue. Our results remain robust to this assumption, as long as the percentage of tips relative to sales does not systematically vary across different schedule types.

worker fixed effects, and u_{ijts} is the idiosyncratic error term. Regression results are shown in Table 4.5.

Table 4.5: Differences in Hourly Earning Potential for RTA and SNA Shifts Relative to Regular Shifts.

Model:	White Male	White Female	Minority Male	Minority Female
<i>Variables</i>				
RTA_s	-0.4398*** (0.0174)	-0.4199*** (0.0135)	-0.4125*** (0.0271)	-0.3981*** (0.0211)
SNA_s	-0.0205*** (0.0055)	-0.0225*** (0.0045)	-0.0360*** (0.0117)	-0.0316*** (0.0079)
Controls and fixed-effects	Yes	Yes	Yes	Yes
Observations	18,338	29,326	4,503	8,292
R ²	0.53449	0.54157	0.55012	0.50174

Clustered ($Worker_i$) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

As shown in Table 4.5, compared to regular shifts, RTA shifts yield an astonishing reduction in hourly earning potential of around 40% for all groups. While SNA shifts offer better earnings than RTA shifts, they still fall short of regular hours by 2-4%. For RTA, the discrepancy arises because RTA shifts typically occur during off-peak periods immediately following lunch and dinner times. On the other hand, both SNA and regular shifts are concentrated during peak times of each day. However, while regular shifts tend to cluster around weekends when demand and earning potential are higher, SNA shifts are more evenly distributed throughout the week as they are primarily driven by worker absences. This results in a lower earning potential for SNA shifts. In sum, both RTA and SNA shifts offer poorer income opportunities than regular shifts, making it even less likely that workers

would prefer last-minute shifts over regular ones.

Finally, as previously noted, while regular shifts are guaranteed as managers do not proactively cancel shifts at our restaurants, last-minute additions often emerge unpredictably from peer absences or sudden demand fluctuations. This inherent uncertainty further diminishes women's likelihood of preferring last-minute shifts over regular ones.

Given the inconvenience, lower earning potential, and lack of guaranteed hours associated with last-minute shifts, we conclude that Mechanism 3 - the notion that women prefer the flexibility afforded by low sufficiency and frequent last-minute scheduling - does not explain the observed disparities.

4.5.2 Ruling Out Mechanism 2 - Preference for Fewer Hours

Between Mechanism 1 and Mechanism 2, we argue that Mechanism 2 is unlikely the driver of the unexplained disparities. The distinction between these mechanisms lies in whether workers desire additional hours beyond the initially offered ones. To examine this, we revisit the sales performance comparison in Table 4.4. If workers did not value the extra hours, last-minute shifts would not represent a positive income opportunity but only an inconvenience due to their unpredictability, which should result in a significant drop in sale performance during both SNA and RTA shifts for female workers. However, the finding that women's sales performance on SNA shifts is comparable to their performance on regular shifts—paralleling the pattern observed for men—suggests that women do value these additional SNA hours. This together with our arguments in Section 4.5.1 implies that women indeed want to work beyond their scheduled shifts but would prefer these additional hours to be scheduled in advance rather than added last minute.

Hereby, we conclude that Mechanism 1 is the driving mechanism behind the unexplained disparities. Managers tend to prioritize assigning hours to male workers in advance, leaving female workers with fewer initially scheduled hours. As a result, women are more often compelled to take on last-minute shifts to compensate for the shortfall, even though these shifts are generally less profitable and more inconvenient.

4.5.3 *Assessing Mechanism 1 - The Preference for More Hours and Its Magnitude*

At this point, we have provided compelling evidence that women indeed desire more hours but are initially under-offered. However, we cannot attribute all the unexplained offered sufficiency disparities in Table 4.2 solely to under-offering, as female workers may indeed have lower availability than male workers, as suggested by Blau and Kahn [2017] and Cook et al. [2021]. Therefore, the next step is to assess how much more they would like to work. While direct measures of availability are lacking as it is the case in almost all scheduling data available in practice, the number of hours that workers actually worked provides a conservative estimate of their availability. To investigate further, we estimate gender-racial gaps in workers' sufficiency based on their worked hours, as shown in Table 4.6, and compare those gaps with the gaps in offered hours.

We observe that White females have a sufficiency gap of 8.9% in worked hours after accounting for tenure, role, intrinsic abilities, and absenteeism, which is 1.8% lower than their sufficiency gap in offered hours as shown in Column 2 of Table 4.2. Meanwhile, minority workers (both male and female) fully close the offered sufficiency gap in the hours they actually work. The smaller *worked* sufficiency gaps relative to *offered* sufficiency gaps for women suggest that they are more available than their offered hours reflect. Specifically, while White females may indeed be less available than White males, the 1.80% reduced gap for White females and the complete closure of the 6.89% offered sufficiency gap for minority females are likely to arise from under-offering rather than unavailability.

Summary. Our analysis in Section 4.5.1 shows that women's lower average offered sufficiency and higher incidence of last-minute shift additions are not driven by a preference for such scheduling. While lower availability may account for part of the disparity in offered sufficiency among white women, it does not explain the gaps observed for minority women. Rather, these disparities appear to stem from systematic under-scheduling at the outset, with women subsequently being more frequently compelled to accept last-minute additions.

Table 4.6: Actually Worked Sufficiency.

Dependent Variable:	$\log(WorkedWeeklyHrs_{ijt} + 1)$
Model:	(1)
<i>Variables</i>	
<i>WhiteFemale_i</i>	-0.0893*** (0.0092)
<i>MinorityMale_i</i>	0.0099 (0.0168)
<i>MinorityFemale_i</i>	-0.0139 (0.0135)
Controls and fixed-effects	Yes
Observations	29,767
R ²	0.18220

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

4.6 Discussions

After documenting scheduling disparities that cannot be fully explained by workers' human capital, preferences, or availability, this section examines the racial patterns within these residual gaps in greater details. We then explore additional factors that may contribute to the unexplained disparities. Finally, we present regression results from a three-year dataset to assess the robustness and generalizability of our findings.

4.6.1 *Racial Disparities*

We find no theoretical basis in the literature for racial differences in availability or flexibility preferences, as is sometimes argued for gender disparities. Nonetheless, minority workers as a group face similar scheduling disparities as women in our data.

Minority males experience a 3.31% offered sufficiency gap relative to White men, which disappears once worked sufficiency is considered. Compared to White women, although minority women generally have higher sufficiency, they face a larger difference in the offered and worked sufficiency gaps. This indicates a greater discrepancy between their offered sufficiency and actual availability. Additionally, minority women encounter more frequent real-time and short-notice additions than their White counterparts, even when they have the same number of scheduled hours in advance.

Combining the racial comparison for both genders, it becomes evident that the same gender disparity pattern exists across races: compared to White workers of the same gender, minority workers are more likely to be under-offered and receive more frequent last-minute additions, both not driven by their preference or availability. In the absence of neither empirical nor theoretical evidence suggesting lower availability or a stronger preference for last-minute additions among minority workers, our findings point to race-based disadvantages. Furthermore, our findings align with previous research suggesting that race and gender disparities often compound, exacerbating the challenges faced by women of color [Crenshaw, 2013, Collins, 2022, Acker, 2006], as demonstrated in studies on wage gaps [Jones and Schmitt, 2016] and customer biases [Kamalahmadi et al., 2023]. Notably, our results are consistent with Storer et al. [2020], who found that women of color experience the worst temporal quality of work schedules nationally.

4.6.2 *Additional Factors*

Our findings indicate substantial gender and racial disparities in scheduling quality that persist even after careful consideration of human capital, worker availability, and preferences. This suggests that additional factors may be influential yet difficult to measure empirically.

One such factor is managerial bias, documented extensively in prior literature across

various workplace settings. Managers may consciously or unconsciously rely on taste-based discrimination, which involves personal prejudice [Becker, 2010]; status-based discrimination, shaped by societal hierarchies and stereotypes [Ridgeway, 2001]; or statistical discrimination, where decisions reflect generalized group-level beliefs rather than individual attributes [Phelps, 1972, Aigner and Cain, 1977]. Such biases, while subtle, could systematically disadvantage specific demographic groups in scheduling decisions.

Furthermore, social pressure and negotiation dynamics may contribute significantly to scheduling disparities. For example, even without managerial imposition, women and minority workers may experience heightened internalized pressure to accept unfavorable scheduling arrangements, conforming to gendered and racial stereotypes emphasizing agreeableness or compliance [Eagly and Wood, 1991, Babcock et al., 2017, Komarraju et al., 2008]. Additionally, disparities in negotiation power—stemming from differential abilities, outside options, or biased managerial responses—could reinforce these inequalities, making it challenging for disadvantaged groups to advocate successfully for preferred scheduling arrangements [Bowles et al., 2007, Hernandez et al., 2019].

These factors—managerial biases, social pressure, and negotiation power—are inherently interrelated and complex, posing substantial challenges to empirical disentanglement. While our study provides *suggestive* evidence that these factors may play a role in the unexplained schedule disparity, further research would be valuable to more *explicitly* explore these nuanced dynamics in scheduling decisions in practice.

4.6.3 Results On The Three-Year Sample

We next show that the gender and racial disparities identified in our main 9-month analysis persist when replicated using the full three-year dataset, as shown in Table C.4. While the main analysis leverages the shorter 9-month window to incorporate detailed controls for workers’ intrinsic abilities, the three-year dataset excludes these detailed ability controls but benefits from a larger number of worker-week observations. Thus, this analysis offers a robust validity check and reinforces the generalizability of our main conclusions.

The three-year regressions in Table 7 echo—and sharpen—the nine-month findings. All

Table 4.7: Schedule Disparities Using 3-yr Full Sample

Dependent Variables:	$\log(\text{OfferedWeeklyHrs}_{ijt} + 1)$	$RTARatio_{ijt}$	$SNARatio_{ijt}$	$\log(\text{WorkedWeeklyHrs}_{ijt} + 1)$
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
<i>WhiteFemale_i</i>	-0.0523*** (0.0046)	0.0104*** (0.0020)	0.0086*** (0.0014)	-0.0363*** (0.0048)
<i>MinorityMale_i</i>	-0.0148* (0.0080)	0.0108*** (0.0034)	0.0159*** (0.0024)	0.0219*** (0.0081)
<i>MinorityFemale_i</i>	-0.0449*** (0.0067)	0.0238*** (0.0030)	0.0187*** (0.0021)	0.0014 (0.0070)
<i>MidTenure_{ijt}</i>	-0.0065 (0.0050)	-0.0119*** (0.0021)	-0.0170*** (0.0015)	-0.0238*** (0.0052)
<i>LongTenure_{ijt}</i>	-0.0497*** (0.0062)	-0.0190*** (0.0027)	-0.0384*** (0.0018)	-0.1011*** (0.0064)
<i>FullAbsenteesim_{ijt}</i>	-1.017*** (0.0233)	-0.4536*** (0.0088)	0.1366*** (0.0097)	-2.014*** (0.0268)
<i>PartialAbsenteesim_{ijt}</i>	-0.4479*** (0.0313)	-0.6296*** (0.0102)	0.0701*** (0.0099)	-0.7090*** (0.0239)
<i>OfferedWeeklyHrs_{ijt}</i>		-0.0050*** (0.0002)	-0.0121*** (0.0001)	
<i>OfferedLTWeeklyHours_{ij}</i>		0.0020*** (0.0002)	0.0091*** (0.0002)	
<i>Fixed-effects</i>				
<i>Store_j</i>	Yes	Yes	Yes	Yes
<i>Year_t</i>	Yes	Yes	Yes	Yes
<i>Week_t</i>	Yes	Yes	Yes	Yes
<i>Role_i</i>	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	112,524	109,926	109,599	112,524
R ²	0.11298	0.10011	0.18126	0.17501

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

coefficients retain the same signs: women and racial minorities continue to receive fewer pre-scheduled hours while shouldering a disproportionate share of short-notice and real-time shifts. Predictability gaps remain largely unchanged but become more statistically significant, particularly for minority men—the smallest demographic group in the sample. Although disparities in offered hours narrow slightly, the gap between offered and worked hours remains virtually unchanged, indicating a persistent and consistent pattern of under-scheduling. Overall, the persistence of these patterns through three years of seasonal and demand swings confirms that the disparities are structural, not era-specific.

4.7 Concluding Remarks

This study provides the first large-scale empirical evidence of gender and racial disparities in the temporal quality of work schedules using field data—an often-overlooked yet critical dimension of job quality. Leveraging granular shift-level, transaction, and customer rating data from a national restaurant chain, we document persistent gaps in scheduling sufficiency and predictability across workers of different race and gender. Even after accounting for a rich set of human capital variables—including role composition, tenure, absenteeism, and multidimensional performance measures—women and minority workers, especially minority women, are systematically offered fewer hours and subjected to more frequent last-minute shift additions compared to White male counterparts.

By ruling out classic explanations such as differences in worker preferences or availability, our analysis attributes much of the unexplained disparity to systematic under-scheduling at the outset, particularly for women. This initial shortfall compels affected workers to rely more heavily on unpredictable, less lucrative last-minute shifts—undermining their earnings potential, stability, and well-being.

Our findings contribute to both the operations and workplace inequality literatures by introducing a robust empirical framework for evaluating schedule disparities using real-world field data. From a managerial perspective, our study highlights the unintended consequences of just-in-time scheduling and underscores the need for more equitable scheduling practices. Firms aiming to retain a diverse and productive workforce must re-evaluate scheduling algorithms and decision-making processes to ensure that flexibility and efficiency do not

come at the cost of fairness and inclusion.

BIBLIOGRAPHY

- Joan Acker. Inequality regimes: Gender, class, and race in organizations. *Gender & society*, 20(4):441–464, 2006.
- Sam Aflaki and Serguei Netessine. Strategic investment in renewable energy sources: The effect of supply intermittency. *Manufacturing & Service Operations Management*, 19(3):489–507, 2017.
- Vishal Agrawal and Deishin Lee. The effect of sourcing policies on suppliers’ sustainable practices. *Production and Operations Management*, 28(4):767–787, 2019.
- Vishal Agrawal and Şafak Yücel. Renewable energy sourcing. *Responsible Business Operations: Challenges and Opportunities*, pages 211–224, 2021.
- Dennis J Aigner and Glen G Cain. Statistical theories of discrimination in labor markets. *Ilr Review*, 30(2):175–187, 1977.
- George A Akerlof. Labor contracts as partial gift exchange. *The quarterly journal of economics*, 97(4):543–569, 1982.
- Joseph G Altonji and Rebecca M Blank. Race and gender in the labor market. *Handbook of labor economics*, 3:3143–3259, 1999.
- Chad Augustine and Nathan Blair. Energy storage futures study: Storage technology modeling input data report. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2021.
- Argyro Avgoustaki and Ioulia Bessa. Examining the link between flexible working arrangement bundles and employee work effort. *Human Resource Management*, 58(4):431–449, 2019.

Linda Babcock and Sara Laschever. Women don't ask: Negotiation and the gender divide. In *Women Don't Ask*. Princeton University Press, 2009.

Linda Babcock, Maria P Recalde, Lise Vesterlund, and Laurie Weingart. Gender differences in accepting and receiving requests for tasks with low promotability. *American Economic Review*, 107(3):714–747, 2017.

Galen L Barbose. Us renewables portfolio standards 2021 status update: Early release. 2021.

John Manuel Barrios, Laura M Giuliano, and Andrew J Leone. In living color: Does in-person screening affect who gets hired? *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-38), 2020.

Gary S Becker. *The economics of discrimination*. University of Chicago press, 2010.

Alon Bergman, Guy David, and Hummy Song. “i quit”: Schedule volatility as a driver of voluntary employee turnover. *Manufacturing & Service Operations Management*, 25(4):1416–1435, 2023.

Marianne Bertrand and Sendhil Mullainathan. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American economic review*, 94(4):991–1013, 2004.

Atul Bhandari, Alan Scheller-Wolf, and Mor Harchol-Balter. An exact and efficient algorithm for the constrained dynamic operator staffing problem for call centers. *Management Science*, 54(2):339–353, 2008.

Bird Bird. Collaborative offsite PPA structures for renewable energy buyers, 2018. URL <https://www.twobirds.com/-/media/pdfs/joiningtheclubcollaborativeoffsiteppastructuresforrenewableener.pdf>.

Tomas Björk, Agatha Murgoci, and Xun Yu Zhou. Mean–variance portfolio optimization with state-dependent risk aversion. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 24(1):1–24, 2014.

- Francine D Blau and Lawrence M Kahn. Gender differences in pay. *Journal of Economic perspectives*, 14(4):75–100, 2000.
- Francine D Blau and Lawrence M Kahn. The gender pay gap: Going, going... but not gone. *The declining significance of gender*, pages 37–66, 2006a.
- Francine D Blau and Lawrence M Kahn. The us gender pay gap in the 1990s: Slowing convergence. *Ilr Review*, 60(1):45–66, 2006b.
- Francine D Blau and Lawrence M Kahn. The gender wage gap: Extent, trends, and explanations. *Journal of economic literature*, 55(3):789–865, 2017.
- Trine Krogh Boomsma, Nigel Meade, and Stein-Erik Fleten. Renewable energy investments under different support schemes: A real options approach. *European Journal of Operational Research*, 220(1):225–237, 2012.
- Hannah Riley Bowles and Kathleen L McGinn. 2 untapped potential in the study of negotiation and gender inequality in organizations. *The Academy of Management Annals*, 2(1):99–132, 2008.
- Hannah Riley Bowles, Linda Babcock, and Lei Lai. Social incentives for gender differences in the propensity to initiate negotiations: Sometimes it does hurt to ask. *Organizational Behavior and human decision Processes*, 103(1):84–103, 2007.
- A Silvia Cancio, T David Evans, and David J Maume Jr. Reconsidering the declining significance of race: Racial differences in early career wages. *American Sociological Review*, pages 541–556, 1996.
- Felipe Caro, Charles J Corbett, Tarkan Tan, and Rob Zuidwijk. Double counting in supply chain carbon footprinting. *Manufacturing & Service Operations Management*, 15(4):545–558, 2013.
- Suchan Chae and Paul Heidhues. Buyers’ alliances for bargaining power. *Journal of economics & management strategy*, 13(4):731–754, 2004.

Li Chen and Hau L Lee. Sourcing under supplier responsibility risk: The effects of certification, audit, and contingency payment. *Management Science*, 63(9):2795–2812, 2017.

Rachel R Chen and Paolo Roma. Group buying of competing retailers. *Production and operations management*, 20(2):181–197, 2011.

Shi Chen, Qinqin Zhang, and Yong-Pin Zhou. Impact of supply chain transparency on sustainability under ngo scrutiny. *Production and Operations Management*, 28(12):3002–3022, 2019.

Joshua Choper, Daniel Schneider, and Kristen Harknett. Uncertain time: Precarious schedules and job turnover in the us service sector. *ILR Review*, 75(5):1099–1132, 2022.

Lisa E Cohen, Joseph P Broschak, and Heather A Haveman. And then there were more? the effect of organizational sex composition on the hiring and promotion of managers. *American Sociological Review*, pages 711–727, 1998.

Philip N Cohen and Matt L Huffman. Working for the woman? female managers and the gender wage gap. *American sociological review*, 72(5):681–704, 2007.

Patricia Hill Collins. *Black feminist thought: Knowledge, consciousness, and the politics of empowerment*. routledge, 2022.

Cody Cook, Rebecca Diamond, Jonathan V Hall, John A List, and Paul Oyer. The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *The Review of Economic Studies*, 88(5):2210–2238, 2021.

Kimberlé Crenshaw. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. In *Feminist legal theories*, pages 23–51. Routledge, 2013.

Amy JC Cuddy, Susan T Fiske, and Peter Glick. When professionals become mothers, warmth doesn’t cut the ice. *Journal of Social issues*, 60(4):701–718, 2004.

- Ruomeng Cui, Jun Li, and Dennis J Zhang. Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on airbnb. *Management Science*, 66(3):1071–1094, 2020.
- Lingxiu Dong. Toward resilient agriculture value chains: challenges and opportunities. *Production and Operations Management*, 30(3):666–675, 2021.
- Lingxiu Dong, Iva Rashkova, and Duo Shi. Food safety audits in developing economies: Decentralization vs. centralization. *Manufacturing & Service Operations Management*, 24(6):2863–2881, 2022.
- Alice H Eagly. *Sex differences in social behavior: A social-role interpretation*. Psychology Press, 2013.
- Alice H Eagly and Wendy Wood. Explaining sex differences in social behavior: A meta-analytic perspective. *Personality and social psychology bulletin*, 17(3):306–315, 1991.
- Fred F Easton. Service completion estimates for cross-trained workforce schedules under uncertain attendance and demand. *Production and Operations Management*, 23(4):660–675, 2014.
- Adel Elomri, Asma Ghaffari, Zied Jemai, and Yves Dallery. Coalition formation and cost allocation for joint replenishment systems. *Production and Operations Management*, 21(6):1015–1027, 2012.
- Nesim Erkip, Warren H Hausman, and Steven Nahmias. Optimal centralized ordering policies in multi-echelon inventory systems with correlated demands. *Management Science*, 36(3):381–392, 1990.
- Ernst Fehr, Lorenz Goette, and Christian Zehnder. A behavioral account of the labor market: The role of fairness concerns. *Annu. Rev. Econ.*, 1(1):355–384, 2009.
- Qi Feng and Lauren Xiaoyuan Lu. Supply chain contracting under competition: Bilateral bargaining vs. stackelberg. *Production and Operations Management*, 22(3):661–675, 2013.

- Qi Feng, Guoming Lai, and Lauren Xiaoyuan Lu. Dynamic bargaining in a supply chain with asymmetric demand information. *Management Science*, 61(2):301–315, 2015.
- Benedikt Finnah, Jochen Gönsch, and Florian Ziel. Integrated day-ahead and intraday self-schedule bidding for energy storage systems using approximate dynamic programming. *European Journal of Operational Research*, 301(2):726–746, 2022.
- Ryan Finnigan and Savannah Hunter. Occupational composition and racial/ethnic inequality in varying work hours in the great recession. In *Race, Identity and Work*. Emerald Publishing Limited, 2018.
- Garrett Fitzgerald, James Mandel, Jesse Morris, and Hervé Touati. The economics of battery energy storage: How multi-use, customer-sited batteries deliver the most services and value to customers and the grid. Technical Report —, Rocky Mountain Institute, October 2015. URL <https://rmi.org/wp-content/uploads/2017/03/RMI-TheEconomicsOfBatteryEnergyStorage-FullReport-FINAL.pdf>. Published October 2015; accessed August 15, 2025.
- Ke Fu, Guoming Lai, Weixin Shang, and Jiayan Xu. Group buying between competitors: Exogenous and endogenous? power structures. *Production and Operations Management*, 33(10):1997–2013, 2024.
- Noah Gans, Ger Koole, and Avishai Mandelbaum. Telephone call centers: Tutorial, review, and research prospects. *Manufacturing & Service Operations Management*, 5(2):79–141, 2003.
- Noah Gans, Haipeng Shen, Yong-Pin Zhou, Nikolay Korolev, Alan McCord, and Herbert Ristock. Parametric forecasting and stochastic programming models for call-center workforce scheduling. *Manufacturing & Service Operations Management*, 17(4):571–588, 2015.
- Zuguang Gao, Nur Sunar, and John R Birge. Designing renewable power purchase agreements: Impact on green energy investment. *Available at SSRN 4670536*, 2023.
- Yashar Ghiassi-Farrokhfal, Rodrigo Belo, Mohammed Reza Hesamzadeh, and Derek Bunn.

- Optimal electricity imbalance pricing for the emerging penetration of renewable and low-cost generators. *Manufacturing & Service Operations Management*, 2023.
- Laura Giuliano, David I Levine, and Jonathan Leonard. Manager race and the race of new hires. *Journal of Labor Economics*, 27(4):589–631, 2009.
- Lonnie Golden. Irregular work scheduling and its consequences. *Work-Life Balance in the Modern Workplace*, page 115, 2015.
- Lonnie Golden. Still falling short on hours and pay: Part-time work becoming new normal. *Available at SSRN 2881673*, 2016.
- Lonnie Golden and Jaeseung Kim. Underemployment just isn’t working for us part-time workers. *The Center for Law and Social Policy*, 2020.
- Claudia Goldin. A grand gender convergence: Its last chapter. *American economic review*, 104(4):1091–1119, 2014.
- Claudia Goldin and Lawrence F Katz. The power of the pill: Oral contraceptives and women’s career and marriage decisions. *Journal of political Economy*, 110(4):730–770, 2002.
- Claudia Goldin and Cecilia Rouse. Orchestrating impartiality: The impact of “blind” auditions on female musicians. *American economic review*, 90(4):715–741, 2000.
- Shadi Goodarzi, Sam Aflaki, and Andrea Masini. Optimal feed-in tariff policies: The impact of market structure and technology characteristics. *Production and Operations Management*, 28(5):1108–1128, 2019.
- Sanjith Gopalakrishnan, Daniel Granot, Frieda Granot, Greys Sošić, and Hailong Cui. Incentives and emission responsibility allocation in supply chains. *Management Science*, 67(7):4172–4190, 2021.
- Greenbiz. Walmart suppliers join forces to buy renewable electricity, 2022. URL <https://www.greenbiz.com/article/>

walmart-suppliers-join-forces-buy-renewable-electricity#:~:text=Under%20Gigaton%20PPA%2C%20a%20diverse,megawatt%2Dhours%20of%20power%20annually.

Jason A Grissom and Lael R Keiser. A supervisor like me: Race, representation, and the satisfaction and turnover decisions of public sector employees. *Journal of Policy Analysis and Management*, 30(3):557–580, 2011.

Kristen Harknett, Daniel Schneider, and Sigrid Luhr. Who cares if parents have unpredictable work schedules?: Just-in-time work schedules and child care arrangements. *Social Problems*, 69(1):164–183, 2022.

Julia R Henly and Susan J Lambert. Unpredictable work timing in retail jobs: Implications for employee work–life conflict. *Ilr Review*, 67(3):986–1016, 2014.

Julia R Henly, H Luke Shaefer, and Elaine Waxman. Nonstandard work schedules: Employer-and employee-driven flexibility in retail jobs. *Social service review*, 80(4):609–634, 2006.

Morela Hernandez, Derek R Avery, Sabrina D Volpone, and Cheryl R Kaiser. Bargaining while black: The role of race in salary negotiations. *Journal of Applied Psychology*, 104(4):581, 2019.

Chengfan Hou, Mengshi Lu, Tianhu Deng, and Zuo-Jun Max Shen. Coordinating project outsourcing through bilateral contract negotiations. *Manufacturing & Service Operations Management*, 23(6):1543–1561, 2021.

Vernon N Hsu, Guoming Lai, Baozhuang Niu, and Wenqiang Xiao. Leader-based collective bargaining: Cooperation mechanism and incentive analysis. *Manufacturing & Service Operations Management*, 19(1):72–83, 2017.

Bin Hu, Izak Duenyas, and Damian R Beil. Does pooling purchases lead to higher profits? *Management Science*, 59(7):1576–1593, 2013.

- Qiaohai Hu, Leroy B Schwarz, and Nelson A Uhan. The impact of group purchasing organizations on healthcare-product supply chains. *Manufacturing & Service Operations Management*, 14(1):7–23, 2012.
- Shanshan Hu, Gilvan C Souza, Mark E Ferguson, and Wenbin Wang. Capacity investment in renewable energy technology with supply intermittency: Data granularity matters! *Manufacturing & Service Operations Management*, 17(4):480–494, 2015.
- Daesik Hur, Vincent A Mabert, and Kurt M Bretthauer. Real-time work schedule adjustment decisions: An investigation and evaluation. *Production and Operations Management*, 13(4):322–339, 2004.
- Intercontinental Exchange (ICE). End-of-day power market data: Pjm western hub (peak, daily). ICE Data Services, 2023. URL <https://www.ice.com/fixed-income-data-services/data-and-analytics/end-of-day-reports>. Daily VWAP from executed bilateral trades; includes daily volume (MWh), trade count, and counterparty count. Author computes trade-volume-weighted annual averages.
- Seyed M Iravani, Mark P Van Oyen, and Katharine T Sims. Structural flexibility: A new perspective on the design of manufacturing and service operations. *Management Science*, 51(2):151–166, 2005.
- Daniel R Jiang and Warren B Powell. Optimal hour-ahead bidding in the real-time electricity market with battery storage using approximate dynamic programming. *INFORMS Journal on Computing*, 27(3):525–543, 2015.
- Per Johansson and Mårten Palme. Assessing the effect of public policy on worker absenteeism. *Journal of Human Resources*, pages 381–409, 2002.
- Janelle Jones and John Schmitt. Trends in job quality for african-american workers, 1979–2011. *The Review of Black Political Economy*, 43(1):1–19, 2016.
- Anne Junor, John O’Brien, and Michael O’Donnell. Welfare wars: public service frontline absenteeism as collective resistance. *Qualitative Research in Accounting & Management*, 6(1/2):26–40, 2009.

- Basak Kalkanci and Erica L Plambeck. Reveal the supplier list? a trade-off in capacity vs. responsibility. *Manufacturing & Service Operations Management*, 22(6):1251–1267, 2020.
- Masoud Kamalahmadi, Qiuping Yu, and Yong-Pin Zhou. Call to duty: Just-in-time scheduling in a restaurant chain. *Management science*, 67(11):6751–6781, 2021.
- Masoud Kamalahmadi, Qiuping Yu, and Yong-Pin Zhou. Racial and gender biases in customer satisfaction surveys: Evidence from a restaurant chain. *University of Miami Business School Research Paper*, (4418420), 2023.
- Bardia Kamrad and Peter Ritchken. Valuing fixed price supply contracts. *European journal of operational research*, 74(1):50–60, 1994.
- Bardia Kamrad and Akhtar Siddique. Supply contracts, profit sharing, switching, and reaction options. *Management Science*, 50(1):64–82, 2004.
- Rachit Kansal. Introduction to the virtual power purchase agreement. *Rocky Mountain Institute*, 2018. URL <https://rmi.org/wp-content/uploads/2018/12/rmi-brc-intro-vppa.pdf>.
- Leonard Kaufman and Peter J. Rousseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis*. John Wiley & Sons, New York, 1990. doi: 10.1002/9780470316801.
- Saravanan Kesavan, Bradley R Staats, and Wendell Gilland. Volume flexibility in services: The costs and benefits of flexible labor resources. *Management Science*, 60(8):1884–1906, 2014.
- Saravanan Kesavan, Susan J Lambert, Joan C Williams, and Pradeep K Pendem. Doing well by doing good: Improving retail store performance with responsible scheduling practices at the gap, inc. *Management Science*, 68(11):7818–7836, 2022.
- Pinar Keskinocak and Seil Savařaneril. Collaborative procurement among competing buyers. *Naval Research Logistics (NRL)*, 55(6):516–540, 2008.

- Javad Khazaei, Michael Coulon, and Warren B Powell. Adapt: A price-stabilizing compliance policy for renewable energy certificates: The case of srec markets. *Operations Research*, 65(6):1429–1445, 2017.
- Jae Ho Kim and Warren B Powell. Optimal energy commitments with storage and intermittent supply. *Operations research*, 59(6):1347–1360, 2011.
- A Gürhan Kök, Kevin Shang, and Şafak Yücel. Impact of electricity pricing policies on renewable energy investments and carbon emissions. *Management Science*, 64(1):131–148, 2018.
- A Gürhan Kök, Kevin Shang, and Şafak Yücel. Investments in renewable and conventional energy: The role of operational flexibility. *Manufacturing & Service Operations Management*, 22(5):925–941, 2020.
- Meera Komarraju, Stephen J Dollinger, and Jennifer L Lovell. Individualism-collectivism in horizontal and vertical directions as predictors of conflict management styles. *International Journal of Conflict Management*, 19(1):20–35, 2008.
- Panos Kouvelis, Xiaole Wu, and Yixuan Xiao. Cash hedging in a supply chain. *Management Science*, 65(8):3928–3947, 2019.
- Tim Kraft, León Valdés, and Yanchong Zheng. Supply chain visibility and social responsibility: Investigating consumers’ behaviors and motives. *Manufacturing & Service Operations Management*, 20(4):617–636, 2018.
- Tali Kristal, Yinon Cohen, and Edo Navot. Benefit inequality among american workers by gender, race, and ethnicity, 1982–2015. *Sociological Science*, 5:461, 2018.
- Jesus Lago, Grzegorz Marcjasz, Bart De Schutter, and Rafał Weron. Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, 293:116983, 2021.
- Guoming Lai, François Margot, and Nicola Secomandi. An approximate dynamic program-

- ming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations research*, 58(3):564–582, 2010.
- Susan J Lambert. Passing the buck: Labor flexibility practices that transfer risk onto hourly workers. *Human relations*, 61(9):1203–1227, 2008.
- Susan J Lambert, Peter J Fugiel, and Julia R Henly. Precarious work schedules among early-career employees in the us: A national snapshot. *Research brief. Chicago: University of Chicago, Employment Instability, Family Well-Being, and Social Policy Network (EINet)*, 2014.
- Susan J Lambert, Julia R Henly, and Jaeseung Kim. Precarious work schedules as a source of economic insecurity and institutional distrust. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 5(4):218–257, 2019.
- Hau L Lee and Christopher S Tang. Socially and environmentally responsible value chain innovations: New operations management research opportunities. *Management Science*, 64(3):983–996, 2018.
- Nils Löhndorf and David Wozabal. The value of coordination in multimarket bidding of grid energy storage. *Operations research*, 71(1):1–22, 2023.
- Nils Löhndorf, David Wozabal, and Stefan Minner. Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research*, 61(4):810–823, 2013.
- Hadas Mandel and Moshe Semyonov. Going back in time? gender differences in trends and sources of the racial pay gap, 1970 to 2010. *American Sociological Review*, 81(5):1039–1068, 2016.
- Bertrand Marianne. New perspectives on gender. In *Handbook of labor economics*, volume 4, pages 1543–1590. Elsevier, 2011.
- Harry M Markowits. Portfolio selection. *Journal of finance*, 7(1):71–91, 1952.

- Harry Markowitz. Mean–variance approximations to expected utility. *European Journal of Operational Research*, 234(2):346–355, 2014.
- Vijay Mehrotra, Ozgür Ozlük, and Robert Saltzman. Intelligent procedures for intra-day updating of call center agent schedules. *Production and operations management*, 19(3):353–367, 2010.
- Luca Mendicino, Daniele Menniti, Anna Pinnarelli, and Nicola Sorrentino. Corporate power purchase agreement: Formulation of the related levelized cost of energy and its application to a real life case study. *Applied Energy*, 253:113577, 2019.
- Robert Christie Mill. Restaurant management: Customers, operations, and employees. (*No Title*), 2007.
- Paul M Muchinsky. Employee absenteeism: A review of the literature. *Journal of Vocational Behavior*, 10(3):316–340, 1977.
- Kurt P Munz, Minah H Jung, and Adam L Alter. Name similarity encourages generosity: A field experiment in email personalization. *Marketing Science*, 39(6):1071–1091, 2020.
- Abhinay Muthoo. *Bargaining theory with applications*. Cambridge University Press, 1999.
- Selvaprabu Nadarajah and Nicola Secomandi. Merchant energy trading in a network. *Operations Research*, 66(5):1304–1320, 2018.
- Mahesh Nagarajan and Yehuda Bassok. A bargaining framework in supply chains: The assembly problem. *Management science*, 54(8):1482–1496, 2008.
- Mahesh Nagarajan and Greys Sošić. Stable farsighted coalitions in competitive markets. *Management Science*, 53(1):29–45, 2007.
- Mahesh Nagarajan and Greys Sošić. Coalition stability in assembly models. *Operations Research*, 57(1):131–145, 2009.
- Mahesh Nagarajan, Greys Sosic, and Hao Zhang. Stable group purchasing organizations. *Marshall School of Business Working Paper No. FBE*, pages 20–10, 2010.

National Health Interview Survey. Work conditions and serious psychological distress among working adults aged 18–64: United states, 2021, 2023. URL https://www.cdc.gov/nchs/products/databriefs/db467.htm#Key_finding. Accessed: 2024-08-08.

National Restaurant Association. Restaurant employee demographics, 2024. URL <https://restaurant.org/getmedia/6f8b55ed-5b3f-40f5-ad04-709ff7ff9f0f/nra-data-brief-restaurant-employee-demographics.pdf>. Accessed: 2024-08-07.

David M Newbery. Reforming competitive electricity markets to meet environmental targets. *Economics of Energy & Environmental Policy*, 1(1):69–82, 2012.

Eric O’Shaughnessy, Jenny Heeter, Chandra Shah, and Sam Koebrich. Corporate acceleration of the renewable energy transition and implications for electric grids. *Renewable and Sustainable Energy Reviews*, 146:111160, 2021.

Eric J O’Shaughnessy, Jenny S Heeter, and Jennifer Sauer. Status and trends in the us voluntary green power market (2017 data). Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2018.

David S Pedulla and Devah Pager. Race and networks in the job search process. *American Sociological Review*, 84(6):983–1012, 2019.

Edmund S Phelps. The statistical theory of racism and sexism. *The american economic review*, 62(4):659–661, 1972.

Evan L Porteus. *Foundations of stochastic inventory theory*. Stanford University Press, 2002.

PV Magazine. WA councils to partner for "australian-first" renewable energy PPA, 2021. URL [https://www.pv-magazine-australia.com/2021/10/19/wa-councils-to-partner-for-australian-first-renewable-energy-ppa/#:~:text=More%20than%2050%20Western%20Australian%20\(WA\)%20local%20governments%20are%20poised,to%20reduce%20their%20overall%20emissions.](https://www.pv-magazine-australia.com/2021/10/19/wa-councils-to-partner-for-australian-first-renewable-energy-ppa/#:~:text=More%20than%2050%20Western%20Australian%20(WA)%20local%20governments%20are%20poised,to%20reduce%20their%20overall%20emissions.)

- RE100. RE100 Annual disclosure report, 2023. URL <https://www.there100.org/our-work/publications/re100-2023-annual-disclosure-report#:~:text=Our%202023%20Annual%20Disclosure%20Report,and%20are%20now%20approaching%20Germany's>.
- Cecilia L Ridgeway. Gender, status, and leadership. *Journal of Social issues*, 57(4):637–655, 2001.
- Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society*, pages 97–109, 1982.
- Daniel F Salas and Warren B Powell. Benchmarking a scalable approximate dynamic programming algorithm for stochastic control of grid-level energy storage. *INFORMS Journal on Computing*, 30(1):106–123, 2018.
- Scott E Sampson and Craig M Froehle. Foundations and implications of a proposed unified services theory. *Production and operations management*, 15(2):329–343, 2006.
- Thomas C Schelling. An essay on bargaining. *The American Economic Review*, 46(3):281–306, 1956.
- Daniel Schneider and Kristen Harknett. How work schedules affect health and wellbeing: The mediating roles of economic insecurity and work-life conflict. In *Meeting of the Population association of America*, 2017.
- Daniel Schneider and Kristen Harknett. Parental exposure to routine work schedule uncertainty and child behavior. 2019a.
- Daniel Schneider and Kristen Harknett. It’s about time: How work schedule instability matters for workers, families, and racial inequality. *The Shift Project Research Brief*, 2019b.
- Daniel Schneider and Kristen Harknett. Hard times: Routine schedule unpredictability and material hardship among service sector workers. *Social Forces*, 99(4):1682–1709, 2021.

- Nicola Secomandi. Optimal commodity trading with a capacitated storage asset. *Management Science*, 56(3):449–467, 2010.
- Nicola Secomandi. On the valuation and monetization roles of the rolling intrinsic policy for merchant commodity storage. *Production and Operations Management*, 34(6):1426–1439, 2025.
- Joachim Seel, Dev Millstein, Andrew D. Mills, Mark Bolinger, and Ryan H. Wiser. Plentiful electricity turns wholesale prices negative. *Advances in Applied Energy*, 4:100073, 2021. doi: 10.1016/j.adapen.2021.100073.
- Scott M Shafer, David A Nembhard, and Mustafa V Uzumeri. The effects of worker learning, forgetting, and heterogeneity on assembly line productivity. *Management Science*, 47(12):1639–1653, 2001.
- Afzal S Siddiqui, Makoto Tanaka, and Yihsu Chen. Are targets for renewable portfolio standards too low? the impact of market structure on energy policy. *European Journal of Operational Research*, 250(1):328–341, 2016.
- Vincent W Slaugh, Alan A Scheller-Wolf, and Sridhar R Tayur. Consistent staffing for long-term care through on-call pools. *Production and Operations Management*, 27(12):2144–2161, 2018.
- Marco Slikker, Jan Cornelis Fransoo, and Marcus Johannes Franciscus Wouters. *Joint ordering in multiple news-vendor problems: a game-theoretical approach*. Technische Universiteit Eindhoven, 2001.
- C Matthew Snipp and Sin Yi Cheung. Changes in racial and gender inequality since 1970. *The ANNALS of the American Academy of Political and Social Science*, 663(1):80–98, 2016.
- Paul E Spector and Suzy Fox. An emotion-centered model of voluntary work behavior: Some parallels between counterproductive work behavior and organizational citizenship behavior. *Human resource management review*, 12(2):269–292, 2002.

- Adam Storer, Daniel Schneider, and Kristen Harknett. What explains racial/ethnic inequality in job quality in the service sector? *American Sociological Review*, 85(4):537–572, 2020.
- Nur Sunar and John R Birge. Strategic commitment to a production schedule with uncertain supply and demand: Renewable energy in day-ahead electricity markets. *Management Science*, 65(2):714–734, 2019.
- Nur Sunar and Erica Plambeck. Allocating emissions among co-products: Implications for procurement and climate policy. *Manufacturing & Service Operations Management*, 18(3):414–428, 2016.
- Nur Sunar and Jayashankar M Swaminathan. Net-metered distributed renewable energy: A peril for utilities? *Management Science*, 67(11):6716–6733, 2021.
- Nur Sunar and Jayashankar M Swaminathan. Socially relevant and inclusive operations management. *Production and Operations Management*, 31(12):4379–4392, 2022.
- Tamarindo. What is an anchor tenant in renewables?, 2019. URL <https://tamarindo.global/articles/what-is-an-anchor-tenant-in-renewables/>.
- Tom Fangyun Tan and Serguei Netessine. When does the devil make work? an empirical study of the impact of workload on worker productivity. *Management Science*, 60(6):1574–1593, 2014.
- Tom Fangyun Tan and Serguei Netessine. When you work with a superman, will you also fly? an empirical study of the impact of coworkers on performance. *Management Science*, 65(8):3495–3517, 2019.
- James Tobin. Liquidity preference as behavior towards risk. *The review of economic studies*, 25(2):65–86, 1958.
- Alessio Trivella, Danial Mohseni-Taheri, and Selvaprabu Nadarajah. Meeting corporate renewable power targets. *Management science*, 69(1):491–512, 2023.

US Bureau of Labor Statistics. Characteristics of minimum wage workers, 2022, 2023. URL <https://www.bls.gov/opub/reports/minimum-wage/2022/home.htm#:~:text=In%202022%2C%2078.7%20million%20workers,all%20wage%20and%20salary%20workers.> Accessed: 2024-08-08.

US Census Bureau. Foodservice and drinking place sales in the u.s. 1992-2022, 2024. URL <https://www.statista.com/statistics/239410/us-food-service-and-drinking-place-sales/>. Accessed: 2024-08-07.

U.S. Energy Information Administration. Form eia-861m (monthly electric power industry report): Detailed data. <https://www.eia.gov/electricity/data/eia861m/>. Accessed: 2025-08-21.

Hao Wang and Baosen Zhang. Energy storage arbitrage in real-time markets via reinforcement learning. In *2018 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2018.

Yulan Wang, Baozhuang Niu, Pengfei Guo, and Jing-Sheng Song. Direct sourcing or agent sourcing? contract negotiation in procurement outsourcing. *Manufacturing & Service Operations Management*, 23(2):294–310, 2021.

JM Wooldridge. *Econometric analysis of cross section and panel data*, 2nd edition, 2010.

Jeffrey M Wooldridge. *Introductory Econometrics: A Modern Approach 6rd ed.* Cengage learning, 2016.

Owen Q Wu and Roman Kapuscinski. Curtailing intermittent generation in electrical systems. *Manufacturing & Service Operations Management*, 15(4):578–595, 2013.

Owen Q Wu, Derek D Wang, and Zhenwei Qin. Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing & Service Operations Management*, 14(3):455–471, 2012.

Junting Ye, Shuchu Han, Yifan Hu, Baris Coskun, Meizhu Liu, Hong Qin, and Steven

- Skiena. Nationality classification using name embeddings. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1897–1906, 2017.
- Ming Yi, Saud Alghumayjan, and Bolun Xu. Perturbed Decision-Focused Learning for Modeling Strategic Energy Storage, June 2024. URL <http://arxiv.org/abs/2406.17085>. arXiv:2406.17085.
- Roberto Zanchi, Mark Porter, and Nicole Miller. The dutch wind consortium, 2017. URL https://www.dsm.com/content/dam/dsm/corporate/en_US/documents/2017-12-07-brc-the-dutch-wind-consortium-case-study.pdf.
- Zeigo. Decarbonizing supply chains through aggregated PPAs, 2023. URL <https://www.zeigo.com/2023/07/17/decarbonizing-supply-chains-through-aggregated-ppas/>.
- Yangfang Zhou, Alan Scheller-Wolf, Nicola Secomandi, and Stephen Smith. Electricity trading and negative prices: Storage vs. disposal. *Management Science*, 62(3):880–898, 2016.
- Yangfang Zhou, Alan Scheller-Wolf, Nicola Secomandi, and Stephen Smith. Managing wind-based electricity generation in the presence of storage and transmission capacity. *Production and Operations Management*, 28(4):970–989, 2019.

Appendix A

SUPPLEMENTARY MATERIAL FOR CHAPTER 2: COLLECTIVE POWER PURCHASE AGREEMENT NEGOTIATION

A.1 *Equilibrium Outcomes Under Base Model*

Let q^* be the common project size characterized in Theorem 2.7. Define the total surplus at q^* by

$$\Delta^{full} \equiv \frac{G_s}{G_1} Q_1 q^* + Q_s q^* - C(q^*) = \frac{G_s}{G_2} Q_2 q^* + Q_s q^* - C(q^*),$$

where the two expressions coincide by Assumption 2.3. Then, under Assumptions 2.2 and 2.3, the equilibrium outcomes for the separate, agent-based, and joint mechanisms are summarized in Table A.1. Parties' equilibrium bargaining gains can be written as the product of a coefficient and Δ^{full} . The latter is the project's net present total profit over its lifetime with the buyers' and the seller's perspectives combined, adjusting for different discount factors. Specifically, $Q_i q^*$ is buyer i 's net present value from the project over the contractual period, and $Q_s q^*$ is the seller's net present value after contract expiration until the physical retirement of the project. For the subsequent analysis, we assume $\Delta^{full} > 0$, implying that the project is perceived as profitable, a prerequisite for the existence of the bargaining equilibrium. In the special case where all parties share the same discount factor, i.e., $G_1 = G_2 = G_s$, the expression for Δ^{full} simplifies to $Q_1 q^* + Q_s q^* - C(q^*)$ or $Q_2 q^* + Q_s q^* - C(q^*)$. Under Assumption 2.3, $G_1 = G_2$ implies $Q_1 = Q_2$, and all discount-related multipliers in Table A.1 become 1. Consequently, differences in the parties' bargaining gains are solely driven by their external bargaining powers.

A.2 *Alternative Bargaining-Power Assumptions*

Assumption 2.8 imposes three simplifying restrictions on the bargaining-power parameters: the buyer-buyer bargaining power is proportional to the buyers' standalone seller-facing bargaining powers, the buyer group's seller-facing bargaining power is additive up to a cap, and

Table A.1: Equilibrium outcomes under different negotiation mechanisms

Outcome	Separate	Agent-Based	Joint
Buyer 1 gain Δ_1	$\frac{(1 - \alpha_{2s}) \alpha_{1s}}{1 - \alpha_{1s} \alpha_{2s}} \frac{G_1}{G_s} \Delta^{full}$	$\alpha_{as} \alpha_{12} \frac{G_1}{G_s} \Delta^{full}$	$\alpha_{bs} \alpha_{12} \frac{G_1}{G_s} \Delta^{full}$
Buyer 2 gain Δ_2	$\frac{(1 - \alpha_{1s}) \alpha_{2s}}{1 - \alpha_{1s} \alpha_{2s}} \frac{G_2}{G_s} \Delta^{full}$	$(1 - \alpha_{12}) \frac{G_2}{G_s} \Delta^{full}$	$\alpha_{bs} (1 - \alpha_{12}) \frac{G_2}{G_s} \Delta^{full}$
Seller gain Δ_s	$\frac{1 - \alpha_{1s} - \alpha_{2s} + \alpha_{1s} \alpha_{2s}}{1 - \alpha_{1s} \alpha_{2s}} \Delta^{full}$	$(1 - \alpha_{as}) \alpha_{12} \Delta^{full}$	$(1 - \alpha_{bs}) \Delta^{full}$
Project size	$q^S = q^A = q^J = q^*, \quad C'(q^*) = \frac{G_s}{G_1} Q_1 + Q_s = \frac{G_s}{G_2} Q_2 + Q_s$		
Project split	Non-unique	Non-unique	$q_1^J = \alpha_{12} q^*, \quad q_2^J = (1 - \alpha_{12}) q^*$
Contract prices	$v_i^S = \frac{1 - \alpha_{is}}{1 - \alpha_{is} \alpha_{js}} \frac{Q_i}{G_i} + \frac{(1 - \alpha_{js}) \alpha_{is}}{1 - \alpha_{is} \alpha_{js}} \left[\frac{C(q^*) - Q_s q^*}{G_s} - \frac{Q_j q_j^S}{G_j} \right] \frac{1}{q_i^S}$	$v_1^A = \alpha_{12} (1 - \alpha_{as}) \frac{Q_1}{G_1} + (1 - \alpha_{12} + \alpha_{as} \alpha_{12}) \frac{C(q^*) - Q_s q^*}{G_s q^*}, \quad v_2^A = \alpha_{12} \frac{Q_2}{G_2} + (1 - \alpha_{12}) \left[\frac{C(q^*) - Q_s q^*}{G_s} - \frac{Q_1 q_1^A}{G_1} \right] \frac{1}{q_2^A}$	$v^J = (1 - \alpha_{bs}) \frac{Q_1}{G_1} + \alpha_{bs} \frac{C(q^*) - Q_s q^*}{G_s q^*}$

Note: Agent-based results are shown for Buyer 1 as the agent with seller-facing bargaining power α_{as} ; the Buyer 2-as-agent case is symmetric and uses the acting agent's seller-facing bargaining power.

the acting agent's seller-facing bargaining power is an interpolation between his standalone power and the buyer group's collective power. This appendix shows how the main buyer-wise viability and collective-mechanism comparison results change when these restrictions are relaxed. The local notation below nests the main-text notation: a_i corresponds to α_{is} , η_i to buyer i 's internal share, $\bar{\alpha}_{bs}$ to the buyer group's seller-facing bargaining power, and $\bar{\alpha}_{as}^{(i)}$ to the acting agent's seller-facing bargaining power.

Let

$$a_i := \alpha_{is}, \quad i \in \{1, 2\},$$

denote buyer i 's standalone bargaining power against the seller. Let

$$\eta_1 := \alpha_{12}, \quad \eta_2 := 1 - \alpha_{12},$$

where η_i is buyer i 's share in the internal buyer-buyer negotiation. Under Assumption 2.8, $\eta_i = a_i/(a_1 + a_2)$, but here we allow η_i to be arbitrary in $(0, 1)$.

Let $\bar{\alpha}_{bs} \in (0, 1]$ denote the buyer group's seller-facing bargaining power under joint negotiation, without imposing

$$\bar{\alpha}_{bs} = \min\{a_1 + a_2, 1\}.$$

Finally, let $\bar{\alpha}_{as}^{(i)} \in [0, 1]$ denote the seller-facing bargaining power of buyer i when buyer i acts as the agent. Unless otherwise stated, we maintain the natural upper bound

$$\bar{\alpha}_{as}^{(i)} \leq \bar{\alpha}_{bs},$$

which says that the acting agent cannot access more seller-facing bargaining power than the buyer group itself.

Because Theorem 2.7 shows that the risk-neutral project size is common across mechanisms, the following comparisons are entirely surplus-allocation comparisons. Define buyer i 's normalized separate-negotiation benchmark coefficient by

$$\kappa_i := \frac{(1 - a_j)a_i}{1 - a_1a_2}, \quad \{i, j\} = \{1, 2\}.$$

Thus,

$$\Delta_i^S = \kappa_i \frac{G_i}{G_s} \Delta^{full}.$$

A.2.1 Buyer-Wise Viability Under General Bargaining Powers

Under joint negotiation, buyer i 's gain is

$$\Delta_i^J = \bar{\alpha}_{bs} \eta_i \frac{G_i}{G_s} \Delta^{full}.$$

Therefore joint negotiation is buyer-wise viable relative to separate negotiation if and only if

$$\bar{\alpha}_{bs} \eta_i \geq \kappa_i, \quad i \in \{1, 2\}.$$

Equivalently,

$$\bar{\alpha}_{bs} \geq \max \left\{ \frac{\kappa_1}{\eta_1}, \frac{\kappa_2}{\eta_2} \right\}.$$

This expression shows directly how relaxing the proportional internal split affects viability. A larger η_1 makes buyer 1's participation constraint easier to satisfy but makes buyer 2's participation constraint harder to satisfy, and vice versa. Similarly, reducing $\bar{\alpha}_{bs}$ below the capped-additive benchmark makes joint negotiation less likely to be buyer-wise viable.

For agent-based negotiation, let $A^{(i)}$ denote the role assignment in which buyer i acts as the agent and buyer j acts as the follower. The agent's gain is

$$\Delta_i^{A^{(i)}} = \bar{\alpha}_{as}^{(i)} \eta_i \frac{G_i}{G_s} \Delta^{full},$$

whereas the follower's gain is

$$\Delta_j^{A^{(i)}} = (1 - \eta_i) \frac{G_j}{G_s} \Delta^{full}.$$

Hence role assignment $A^{(i)}$ is buyer-wise viable if and only if

$$\bar{\alpha}_{as}^{(i)} \eta_i \geq \kappa_i \quad \text{and} \quad 1 - \eta_i \geq \kappa_j.$$

Equivalently,

$$\bar{\alpha}_{as}^{(i)} \geq \frac{\kappa_i}{\eta_i}, \quad \eta_i \leq 1 - \kappa_j.$$

The first condition is the agent's participation constraint: the acting agent must receive enough seller-facing bargaining power to prefer serving as agent rather than negotiating separately. The second condition is the follower's participation constraint: the internal buyer-buyer split must give the follower enough surplus relative to his separate-negotiation outside option.

These generalized conditions also show that the nested-viability implication in the main text depends on Assumption 2.8. Under the maintained assumption, buyer-wise viability of an agent-based role assignment implies buyer-wise viability of joint negotiation. Under fully general bargaining-power parameters, this implication need not hold automatically. In particular, the follower receives $(1 - \eta_i)$ under agent-based negotiation, but only $\bar{\alpha}_{bs}(1 - \eta_i)$ under joint negotiation. Thus, when $\bar{\alpha}_{bs} < 1$, the follower's participation constraint under agent-based negotiation does not by itself imply the follower's participation constraint under joint negotiation.

A.2.2 Agent Role Choice Under General Bargaining Powers

For role assignment $A^{(i)}$, the buyer group's total gain is

$$\Delta_b^{A^{(i)}} = \frac{\Delta^{full}}{G_s} \left[G_i \eta_i \bar{\alpha}_{as}^{(i)} + G_j (1 - \eta_i) \right].$$

Therefore,

$$\Delta_b^{A^{(1)}} - \Delta_b^{A^{(2)}} = \frac{\Delta^{full}}{G_s} \left[G_2 \eta_2 (1 - \bar{\alpha}_{as}^{(2)}) - G_1 \eta_1 (1 - \bar{\alpha}_{as}^{(1)}) \right].$$

Thus buyer 1 should be selected as the agent if and only if

$$G_1 \eta_1 (1 - \bar{\alpha}_{as}^{(1)}) \leq G_2 \eta_2 (1 - \bar{\alpha}_{as}^{(2)}).$$

This generalizes Lemma 2.11. The buyer group should assign the agent role to the buyer whose retained internal share creates the smaller seller-extraction loss.

A.2.3 Joint Versus Agent-Based Negotiation Under General Bargaining Powers

For any viable role assignment $A^{(i)}$, the buyer group's gain under joint negotiation can be written as

$$\Delta_b^J = \frac{\Delta^{full}}{G_s} \bar{\alpha}_{bs} [G_i \eta_i + G_j (1 - \eta_i)].$$

Using the expression for $\Delta_b^{A^{(i)}}$ above,

$$\Delta_b^{A^{(i)}} - \Delta_b^J = \frac{\Delta^{full}}{G_s} \left[G_j (1 - \eta_i) (1 - \bar{\alpha}_{bs}) - G_i \eta_i (\bar{\alpha}_{bs} - \bar{\alpha}_{as}^{(i)}) \right].$$

Therefore, role assignment $A^{(i)}$ gives the buyer group a higher total gain than joint negotiation if and only if

$$G_j (1 - \eta_i) (1 - \bar{\alpha}_{bs}) > G_i \eta_i (\bar{\alpha}_{bs} - \bar{\alpha}_{as}^{(i)}).$$

The left-hand side is the value of shielding the follower's share from the seller-facing negotiation. The right-hand side is the delegation loss from using the agent rather than the full buyer group in the seller-facing negotiation.

This expression nests the main intuition. Agent-based negotiation becomes more attractive when the buyer group's seller-facing bargaining power $\bar{\alpha}_{bs}$ is low, because joint negotiation exposes the entire buyer-side surplus to seller extraction. Joint negotiation becomes more attractive when $\bar{\alpha}_{bs}$ is high, especially when the acting agent cannot fully access

the buyer group's seller-facing bargaining power. In the limiting case $\bar{\alpha}_{bs} = 1$, joint negotiation weakly dominates agent-based negotiation whenever $\bar{\alpha}_{as}^{(i)} \leq 1$, with equality only when $\bar{\alpha}_{as}^{(i)} = 1$.

A.2.4 Connection to Assumption 2.8

Assumption 2.8 is recovered by setting

$$\eta_i = \frac{a_i}{a_1 + a_2}, \quad \bar{\alpha}_{bs} = \min\{a_1 + a_2, 1\},$$

and

$$\bar{\alpha}_{as}^{(i)} = a_i + \lambda(\bar{\alpha}_{bs} - a_i).$$

Under this specialization,

$$\frac{\kappa_i}{\eta_i} = \frac{(1 - a_j)(a_1 + a_2)}{1 - a_1 a_2} = \tau_i,$$

so the generalized joint-viability condition reduces to Proposition 2.9, and the generalized agent-viability condition reduces to Proposition 2.10. The generalized role-choice expression reduces to Lemma 2.11, and the generalized joint-versus-agent comparison reduces to Theorem 2.12.

A.3 Kalai–Smorodinsky Negotiation Framework

In this section we show that, under the alternative Kalai–Smorodinsky (KS) bargaining framework, separate and joint negotiation require the common annualized valuation condition in Assumption 2.3 for project-size decoupling. By contrast, in the agent-based negotiation mechanism, the agent-seller stage is pinned down by the agent buyer's valuation. The KS framework states that a two-party bargaining problem consists of a pair $(\mathcal{F}, \mathbf{d})$ where $\mathcal{F}$ is the feasible agreements set and $\mathbf{d} = (d_1, d_2)$ disagreement points (outside option). Denote F_i^U as the utopia payoff for party $i \in \{1, 2\}$, which is the best value that i can expect to get in the feasible set $\mathcal{F}$. The solution payoff (μ_1, μ_2) to KS bargaining is given by $\frac{\mu_1 - d_1}{F_1^U - d_1} = \frac{\mu_2 - d_2}{F_2^U - d_2}$.

Separate Negotiation. Apply this to the separate negotiation. Consider the bilateral negotiation between buyer i and seller s , with the other buyer ($j \neq i$) fixed. Fix (q_j, v_j) and

write $q = q_i + q_j$. Feasibility requires $\Delta_i \geq 0$, $\Delta_s \geq 0$. For buyer i , F_i^U is achieved when $\Delta_s = 0$, i.e. seller break-even, $\Delta_s = 0 \implies (v_i q_i)_i^U = \frac{C(q) - Q_s q - G_s v_j q_j}{G_s}$. Plug this into Δ_i to derive the corresponding best feasible Δ_i^U , there is $\Delta_i^U = Q_i q_i - G_i \frac{C(q) - Q_s q - G_s v_j q_j}{G_s}$, which is just a function of q_i . Taking derivative w.r.t. q_i (recall that $q = q_i + q_j$ and q_j is fixed) gives FOC

$$\frac{d\Delta_i^U}{dq_i} = Q_i - \frac{G_i}{G_s} (C'(q) - Q_s) = 0 \implies C'(q^*) = Q_s + \frac{G_s}{G_i} Q_i. \quad (\text{A.1})$$

For seller s , F_s^U is achieved when $\Delta_i = 0$, i.e. buyer break-even, $\Delta_i = 0 \implies (v_i q_i)_s^U = \frac{Q_i q_i}{G_i}$. The corresponding seller payoff is $\Delta_s^U(q_i) = G_s \left(\frac{Q_i}{G_i} q_i + v_j q_j \right) + Q_s q - C(q)$. Derivative w.r.t. q_i yields

$$\frac{d\Delta_s^U}{dq_i} = \frac{G_s}{G_i} Q_i + Q_s - C'(q) = 0 \implies C'(q^*) = Q_s + \frac{G_s}{G_i} Q_i. \quad (\text{A.2})$$

Hence in both best feasible scenarios the equilibrium project sizes are the same. The Pareto frontier of $\mathcal{F}$ is thus a linear segment connecting the endpoints $(\Delta_i^U, 0)$ and $(0, \Delta_s^U)$, along which project size is fixed at q^* and therefore the total “pie size” is constant. Since the equilibrium project size condition in (A.1) and (A.2) yields the same implicit equation as in the Nash bargaining framework for both buyers, separate negotiation still requires the common annualized valuation condition in Assumption 2.3.

Agent-based Negotiation. Consider buyer 1 as agent. Fix (v_2, q_2) from stage one and write $q = q_1 + q_2$. In the stage-two bilateral problem between the agent and the seller, feasibility requires $\Delta_1 \geq 0$, $\Delta_s \geq 0$, with disagreement at $(0, 0)$. As in the separate case, we compute each party’s utopia by setting the other party’s payoff to zero. Seller break-even ($\Delta_s = 0$) implies $(v_1 q)_1^U = \frac{C(q) - Q_s q}{G_s}$. Plugging into Δ_1 gives $\Delta_1^U(q_1) = Q_1 q_1 + G_1 v_2 q_2 - \frac{G_1}{G_s} (C(q) - Q_s q)$, which depends on q_1 only through $q = q_1 + q_2$. Differentiating w.r.t. q_1 ,

$$\frac{d\Delta_1^U}{dq_1} = Q_1 - \frac{G_1}{G_s} (C'(q) - Q_s) = 0 \implies C'(q^*) = Q_s + \frac{G_s}{G_1} Q_1. \quad (\text{A.3})$$

Agent break-even ($\Delta_1 = 0$) implies $(v_1 q)_s^U = \frac{Q_1}{G_1} q_1 + v_2 q_2$. Thus $\Delta_s^U(q_1) = G_s \left(\frac{Q_1}{G_1} q_1 + v_2 q_2 \right) + Q_s q - C(q)$, and

$$\frac{d\Delta_s^U}{dq_1} = \frac{G_s}{G_1} Q_1 + Q_s - C'(q) = 0 \implies C'(q^*) = Q_s + \frac{G_s}{G_1} Q_1, \quad (\text{A.4})$$

identical to (A.3). Hence, along the stage-two Pareto frontier the project size is fixed at q^* , and the “pie size” is constant. The FOC (A.3)–(A.4) pins down q^* independently of (v_2, q_2) and independently of Q_2, G_2 . Therefore, under the KS framework, the agent-seller stage does *not* require Assumption 2.3 to determine the equilibrium project size.

Joint Negotiation. Let $\Delta_b := \Delta_1 + \Delta_2 = [(Q_1 z_1 + Q_2 z_2) - (G_1 z_1 + G_2 z_2)v] q$ denote the coalition’s gain. Feasibility in stage two requires $\Delta_b \geq 0$, $\Delta_s \geq 0$, with disagreement at $(0, 0)$. Seller break-even ($\Delta_s = 0$) gives $(vq)_b^U = \frac{C(q) - Q_s q}{G_s}$. Thus the coalition’s utopia payoff is $\Delta_b^U(q) = (Q_1 z_1 + Q_2 z_2) q - \frac{G_1 z_1 + G_2 z_2}{G_s} (C(q) - Q_s q)$. Differentiating w.r.t. q ,

$$\frac{d\Delta_b^U}{dq} = (Q_1 z_1 + Q_2 z_2) - \frac{G_1 z_1 + G_2 z_2}{G_s} (C'(q) - Q_s) = 0 \implies C'(q^*) = Q_s + \frac{G_s (Q_1 z_1 + Q_2 z_2)}{G_1 z_1 + G_2 z_2}. \quad (\text{A.5})$$

Coalition break-even ($\Delta_b = 0$) implies $v^U = \frac{Q_1 z_1 + Q_2 z_2}{G_1 z_1 + G_2 z_2}$. Hence $\Delta_s^U(q) = \left[\frac{G_s (Q_1 z_1 + Q_2 z_2)}{G_1 z_1 + G_2 z_2} + Q_s \right] q - C(q)$, and

$$\frac{d\Delta_s^U}{dq} = \frac{G_s (Q_1 z_1 + Q_2 z_2)}{G_1 z_1 + G_2 z_2} + Q_s - C'(q) = 0, \quad (\text{A.6})$$

which coincides with (A.5). Therefore, along the stage-two Pareto frontier the project size is fixed at q^* , and the pie size is constant. Unlike the agent-based case, the FOC (A.5) depends on (z_1, z_2) through the ratio $\frac{Q_1 z_1 + Q_2 z_2}{G_1 z_1 + G_2 z_2}$. Therefore the stage-one allocation (z_1, z_2) generally *feeds back* into stage two and alters q^* . A clean decoupling, i.e., a q^* that is independent of (z_1, z_2) , occurs if and only if $\frac{Q_1}{G_1} = \frac{Q_2}{G_2}$, as in Assumption 2.3. Hence, under KS, joint negotiation requires Assumption 2.3 to decouple stage-one sharing from the equilibrium project size.

A.4 Equilibrium Outcomes with Risk-Averse Buyers

Define

$$A_i = Q_i + 2 \rho_i (V_{gri} + V_{dpci}), \quad B_i = \rho_i (V_{pi} + V_{ri}),$$

and

$$\tilde{\Delta}^{\text{full}} = \sum_{i=1}^2 \frac{G_s}{G_i} (A_i \tilde{q}_i^{S,A} - B_i (\tilde{q}_i^{S,A})^2) + Q_s \tilde{q}^{S,A} - C(\tilde{q}^{S,A}).$$

Under separate and agent-based negotiation, the equilibrium gains are given in Table A.2, consistent with the profit split in the risk-neutral case reported in Appendix A.1. The

agent-based column reports the Buyer 1-as-agent case. For joint negotiation, a closed-form solution for the equilibrium project split ratio z_1 and the corresponding gains for each party is not attainable. Finally, these results are obtained **without** Assumption 2.3.

Table A.2: Equilibrium outcomes with risk-averse buyers

Outcome	Separate	Agent-Based
Buyer 1 gain $\tilde{\Delta}_1$	$\frac{(1 - \alpha_{2s}) \alpha_{1s}}{1 - \alpha_{1s} \alpha_{2s}} \frac{G_1}{G_s} \tilde{\Delta}^{\text{full}}$	$\alpha_{as} \alpha_{12} \frac{G_1}{G_s} \tilde{\Delta}^{\text{full}}$
Buyer 2 gain $\tilde{\Delta}_2$	$\frac{(1 - \alpha_{1s}) \alpha_{2s}}{1 - \alpha_{1s} \alpha_{2s}} \frac{G_2}{G_s} \tilde{\Delta}^{\text{full}}$	$(1 - \alpha_{12}) \frac{G_2}{G_s} \tilde{\Delta}^{\text{full}}$
Seller gain Δ_s	$\frac{1 - \alpha_{1s} - \alpha_{2s} + \alpha_{1s} \alpha_{2s}}{1 - \alpha_{1s} \alpha_{2s}} \tilde{\Delta}^{\text{full}}$	$(1 - \alpha_{as}) \alpha_{12} \tilde{\Delta}^{\text{full}}$

Note: Results for Agent-based negotiation are for the scenario where Buyer 1 acts as agent with seller-facing bargaining power α_{as} . Buyer 2-as-agent results are symmetric.

A.5 Heuristic REC Banking Policies for Risk-Averse Buyers

Corollary A.1 (Myopic Policy). *For a two-period problem ($T = 2$), the optimal target banking levels are given by:*

$$s_T^* = 0, \quad s_{T-1}^* = \max \left\{ 0, \min \left\{ g_T - q + \frac{\gamma \mathbb{E}[r_T] - r_{T-1}}{2\rho\gamma^{T+1} \text{Var}[r_T]}, [I_{T-1} + q - g_{T-1}]^+ \right\} \right\}.$$

Corollary A.1 indicates that in the terminal period it is always optimal to sell all RECs. In the first period, by contrast, the optimal banking decision is an unconstrained target bounded by the feasible interval $[0, [I_{T-1} + q - g_{T-1}]^+]$. The unconstrained target has two components. The first is $g_T - q$, which reflects the REC deficit in the terminal period. If $g_T < q$, there will be sufficient RECs to meet the target in period T , so the buyer should bank less (a negative level if possible). Conversely, if $g_T > q$, there will be a shortfall in the terminal period, implying a higher banking target in period $T - 1$ (period 1 here). In other words, this component aligns period $T - 1$ banking with the anticipated need to satisfy the renewable target in period T . The second component is an adjustment term,

$\frac{\gamma\mathbb{E}[r_T]-r_{T-1}}{2\rho\gamma^{T+1}\text{Var}[r_T]}$, which reflects intertemporal price considerations and risk preferences. When $\gamma\mathbb{E}[r_T] < r_{T-1}$, REC prices are expected to decline, so the buyer reduces banking; when $\gamma\mathbb{E}[r_T] > r_{T-1}$, prices are expected to rise, so the buyer increases the banking level. The magnitude of this adjustment is scaled by the product $\rho\text{Var}[r_T]$, which represents the risk factor. A lower risk factor (more tolerance or less uncertainty) leads to stronger adjustments based on price forecasts, while a higher risk factor dampens the adjustment, making the buyer more conservative.

Extending the two-period problem to a general T -period setting is analytically intractable. The challenge stems from the truncation of the optimal banking target, which projects the unconstrained solution back into the feasible interval $[0, [I_t + q - g_t]^+]$. This truncation propagates backward, so each period's feasible set depends on whether future targets are capped or floored. This results in a case-dependent, piecewise recursion that prevents a closed-form expression, leaving the solution only characterizable recursively.

Proposition A.2 (Unconstrained Interior Policy Under Risk Aversion). *For a generic T -period problem, assuming that the future decision s_{t+1}^* is known, the unconstrained optimal target inventory policy s_t^{unc} at any period $t < T$ is given by:*

$$s_t^{\text{unc}} = s_{t+1}^* + g_{t+1} - q + \frac{\gamma\mathbb{E}[r_{t+1}] - r_t}{2\rho\gamma^{t+2}\text{Var}[r_{t+1}]}.$$
 (A.7)

Proposition A.2 characterizes the unconstrained optimal banking target in period t , given that the next-period decision s_{t+1}^* is known. Its structure parallels the myopic policy. Specifically, $s_{t+1}^* + g_{t+1}$ represents the total REC requirement in period $t + 1$ (renewable target plus banking target), so $s_{t+1}^* + g_{t+1} - q$ is the amount that must be carried over from period t . The adjustment term, $\frac{\gamma\mathbb{E}[r_{t+1}] - r_t}{2\rho\gamma^{t+2}\text{Var}[r_{t+1}]}$, has the same interpretation as before: it modifies the baseline carryover according to expected intertemporal price differences, scaled by risk aversion and price uncertainty.

Since the problem is analytically intractable, it must be solved numerically. We consider two approaches. The first builds on (A.7); since s_{t+1}^* depends on the unknown r_{t+1} , we can approximate it by replacing s_{t+1}^* with its expectation, averaged over all possible realizations of r_{t+1} . The second applies the myopic policy from Corollary A.1 on a rolling-horizon basis as a heuristic. Embedding this into the contract negotiation Nash bargaining problem yields

no analytical solution and leads to a three-stage hierarchical optimization in which both the mean and variance of REC banking profits must be considered. Solving this problem is numerically challenging. We therefore present the operational banking policy above as an illustrative and tractable component of the RA + banking setting.

A.6 Proof of Statements

For simplicity of display, throughout this section we write $q = q_1 + q_2$.

A.6.1 Proofs of Lemmas 2.4–2.6 and Theorem 2.7

Separate Negotiation

We first solve the Nash product in (2.8), treating (v_1, q_1) as fixed. The first-order condition with respect to v_2 is

$$\frac{-\alpha_{2s}G_2q_2}{Q_2q_2 - G_2v_2q_2} + \frac{(1 - \alpha_{2s})G_sq_2}{G_s(q_1v_1 + q_2v_2) + Q_sq - C(q)} = 0,$$

which gives

$$v_2^S(q_2; q_1, v_1) = (1 - \alpha_{2s})\frac{Q_2}{G_2} + \alpha_{2s}\frac{C(q) - Q_sq - G_sq_1v_1}{G_sq_2}. \quad (\text{A.8})$$

The second derivative of $\log \Omega_{2s}^S$ with respect to v_2 is negative. Substituting (A.8) into the Nash product and applying the envelope theorem with respect to q_2 yields

$$Q_2 + \frac{G_2}{G_s}(Q_s - C'(q)) = 0,$$

and hence

$$C'(q^S) = \frac{G_s}{G_2}Q_2 + Q_s. \quad (\text{A.9})$$

By symmetry, solving (2.7) gives

$$v_1^S(q_1; q_2, v_2) = (1 - \alpha_{1s})\frac{Q_1}{G_1} + \alpha_{1s}\frac{C(q) - Q_sq - G_sq_2v_2}{G_sq_1}, \quad (\text{A.10})$$

and

$$C'(q^S) = \frac{G_s}{G_1}Q_1 + Q_s. \quad (\text{A.11})$$

Under Assumption 2.3, (A.9) and (A.11) impose the same condition on total project size. Therefore separate negotiation pins down q^S but not a unique split (q_1^S, q_2^S) .

Solving the two linear equations (A.8)–(A.10) gives, for $\{i, j\} = \{1, 2\}$,

$$v_i^S = \frac{1 - \alpha_{is}}{1 - \alpha_{is}\alpha_{js}} \frac{Q_i}{G_i} + \frac{(1 - \alpha_{js})\alpha_{is}}{1 - \alpha_{is}\alpha_{js}} \left[\frac{C(q^S) - Q_s q^S}{G_s} - \frac{Q_j q_j^S}{G_j} \right] \frac{1}{q_i^S}.$$

Plugging these prices into the bargaining gains yields

$$\Delta_i^S = \frac{(1 - \alpha_{js})\alpha_{is}}{1 - \alpha_{is}\alpha_{js}} \left[Q_i q^S + \frac{G_i}{G_s} (Q_s q^S - C(q^S)) \right], \quad \{i, j\} = \{1, 2\},$$

and

$$\Delta_s^S = \frac{1 - \alpha_{1s} - \alpha_{2s} + \alpha_{1s}\alpha_{2s}}{1 - \alpha_{1s}\alpha_{2s}} \left[\frac{G_s}{G_1} Q_1 q^S + Q_s q^S - C(q^S) \right].$$

Using the definition of Δ^{full} and $q^S = q^*$ gives Lemma 2.4.

Joint Negotiation

We solve the joint negotiation problem by backward induction. Fix (z_1, z_2) with $z_2 = 1 - z_1$.

The stage-two Nash product depends on the buyer group gain

$$\Delta_1 + \Delta_2 = (Q_1 z_1 + Q_2 z_2 - (G_1 z_1 + G_2 z_2)v)q$$

and the seller gain $\Delta_s = G_s v q + Q_s q - C(q)$. The first-order condition with respect to v gives

$$v^J(q; z_1, z_2) = (1 - \alpha_{bs}) \frac{Q_1 z_1 + Q_2 z_2}{G_1 z_1 + G_2 z_2} + \alpha_{bs} \frac{C(q) - Q_s q}{G_s q}.$$

Substituting this expression and optimizing over q gives

$$C'(q^J) = G_s \frac{Q_1 z_1 + Q_2 z_2}{G_1 z_1 + G_2 z_2} + Q_s.$$

Under Assumption 2.3, the right-hand side reduces to

$$\frac{G_s}{G_1} Q_1 + Q_s = \frac{G_s}{G_2} Q_2 + Q_s,$$

so the stage-two price and project size are independent of the stage-one split. The second-order condition follows from $C''(q) > 0$.

The stage-one problem therefore reduces to maximizing

$$z_1^{\alpha_{12}} (1 - z_1)^{1 - \alpha_{12}},$$

which yields $z_1^J = \alpha_{12}$ and $z_2^J = 1 - \alpha_{12}$. Substituting the resulting price and split into the gains gives

$$\Delta_1^J = \alpha_{bs}\alpha_{12}\frac{G_1}{G_s}\Delta^{full}, \quad \Delta_2^J = \alpha_{bs}(1 - \alpha_{12})\frac{G_2}{G_s}\Delta^{full},$$

and

$$\Delta_s^J = (1 - \alpha_{bs})\Delta^{full}.$$

This proves Lemma 2.5.

Agent-Based Negotiation

Consider the case where buyer 1 is the agent and buyer 2 is the follower. The case with buyer 2 as agent is symmetric. In stage two, (v_2, q_2) are fixed and buyer 1 bargains with the seller. The first-order condition with respect to v_1 gives

$$v_1^A(q_1; q_2, v_2) = (1 - \alpha_{as})\frac{Q_1q_1 + G_1v_2q_2}{G_1q} + \alpha_{as}\frac{C(q) - Q_sq}{G_sq}. \quad (\text{A.12})$$

Substituting (A.12) into the stage-two Nash product and optimizing over q_1 gives

$$C'(q^A) = \frac{G_s}{G_1}Q_1 + Q_s. \quad (\text{A.13})$$

The second-order condition follows from $C''(q) > 0$.

Next solve stage one. Since (A.13) fixes total project size, write $q_1 = q^A - q_2$. Substituting (A.12) into the buyer-buyer Nash product, the first-order condition with respect to v_2 yields

$$v_2^A(q_2) = \alpha_{12}\frac{Q_2}{G_2} + (1 - \alpha_{12})\left[\frac{C(q^A) - Q_sq^A}{G_sq_2} - \frac{Q_1(q^A - q_2)}{G_1q_2}\right].$$

Substituting this expression into the stage-one Nash product and optimizing over q_2 imposes

$$\frac{Q_1}{G_1} = \frac{Q_2}{G_2},$$

which holds by Assumption 2.3. Equivalently, stage one is consistent with

$$C'(q^A) = \frac{G_s}{G_2}Q_2 + Q_s.$$

Thus the total project size is pinned down, but the split (q_1^A, q_2^A) is not unique.

The equilibrium gains are unique. Substituting v_1^A and v_2^A into the gain functions gives

$$\Delta_2^A = (1 - \alpha_{12}) \left[Q_2 q^A + \frac{G_2}{G_s} (Q_s q^A - C(q^A)) \right],$$

$$\Delta_1^A = \alpha_{as} \alpha_{12} \left[Q_1 q^A + \frac{G_1}{G_s} (Q_s q^A - C(q^A)) \right],$$

and

$$\Delta_s^A = (1 - \alpha_{as}) \alpha_{12} \left[\frac{G_s}{G_1} Q_1 q^A + Q_s q^A - C(q^A) \right].$$

Using $q^A = q^*$ and the definition of Δ^{full} proves Lemma 2.6.

Common Project Size

The three preceding derivations show that separate, joint, and agent-based negotiation all imply

$$C'(q) = \frac{G_s}{G_1} Q_1 + Q_s = \frac{G_s}{G_2} Q_2 + Q_s.$$

By Assumption 2.2, $C'(q)$ is strictly increasing, so this equation has at most one solution.

Denote that solution by q^* . Hence

$$q^S = q^J = q^A = q^*,$$

which proves Theorem 2.7.

A.6.2 Proof of Proposition 2.9

Under Assumption 2.8,

$$\alpha_{12} = \frac{\alpha_{1s}}{\alpha_{1s} + \alpha_{2s}}, \quad 1 - \alpha_{12} = \frac{\alpha_{2s}}{\alpha_{1s} + \alpha_{2s}}.$$

Using Lemmas 2.4 and 2.5, for $\{i, j\} = \{1, 2\}$,

$$\Delta_i^J - \Delta_i^S = \frac{G_i}{G_s} \Delta^{full} \frac{\alpha_{is}}{\alpha_{1s} + \alpha_{2s}} (\alpha_{bs} - \tau_i),$$

where τ_i is defined in (2.15). Since $G_i > 0$ and $\Delta^{full} > 0$, buyer i weakly prefers joint negotiation to separate negotiation if and only if $\alpha_{bs} \geq \tau_i$.

If $\alpha_{1s} + \alpha_{2s} \leq 1$, then $\alpha_{bs} = \alpha_{1s} + \alpha_{2s}$. For each i , $\alpha_{bs} \geq \tau_i$ reduces to

$$1 - \alpha_{1s} \alpha_{2s} \geq 1 - \alpha_{js},$$

which holds because $0 < \alpha_{is} < 1$. Thus joint negotiation is always buyer-wise viable in the uncapped region.

If $\alpha_{1s} + \alpha_{2s} > 1$, then $\alpha_{bs} = 1$. The condition $1 \geq \tau_1$ is equivalent to

$$\alpha_{1s} + \alpha_{2s} \leq 1 + \alpha_{2s}^2,$$

and $1 \geq \tau_2$ is equivalent to

$$\alpha_{1s} + \alpha_{2s} \leq 1 + \alpha_{1s}^2.$$

Both buyer participation constraints hold if and only if

$$\alpha_{1s} + \alpha_{2s} \leq 1 + \min\{\alpha_{1s}^2, \alpha_{2s}^2\}.$$

This proves Proposition 2.9.

A.6.3 Proof of Proposition 2.10

Let buyer i be the agent and buyer j the follower, with $\{i, j\} = \{1, 2\}$. By Lemma 2.6 and Assumption 2.8, the agent's gain under role assignment $A^{(i)}$ is

$$\Delta_i^{A^{(i)}} = \frac{G_i}{G_s} \Delta^{full} \frac{\alpha_{is}}{\alpha_{1s} + \alpha_{2s}} \alpha_{as}^{(i)}(\lambda).$$

Comparing with Lemma 2.4,

$$\Delta_i^{A^{(i)}} - \Delta_i^S = \frac{G_i}{G_s} \Delta^{full} \frac{\alpha_{is}}{\alpha_{1s} + \alpha_{2s}} \left(\alpha_{as}^{(i)}(\lambda) - \tau_i \right).$$

Thus the agent participates if and only if $\alpha_{as}^{(i)}(\lambda) \geq \tau_i$.

The follower's gain under $A^{(i)}$ is

$$\Delta_j^{A^{(i)}} = \frac{G_j}{G_s} \Delta^{full} \frac{\alpha_{js}}{\alpha_{1s} + \alpha_{2s}}.$$

Therefore

$$\Delta_j^{A^{(i)}} - \Delta_j^S = \frac{G_j}{G_s} \Delta^{full} \frac{\alpha_{js}(1 - \alpha_{1s} - \alpha_{2s} + \alpha_{is}^2)}{(\alpha_{1s} + \alpha_{2s})(1 - \alpha_{1s}\alpha_{2s})}.$$

Since the denominator is positive, the follower participates if and only if

$$\alpha_{1s} + \alpha_{2s} \leq 1 + \alpha_{is}^2.$$

Combining the agent and follower participation constraints proves Proposition 2.10.

Nested viability. We also note that buyer-wise viability of an agent-based role assignment implies buyer-wise viability of joint negotiation. If $A^{(i)}$ is viable, then the agent's participation constraint gives

$$\alpha_{as}^{(i)}(\lambda) \geq \tau_i.$$

Because Assumption 2.8 implies

$$\alpha_{as}^{(i)}(\lambda) \leq \alpha_{bs},$$

joint negotiation satisfies buyer i 's participation constraint. The follower constraint under $A^{(i)}$ is

$$\alpha_{1s} + \alpha_{2s} \leq 1 + \alpha_{is}^2,$$

which is exactly the corresponding joint-participation condition for buyer j in the capped region; in the uncapped region joint negotiation is always viable by Proposition 2.9. Hence whenever an agent-based role assignment is buyer-wise viable, joint negotiation is also buyer-wise viable.

A.6.4 Proof of Lemma 2.11

For role assignment $A^{(i)}$, the buyer group's total gain is

$$\Delta_b^{A^{(i)}} = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \left[G_i \alpha_{is} \alpha_{as}^{(i)}(\lambda) + G_j \alpha_{js} \right].$$

Equivalently,

$$\Delta_b^{A^{(i)}} = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \left[G_1 \alpha_{1s} + G_2 \alpha_{2s} - G_i \alpha_{is} (1 - \alpha_{as}^{(i)}(\lambda)) \right].$$

Subtracting the buyer-group gain under $A^{(2)}$ from that under $A^{(1)}$ gives

$$\Delta_b^{A^{(1)}} - \Delta_b^{A^{(2)}} = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \left[G_2 \alpha_{2s} (1 - \alpha_{as}^{(2)}(\lambda)) - G_1 \alpha_{1s} (1 - \alpha_{as}^{(1)}(\lambda)) \right].$$

Thus buyer 1 should be selected as the agent exactly when this difference is nonnegative.

Using Assumption 2.8, if $\alpha_{1s} + \alpha_{2s} \leq 1$, then

$$\alpha_{as}^{(1)}(\lambda) = \alpha_{1s} + \lambda \alpha_{2s}, \quad \alpha_{as}^{(2)}(\lambda) = \alpha_{2s} + \lambda \alpha_{1s}.$$

If $\alpha_{1s} + \alpha_{2s} > 1$, then

$$\alpha_{as}^{(i)}(\lambda) = \alpha_{is} + \lambda(1 - \alpha_{is}).$$

Substituting these expressions into the previous difference yields the two cases stated in Lemma 2.11.

A.6.5 Proof of Theorem 2.12

For any viable role assignment $A^{(i)}$, compare the buyer-group gain under agent-based negotiation with that under joint negotiation. From Lemmas 2.5 and 2.6,

$$\Delta_b^J = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \alpha_{bs} (G_1 \alpha_{1s} + G_2 \alpha_{2s}),$$

and

$$\Delta_b^{A^{(i)}} = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \left(G_i \alpha_{is} \alpha_{as}^{(i)}(\lambda) + G_j \alpha_{js} \right).$$

Subtracting gives

$$\Delta_b^{A^{(i)}} - \Delta_b^J = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \left[G_j \alpha_{js} (1 - \alpha_{bs}) - G_i \alpha_{is} (\alpha_{bs} - \alpha_{as}^{(i)}(\lambda)) \right].$$

If $\alpha_{1s} + \alpha_{2s} \leq 1$, then $\alpha_{bs} = \alpha_{1s} + \alpha_{2s}$ and $\alpha_{bs} - \alpha_{as}^{(i)}(\lambda) = (1 - \lambda) \alpha_{js}$. Hence

$$\Delta_b^{A^{(i)}} - \Delta_b^J = \frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} \alpha_{js} [G_j (1 - \alpha_{1s} - \alpha_{2s}) - G_i \alpha_{is} (1 - \lambda)].$$

If $\alpha_{1s} + \alpha_{2s} > 1$, then $\alpha_{bs} = 1$ and $\alpha_{bs} - \alpha_{as}^{(i)}(\lambda) = (1 - \lambda)(1 - \alpha_{is})$, so

$$\Delta_b^{A^{(i)}} - \Delta_b^J = -\frac{\Delta^{full}}{G_s(\alpha_{1s} + \alpha_{2s})} G_i \alpha_{is} (1 - \lambda)(1 - \alpha_{is}).$$

These are exactly the two cases in Theorem 2.12.

A.6.6 Proof of Proposition 2.14 and Table A.2

Define

$$A_i = Q_i + 2\rho_i(V_{gri} + V_{dpci}), \quad B_i = \rho_i(V_{pi} + V_{ri}), \quad i \in \{1, 2\}.$$

Then buyer i 's risk-adjusted gain is

$$\tilde{\Delta}_i = (A_i - G_i v_i) q_i - B_i q_i^2,$$

with the agent-based modification in (2.23) when buyer i acts as the agent.

Separate negotiation. Consider buyer 2's bilateral Nash problem, treating (v_1, q_1) as fixed. The first-order condition with respect to v_2 gives

$$\frac{-\alpha_{2s}G_2q_2}{\tilde{\Delta}_2} + \frac{(1-\alpha_{2s})G_sq_2}{\Delta_s} = 0 \implies \frac{\tilde{\Delta}_2}{\Delta_s} = \frac{\alpha_{2s}G_2}{(1-\alpha_{2s})G_s}, \quad (\text{A.14})$$

and

$$v_2^S(q_2; q_1, v_1) = (1-\alpha_{2s})\frac{A_2 - B_2q_2}{G_2} + \alpha_{2s}\frac{C(q) - Q_sq - G_sv_1q_1}{G_sq_2}.$$

The second-order condition with respect to v_2 is negative. Substituting this price and applying the envelope theorem to q_2 gives

$$A_2 - 2B_2q_2 + \frac{G_2}{G_s}[Q_s - C'(q^S)] = 0. \quad (\text{A.15})$$

Repeating the same steps for buyer 1 gives

$$A_1 - 2B_1q_1 + \frac{G_1}{G_s}[Q_s - C'(q^S)] = 0. \quad (\text{A.16})$$

Equations (A.15)–(A.16) imply

$$C'(q^S) = Q_s + \frac{G_s}{G_1}(A_1 - 2B_1q_1) = Q_s + \frac{G_s}{G_2}(A_2 - 2B_2q_2).$$

Using $q_1 = q - q_2$ and solving for q_2 gives

$$q_2 = \frac{G_1A_2 - G_2A_1 + 2G_2B_1q}{2(G_2B_1 + G_1B_2)}.$$

Substituting this expression back into the first equality yields

$$C'(q^S) = Q_s + \frac{G_s(A_1B_2 + A_2B_1 - 2B_1B_2q^S)}{G_2B_1 + G_1B_2}.$$

The same gain-splitting algebra as in the risk-neutral proof gives

$$\tilde{\Delta}_1^S = \frac{(1-\alpha_{2s})\alpha_{1s}}{1-\alpha_{1s}\alpha_{2s}} \frac{G_1}{G_s} \tilde{\Delta}^{\text{full}}, \quad \tilde{\Delta}_2^S = \frac{(1-\alpha_{1s})\alpha_{2s}}{1-\alpha_{1s}\alpha_{2s}} \frac{G_2}{G_s} \tilde{\Delta}^{\text{full}},$$

and

$$\Delta_s^S = \frac{1-\alpha_{1s}-\alpha_{2s}+\alpha_{1s}\alpha_{2s}}{1-\alpha_{1s}\alpha_{2s}} \tilde{\Delta}^{\text{full}}.$$

Agent-based negotiation. Take buyer 1 as the agent. The opposite role assignment is symmetric. The stage-two problem gives

$$v_1^A = (1-\alpha_{as})\frac{A_1q_1 - B_1q_1^2 + G_1v_2q_2}{G_1q} + \alpha_{as}\frac{C(q) - Q_sq}{G_sq},$$

and the envelope condition gives

$$q_1^A = \frac{A_1 + \frac{G_1}{G_s} [Q_s - C'(q^A)]}{2B_1}.$$

In stage one, substituting the stage-two response and applying the first-order conditions gives the same condition for buyer 2 as in separate negotiation,

$$q_2^A = \frac{A_2 + \frac{G_2}{G_s} [Q_s - C'(q^A)]}{2B_2}.$$

Combining the two project-share equations yields the same total-size condition as separate negotiation:

$$C'(q^A) = Q_s + \frac{G_s (A_1 B_2 + A_2 B_1 - 2B_1 B_2 q^A)}{G_2 B_1 + G_1 B_2}.$$

The equilibrium gains are

$$\tilde{\Delta}_1^A = \alpha_{as} \alpha_{12} \frac{G_1}{G_s} \tilde{\Delta}^{\text{full}}, \quad \tilde{\Delta}_2^A = (1 - \alpha_{12}) \frac{G_2}{G_s} \tilde{\Delta}^{\text{full}},$$

and

$$\Delta_s^A = (1 - \alpha_{as}) \alpha_{12} \tilde{\Delta}^{\text{full}}.$$

These coefficient formulas are the risk-averse counterparts of the base-model splits reported in Table A.2, with α_{as} denoting the acting agent's seller-facing bargaining power.

Joint negotiation. For fixed z_1 and $z_2 = 1 - z_1$, define

$$\bar{A}(z_1) = z_1 A_1 + (1 - z_1) A_2, \quad \bar{B}(z_1) = z_1^2 B_1 + (1 - z_1)^2 B_2, \quad \bar{G}(z_1) = z_1 G_1 + (1 - z_1) G_2.$$

The aggregate buyer gain in stage two is

$$\tilde{\Delta}_{12} = [\bar{A}(z_1) - \bar{G}(z_1)v]q - \bar{B}(z_1)q^2,$$

and the seller gain is $\Delta_s = G_s v q + Q_s q - C(q)$. The first-order conditions for

$$\alpha_{bs} \log \tilde{\Delta}_{12} + (1 - \alpha_{bs}) \log \Delta_s$$

imply

$$\frac{\Delta_s}{\tilde{\Delta}_{12}} = \frac{1 - \alpha_{bs}}{\alpha_{bs}} \frac{G_s}{\bar{G}(z_1)}.$$

Substituting this ratio into the first-order condition with respect to q gives

$$C'(q^J) = Q_s + \frac{G_s [\bar{A}(z_1) - 2\bar{B}(z_1)q^J]}{\bar{G}(z_1)},$$

or equivalently

$$C'(q^J) = Q_s + \frac{G_s [z_1 A_1 + (1 - z_1)A_2 - 2(z_1^2 B_1 + (1 - z_1)^2 B_2)q^J]}{z_1 G_1 + (1 - z_1)G_2}.$$

This proves Proposition 2.14. A closed-form expression for the stage-one optimizer z_1 is generally unavailable.

A.6.7 Proof of Proposition 2.15

Recall the risk-averse first-order condition under separate and agent-based negotiation:

$$C'(\tilde{q}^{S,A}) = Q_s + \frac{G_s (A_1 B_2 + A_2 B_1 - 2B_1 B_2 \tilde{q}^{S,A})}{G_2 B_1 + G_1 B_2}.$$

Let

$$h(q) = \frac{G_s (A_1 B_2 + A_2 B_1 - 2B_1 B_2 q)}{G_2 B_1 + G_1 B_2}.$$

Since $B_1, B_2, G_s > 0$, $h(q)$ is strictly decreasing. The risk-neutral project size satisfies

$$C'(q^*) = Q_s + \frac{G_s}{G_1} Q_1,$$

and q^+ is defined so that

$$h(q^+) = \frac{G_s}{G_1} Q_1.$$

Solving this equality gives the expression for q^+ in the proposition. Since $C'(q) - Q_s$ is strictly increasing and $h(q)$ is strictly decreasing, the crossing point of $C'(q) - Q_s$ with $h(q)$ lies to the right of q^* if and only if $h(q^*) > h(q^+)$, which is equivalent to $q^+ > q^*$. Thus $\tilde{q}^{S,A} > q^*$ if $q^+ > q^*$, and the reverse argument gives $\tilde{q}^{S,A} < q^*$ if $q^+ < q^*$.

For joint negotiation, fix $z_1 \in [0, 1]$ and define

$$h_J(q; z_1) = \frac{G_s [z_1 A_1 + (1 - z_1)A_2 - 2B_J(z_1)q]}{D_J},$$

where

$$D_J = z_1 G_1 + (1 - z_1)G_2, \quad B_J(z_1) = z_1^2 B_1 + (1 - z_1)^2 B_2.$$

Again, $h_J(q; z_1)$ is strictly decreasing in q , while $C'(q) - Q_s$ is strictly increasing. The quantity $q_J^+(z_1)$ is defined by

$$h_J(q_J^+(z_1); z_1) = \frac{G_s}{G_1} Q_1,$$

which gives the expression in the proposition. The same monotonicity argument implies

$$\tilde{q}^J > q^* \quad \text{if } q_J^+(z_1) > q^*, \quad \tilde{q}^J < q^* \quad \text{if } q_J^+(z_1) < q^*.$$

This proves Proposition 2.15.

A.6.8 Proof of Proposition 2.16

The result follows by backward induction. In period T , there is no continuation value. The buyer solves

$$J_T(I_T) = \max_{0 \leq s_T \leq [I_T + q - g_T]^+} r_T(I_T - s_T),$$

so $s_T^* = 0$ and $J_T(I_T) = r_T I_T$, which is convex and piecewise linear.

Assume $J_{t+1}(I_{t+1})$ is convex and piecewise linear in I_{t+1} . In period t ,

$$J_t(I_t) = \max_{0 \leq s_t \leq [I_t + q - g_t]^+} \{r_t(I_t - s_t) + \gamma \mathbb{E}[J_{t+1}(s_t)]\}.$$

The conditional expectation of a convex piecewise-linear function is still convex and piecewise linear. Adding the linear term $r_t(I_t - s_t)$ preserves piecewise linearity in s_t . Therefore the objective is convex in the choice variable, so an optimum over the interval occurs at an endpoint. The buyer either sells all available RECs, $s_t^* = 0$, or banks the entire feasible amount, $s_t^* = [I_t + q - g_t]^+$.

The endpoint comparison is governed by the current price r_t and the discounted expected marginal continuation value. Thus

$$s_t^* = 0 \quad \text{if } r_t \geq \gamma \frac{d \mathbb{E}[J_{t+1}(I_{t+1})]}{d I_{t+1}},$$

and

$$s_t^* = [I_t + q - g_t]^+ \quad \text{if } r_t < \gamma \frac{d \mathbb{E}[J_{t+1}(I_{t+1})]}{d I_{t+1}}.$$

Serial correlation changes the conditional expectation in the continuation value but not this endpoint logic.

Finally, increasing q weakly expands the feasible upper bound $[I_t + q - g_t]^+$, while increasing g_t weakly contracts it. Hence J_t is non-decreasing in q and non-increasing in g_t . This proves Proposition 2.16.

A.6.9 Proof of Corollary 2.17

With REC banking, buyer i 's bargaining gain includes the additional term $\hat{J}_i(q_i)$. In separate negotiation, the same first-order argument as in the base model gives

$$C'(q^*) = \frac{G_s}{G_i}(Q_i + \hat{J}'_i(q_i^*)) + Q_s, \quad i \in \{1, 2\}.$$

Thus the equilibrium under separate negotiation satisfies

$$C'(q^*) = \frac{G_s}{G_1}(Q_1 + \hat{J}'_1(q_1^*)) + Q_s = \frac{G_s}{G_2}(Q_2 + \hat{J}'_2(q_2^*)) + Q_s.$$

The agent-based derivation is identical after accounting for the agent's pass-through payment, so it yields the same first-order characterization.

Under joint negotiation, with $q_i = z_i q$, the buyer group's marginal value is the weighted average of the two buyers' marginal values. The stage-two first-order condition becomes

$$C'(q^*) = G_s \frac{(Q_1 + \hat{J}'_1(q_1^*))z_1^* + (Q_2 + \hat{J}'_2(q_2^*))z_2^*}{G_1 z_1^* + G_2 z_2^*} + Q_s.$$

Let q^0 denote the no-banking equilibrium project size and let q^B denote the equilibrium project size when REC banking is allowed. By Proposition 2.16, $\hat{J}_i$ is non-decreasing in project size, so the banking marginal-value terms in the FOCs are nonnegative. Therefore, at q^0 , the right-hand side of the banking FOC is weakly larger than the no-banking marginal value. Since $C'(q)$ is strictly increasing by Assumption 2.2, the banking equilibrium satisfies $q^B \geq q^0$, with strict inequality whenever the relevant banking subgradient is strictly positive at q^0 . This proves Corollary 2.17.

A.6.10 Proofs of Corollary A.1 and Proposition A.2

In the terminal period, there is no continuation value. The buyer solves

$$\max_{0 \leq s_T \leq [I_T + q - g_T]^+} \gamma^T [m_T + r_T(I_T - s_T)].$$

Since $r_T \geq 0$, the objective is weakly decreasing in s_T , so $s_T^* = 0$.

Now consider period $t < T$ and suppose the next-period target s_{t+1}^* is known. Terms from period $t + 2$ onward do not depend on s_t , so the relevant objective is

$$\begin{aligned} \max_{0 \leq s_t \leq [I_t + q - g_t]^+} & \gamma^t [m_t + r_t(I_t - s_t)] + \gamma^{t+1} \mathbb{E}[m_{t+1} + r_{t+1}(s_t - s_{t+1}^*)] \\ & - \rho \gamma^{2(t+1)} \text{Var}[r_{t+1}] [(s_t - s_{t+1}^*)^2 - 2(g_{t+1} - q)(s_t - s_{t+1}^*) + (g_{t+1} - q)^2]. \end{aligned}$$

The first-order condition is

$$-\gamma^t r_t + \gamma^{t+1} \mathbb{E}[r_{t+1}] - 2\rho \gamma^{2(t+1)} \text{Var}[r_{t+1}] [s_t - s_{t+1}^* + (q - g_{t+1})] = 0.$$

Solving gives the unconstrained target

$$s_t^{\text{unc}} = s_{t+1}^* + g_{t+1} - q + \frac{\gamma \mathbb{E}[r_{t+1}] - r_t}{2\rho \gamma^{t+2} \text{Var}[r_{t+1}]}.$$

The expression above is the unconstrained interior policy in Proposition A.2. When the feasibility bounds are imposed, the corresponding feasible target is obtained by projecting the unconstrained target onto the feasible interval:

$$s_t^* = \max \{0, \min \{s_t^{\text{unc}}, [I_t + q - g_t]^+\} \}.$$

This proves Proposition A.2 and gives the constrained policy used in the two-period case. Specifically, set $t = T - 1$ and use $s_T^* = 0$. The unconstrained target becomes

$$s_{T-1}^{\text{unc}} = g_T - q + \frac{\gamma \mathbb{E}[r_T] - r_{T-1}}{2\rho \gamma^{T+1} \text{Var}[r_T]}.$$

Projecting this target onto the feasible interval $[0, [I_{T-1} + q - g_{T-1}]^+]$ gives

$$s_{T-1}^* = \max \left\{ 0, \min \left\{ g_T - q + \frac{\gamma \mathbb{E}[r_T] - r_{T-1}}{2\rho \gamma^{T+1} \text{Var}[r_T]}, [I_{T-1} + q - g_{T-1}]^+ \right\} \right\}.$$

Together with $s_T^* = 0$, this proves Corollary A.1.

A.7 Numerical Study

This section uses industry data to examine the negotiation equilibrium across three settings. We refer to the base model in Section 2.3 as the *RN No-Banking* setting and to the

two extensions as *RA No-Banking* (Section 2.4.1) and *RN Banking* (Section 2.5.1). The numerical experiments ask whether the base-model conclusions on project size, viability, and collective-mechanism preference continue to hold when buyers are risk-averse or when RECs can be banked.

A.7.1 Data and Parameters

We assume a 22-year operational life ($L = 22$) and a 10-year PPA ($T = 10$). The baseline annual discount factor is 0.95. Output q is measured in MWh per year, and costs follow $C(q) = \frac{1}{2}q^2$. To derive Q_1 , Q_2 , and Q_s , we use daily weighted-average electricity prices at the PJM Hub from January 1, 2001, to December 31, 2022, obtained from Intercontinental Exchange (ICE) [2023]. Assuming annual PPA payments, we aggregate the daily series to annual averages and use them as expected annual prices. Denote the average price in year t by p_t .

For REC prices, we use the national retail REC price trajectory reported for 2012–2018 by O’Shaughnessy et al. [2018]. Because U.S. REC markets are fragmented and lack a centralized public database, this series serves as an indicative benchmark for REC price dynamics. We extrapolate a 22-period trajectory by repeating the observed trend. For robustness, we also test alternative trajectories, including linear, geometric, and monotonic increasing or decreasing paths; our qualitative conclusions are unchanged.

The final data input is the buyers’ local retail price c_{it} . This term cancels out in the RN settings but matters in the RA setting because it may correlate with the market price. We use the average U.S. retail electricity price to end users from 2001 to 2022 [U.S. Energy Information Administration]. All three price streams are plotted in Figure A.1.

To capture price-forecast risk, we detrend historical annual average prices and use the residual volatility as the baseline perceived annual price volatility. Table A.3 summarizes these inputs. For the risk-aversion parameters ρ_1 and ρ_2 , we test values in the range 10^{-4} to 10^{-3} . For the exogenous seller-facing bargaining powers, we test $\alpha_{1s}, \alpha_{2s} \in \{0.1, 0.2, \dots, 0.8\}$.

In sum, the inputs for market value, cost, and buyers’ perceived market risk are grounded




Figure A.1: Electricity and REC Price Series

Table A.3: Annual Price Series Summary Statistics

Series	Mean	Std	Detrended Std
PJM (\$/MWh)	50.63	16.52	16.26
REC (\$/MWh)	0.76	0.27	0.26
Retail (\$/MWh)	95.98	13.28	4.62
Correlation (PJM vs. Retail): 0.067			

in empirical price series. For parameters such as risk aversion and bargaining power, we test a broad grid to cover plausible scenario variation.

A.7.2 Analysis and Results

For the RN No-Banking setting, the closed-form solutions give a direct benchmark. We also solve the corresponding optimization problems numerically using SciPy's `minimize` routine; in all tested instances, the numerical solutions match the analytical results.

For the RA No-Banking setting, the numerical approach follows the stage structure of the model. For joint negotiation, where closed-form buyer gains are unavailable, the outer

layer searches over candidate project split ratios $z_1 \in [0, 1]$ and $z_2 = 1 - z_1$. For each split, the inner layer solves the stage-two bargaining problem for the equilibrium contract price and project size. For separate and agent-based negotiation, we use the closed-form characterizations in Proposition 2.14 and validate them against the corresponding numerical programs.

Under RN Banking, REC banking adds the dynamic program introduced in Section 2.5.1. The dynamic program admits a bang–bang optimal policy, so for any fixed project size the associated expected banking profit can be evaluated efficiently. However, nesting this dynamic program inside the negotiation problem creates a non-smooth mapping from project size to payoffs. To improve numerical stability, we pre-solve the REC-banking problem on a dense grid of project sizes for each buyer and fit a quadratic response surface to the resulting payoff function. This smooth surrogate is then embedded in the negotiation problem.




Figure A.2: Equilibrium Project Size Difference Relative to RN No-Banking Setting

Note: μ denotes the sample mean percentage difference in equilibrium project size relative to the RN No-Banking benchmark.


Figure A.2 reports the percentage difference in equilibrium project size relative to the RN No-Banking benchmark. Under RA No-Banking, each box summarizes the distribution

across the grid of risk-aversion pairs (ρ_1, ρ_2) . Consistent with Proposition 2.15, separate and agent-based negotiation yield the same project-size outcome, while joint negotiation can differ because the joint mechanism aggregates both buyers' risk-adjusted marginal values. Under the empirical parameters, all tested RA No-Banking instances produce smaller project sizes than the RN No-Banking benchmark. This occurs because REC prices are substantially lower than electricity prices and the correlation between market electricity prices and local retail costs is weak, so the hedging terms in A_i are not strong enough to offset the risk-exposure terms in B_i . In a high-hedging calibration, with the standard deviation of local retail price increased to \$25/MWh and the correlation between market price and local cost set to one, project size increases by 4–19% relative to the RN No-Banking benchmark, consistent with Proposition 2.15.


Under RN Banking, all three mechanisms generally yield slightly larger project sizes than in RN No-Banking, consistent with Corollary 2.17. The increase is small under the baseline parameters because the REC price level is low. Joint negotiation produces marginally larger project sizes than separate and agent-based negotiation because it internalizes both buyers' banking benefits in a single buyer-side bargaining problem.

Figure A.3 maps collective-mechanism viability and buyer-group preference in the RA No-Banking and RN Banking setting. First, each point in the $(\alpha_{1s}, \alpha_{2s})$ space is classified by whether joint negotiation, agent-based negotiation, both, or neither is buyer-wise viable relative to separate negotiation. Second, when both collective mechanisms are viable, the buyer group chooses the mechanism with the higher total buyer gain.

The map supports the main analytical story under this set of numerical settings. Joint negotiation remains the broader participation-safe collective mechanism in that there is no region in which agent-based negotiation is viable while joint negotiation is not. Among the both-viable points, agent-based negotiation is preferred when the value of shielding the follower's share from seller extraction exceeds the delegation loss from sending only the agent into the seller-facing negotiation. Joint negotiation is preferred when pooled bargaining power is sufficiently strong, especially near the capped region where the buyer group can exploit the full seller-facing bargaining power available to the coalition. REC banking changes agent-based negotiation viability by adding a buyer-specific option value to



(a) RA No-Banking



(b) RN Banking

Figure A.3: Mechanism Viability and Buyer-Group Preference Maps

project output. In particular, banking can make agent-based negotiation viable in some low-bargaining-power regions where it is not viable in the no-banking benchmark, because the follower's internally agreed share includes banking value that is shielded from the seller-facing bargain. As bargaining powers increase, however, the buyers' separate-negotiation outside options also improve, so agent-based viability can narrow before reappearing when the acting agent receives enough delegated seller-facing power. Thus, REC banking amplifies the shielding benefit of agent-based negotiation, while joint negotiation continues to provide the more robust participation-safe benchmark by pooling seller-facing bargaining power. The numerical patterns therefore reinforce the central tradeoff, where agent-based negotiation can outperform joint negotiation when protected follower surplus is valuable and delegation loss is limited, whereas joint negotiation is preferred when pooled bargaining power dominates.

In summary, the numerical results reveal two consistent patterns. First, with respect to project size, agent-based negotiation tracks the separate-negotiation benchmark closely, while joint negotiation is more sensitive to mechanisms that aggregate both buyers' risk or banking benefits. Second, with respect to mechanism choice, joint negotiation is the more robust viability baseline, while agent-based negotiation becomes attractive when the agent is sufficiently empowered and the shielding benefit for the follower is large.

Appendix B

SUPPLEMENTARY MATERIAL FOR CHAPTER 3: BATTERY DISPATCH IN ELECTRICITY MARKETS WITH NEGATIVE PRICES

B.1 Proofs of Propositions and Corollaries

B.1.1 Proof of Proposition 3.1

We prove the result by backward induction on k .

For $k = N$, the terminal condition is

$$\tilde{J}_N(\cdot, \mathbf{x}) \equiv 0 \quad \text{for every } \mathbf{x} \in \mathcal{X}_N,$$

so the claim is immediate.

Now fix $k \in \{0, \dots, N-1\}$ and $\mathbf{x}_k \in \mathcal{X}_k$, and assume that for every $\mathbf{x}' \in \mathcal{X}_{k+1}$, the function $\tilde{J}_{k+1}(\cdot, \mathbf{x}')$ is continuous and piecewise affine on $[0, M^e]$. Then the continuation function

$$\tilde{\tilde{J}}_{k+1}(b'; \mathbf{x}_k) = \sum_{\mathbf{x}' \in \mathcal{X}_{k+1}} P_k(\mathbf{x}_k, \mathbf{x}') \tilde{J}_{k+1}(b', \mathbf{x}')$$

is also continuous and piecewise affine, because $\mathcal{X}_{k+1}$ is finite and a finite weighted sum of continuous piecewise-affine functions preserves both properties.

Let

$$0 = s_0 < s_1 < \dots < s_n = M^e$$

be a common refinement of the breakpoints of $\tilde{\tilde{J}}_{k+1}(\cdot; \mathbf{x}_k)$. On each interval

$$I_j = [s_{j-1}, s_j], \quad j = 1, \dots, n,$$

there exist constants a_j and c_j such that

$$\tilde{\tilde{J}}_{k+1}(b'; \mathbf{x}_k) = a_j b' + c_j, \quad b' \in I_j.$$

Consider first the charging problem in post-decision SoC:

$$\tilde{V}_k^c(b_k, \mathbf{x}_k) = \max_{b' \in [b_k, \min\{M^e, b_k + M^c\}]} \left\{ -\frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c}(b' - b_k) + \tilde{J}_{k+1}(b'; \mathbf{x}_k) \right\}.$$

Restricting to one continuation segment I_j , the objective becomes

$$-\frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c}(b' - b_k) + a_j b' + c_j = \left(a_j - \frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c} \right) b' + \frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c} b_k + c_j,$$

which is affine in b' . Hence, whenever the feasible charging window intersects I_j , the maximum over that overlap is attained at one of its two endpoints,

$$L_j^c(b_k) = \max\{b_k, s_{j-1}\}, \quad U_j^c(b_k) = \min\{M^e, b_k + M^c, s_j\}.$$

Exactly the same argument applies to the discharging problem

$$\tilde{V}_k^d(b_k, \mathbf{x}_k) = \max_{b' \in [\max\{0, b_k - M^d\}, b_k]} \left\{ -\eta_d \hat{p}_k(\mathbf{x}_k)(b' - b_k) + \tilde{J}_{k+1}(b'; \mathbf{x}_k) \right\},$$

because on I_j its objective is

$$-\eta_d \hat{p}_k(\mathbf{x}_k)(b' - b_k) + a_j b' + c_j = (a_j - \eta_d \hat{p}_k(\mathbf{x}_k)) b' + \eta_d \hat{p}_k(\mathbf{x}_k) b_k + c_j,$$

again affine in b' . Therefore any optimizer over the overlap with I_j is attained at one of the two endpoints

$$L_j^d(b_k) = \max\{0, b_k - M^d, s_{j-1}\}, \quad U_j^d(b_k) = \min\{b_k, s_j\}.$$

Thus the Bellman update is generated by finitely many endpoint candidates. Each endpoint map above is obtained by maxima and minima of affine functions of b_k , so its formula can change only when b_k crosses one of the values

$$\{s_j\}_{j=0}^n \cup \{s_j - M^c\}_{j=0}^n \cup \{s_j + M^d\}_{j=0}^n,$$

after clipping to $[0, M^e]$. On every open interval between two consecutive such values, each endpoint map is affine in b_k ; consequently each candidate value is affine in b_k . Since $\tilde{J}_k(\cdot, \mathbf{x}_k)$ is the maximum of finitely many affine candidates on each such interval, it is piecewise affine on $[0, M^e]$.

Continuity follows from Berge's maximum theorem. For fixed $\mathbf{x}_k$, the Bellman objective is continuous in (b_k, u_k) , and the feasible set

$$U(b_k) = [-M^d, M^c] \cap [-b_k, M^e - b_k]$$

is a nonempty compact interval that varies continuously with b_k . Hence $\tilde{J}_k(\cdot, \mathbf{x}_k)$ is continuous.

Finally, each endpoint candidate $b'^*(b_k)$ induces a candidate control

$$u^*(b_k) = b'^*(b_k) - b_k,$$

which is affine on each event interval. Because only finitely many such candidates are involved, their pairwise equality points form a finite set. Refining the event partition by those switching points yields a finite partition of $[0, M^e]$ on which one affine candidate control is optimal. Therefore there exists an optimal policy $u_k^*(\cdot, \mathbf{x}_k)$ that is piecewise affine in b_k .

Since $\mathbf{x}_k \in \mathcal{X}_k$ was arbitrary, the induction is complete.

B.1.2 Proof of Proposition 3.2

Fix k and $\mathbf{x}_k \in \mathcal{X}_k$, and abbreviate

$$\hat{p}_k := \hat{p}_k(\mathbf{x}_k), \quad \tilde{J}(b') := \tilde{J}_{k+1}(b'; \mathbf{x}_k).$$

Let $I = (e_r, e_{r+1})$ be an open interval between two consecutive points of the event grid $E_k(\mathbf{x}_k)$.

Because $E_k(\mathbf{x}_k)$ contains every continuation breakpoint s_j , every translated entry point $s_j - M^c$, and every translated exit point $s_j + M^d$, nothing qualitative changes as b_k varies within I :

1. b_k stays in the same continuation segment of $\tilde{J}$;
2. the upper charging boundary

$$u^c(b_k) := \min\{M^e, b_k + M^c\}$$

stays in the same continuation segment;

3. the lower discharging boundary

$$u^d(b_k) := \max\{0, b_k - M^d\}$$

stays in the same continuation segment; and

4. the sets of reachable interior breakpoints for charging and discharging remain fixed.

For $b_k \in I$, define

$$S^c(b_k) := \{s_j : b_k < s_j < u^c(b_k)\}, \quad S^d(b_k) := \{s_j : u^d(b_k) < s_j < b_k\}.$$

By construction of the event grid, these sets are the same for every $b_k \in I$. Denote the common sets by S_I^c and S_I^d .

We now examine the only candidates that can be optimal on I .

Stay-put candidate. Choosing $b' = b_k$ yields the value

$$\tilde{J}(b_k).$$

Since b_k remains in one continuation segment over I , this candidate is affine in b_k .

Charging-boundary candidate. Choosing the upper boundary of the feasible charging window yields

$$-\frac{\hat{p}_k}{\eta_c}(u^c(b_k) - b_k) + \tilde{J}(u^c(b_k)).$$

Over I , the map $u^c(b_k)$ is affine in b_k and remains in one continuation segment, so this candidate is affine in b_k .

Best interior charging candidate. For any fixed reachable breakpoint $s_j \in S_I^c$,

$$-\frac{\hat{p}_k}{\eta_c}(s_j - b_k) + \tilde{J}(s_j) = \frac{\hat{p}_k}{\eta_c}b_k + \left(\tilde{J}(s_j) - \frac{\hat{p}_k}{\eta_c}s_j\right).$$

All such candidates have the same slope $\hat{p}_k/\eta_c$, so their pointwise maximum is again a single affine function:

$$\frac{\hat{p}_k}{\eta_c}b_k + \max_{s_j \in S_I^c} \left(\tilde{J}(s_j) - \frac{\hat{p}_k}{\eta_c}s_j\right),$$

or is absent if $S_I^c = \emptyset$.

Discharging-boundary candidate. Choosing the lower boundary of the feasible discharging window yields

$$-\eta_d \hat{p}_k (u^d(b_k) - b_k) + \tilde{J}(u^d(b_k)).$$

Over I , the map $u^d(b_k)$ is affine in b_k and remains in one continuation segment, so this candidate is affine in b_k .

Best interior discharging candidate. For any fixed reachable breakpoint $s_j \in S_I^d$,

$$-\eta_d \hat{p}_k (s_j - b_k) + \tilde{J}(s_j) = \eta_d \hat{p}_k b_k + \left(\tilde{J}(s_j) - \eta_d \hat{p}_k s_j \right).$$

Again all such candidates have the same slope $\eta_d \hat{p}_k$, so their pointwise maximum is a single affine function:

$$\eta_d \hat{p}_k b_k + \max_{s_j \in S_I^d} \left(\tilde{J}(s_j) - \eta_d \hat{p}_k s_j \right),$$

or is absent if $S_I^d = \emptyset$.

Therefore, on the interval I , the Bellman update is the upper envelope of at most five affine functions: stay-put, charging boundary, best interior charging, discharging boundary, and best interior discharging.

The final statement follows immediately. A new breakpoint of $\tilde{J}_k(\cdot, \mathbf{x}_k)$ can arise only when one of these candidates changes formula, which happens at an event-grid point, or when two candidates intersect within the event interval.

B.1.3 Proof of Corollary 3.3

Suppose $\tilde{J}_{k+1}(\cdot; \mathbf{x}_k)$ has n affine segments. Then it has breakpoints

$$0 = s_0 < s_1 < \cdots < s_n = M^e.$$

The three sorted lists

$$\{s_j\}_{j=0}^n, \quad \{s_j - M^e\}_{j=0}^n, \quad \{s_j + M^d\}_{j=0}^n$$

contain at most $3n + 3$ values in total. After clipping to $[0, M^e]$ and deleting duplicates, the event grid $E_k(\mathbf{x}_k)$ therefore has at most $3n + 3$ points.

To compute the Bellman update, first form the event grid by a linear-time merge of the three sorted lists. Then sweep the resulting event intervals from left to right.

For the interior-breakpoint candidates, define the breakpoint scores

$$\alpha_j^c := \tilde{J}_{k+1}(s_j; \mathbf{x}_k) - \frac{\hat{p}_k(\mathbf{x}_k)}{\eta_c} s_j, \quad \alpha_j^d := \tilde{J}_{k+1}(s_j; \mathbf{x}_k) - \eta_d \hat{p}_k(\mathbf{x}_k) s_j.$$

As b_k increases, the reachable charging breakpoints are exactly those with $s_j \in [b_k, b_k + M^c]$, and the reachable discharging breakpoints are exactly those with $s_j \in [b_k - M^d, b_k]$. Because the breakpoints s_j are ordered, these are moving index windows over the two sequences $\{\alpha_j^c\}$ and $\{\alpha_j^d\}$. Their maxima can be maintained in amortized $O(1)$ time per event interval using standard monotone dequeues, since each breakpoint enters and leaves each window at most once over the full sweep.

By Proposition 3.2, once those two window maxima are known, the update on a given event interval is the upper envelope of at most five affine functions. Computing the local envelope of a constant number of affine functions takes $O(1)$ time.

Hence the entire Bellman update for a fixed exogenous state $\mathbf{x}_k$ takes $O(n)$ time. The algorithm stores the event grid, the breakpoint scores, and the resulting affine pieces, so the memory requirement is $O(n)$.

Finally, each event interval contributes only $O(1)$ affine pieces to the updated envelope. Since the number of event intervals is $O(n)$, the updated value function $\tilde{J}_k(\cdot, \mathbf{x}_k)$ also has $O(n)$ affine segments.

B.1.4 Proof of Proposition 3.4

The statement concerns the backward pass conditional on the reduced exogenous graph constructed in Algorithm 1. Thus, in this proof, the representative state sets $\mathcal{X}_k$, the projection maps Π_k , and the reduced transition kernels P_k are treated as already available. The cost of constructing the reduced graph is not included in the backward-pass complexity bound. The shock discretization sizes Q_k are also treated as fixed design parameters, as stated in Proposition 3.4. We keep the factor Q_k in the bound to show how the backward pass scales with the chosen shock discretization.

Fix a period $k \in \{0, \dots, N-1\}$ and a representative exogenous state $\mathbf{x}_k \in \mathcal{X}_k$. Under the discretized shock support

$$\{(\epsilon_{kq}, \pi_{kq})\}_{q=1}^{Q_k}$$

and the projection map Π_k , the continuation function in the implemented Bellman recursion is

$$\tilde{J}_{k+1}(\cdot; \mathbf{x}_k) = \sum_{q=1}^{Q_k} \pi_{kq} \tilde{J}_{k+1}(\cdot, \Pi_k(\mathbf{x}_k, \epsilon_{kq})).$$

Equivalently, using the reduced transition kernel,

$$\tilde{J}_{k+1}(\cdot; \mathbf{x}_k) = \sum_{\mathbf{x}' \in \mathcal{X}_{k+1}} P_k(\mathbf{x}_k, \mathbf{x}') \tilde{J}_{k+1}(\cdot, \mathbf{x}').$$

The first representation shows that, for a fixed $\mathbf{x}_k$, the continuation function is assembled from at most Q_k child value functions, regardless of the total size of $\mathcal{X}_{k+1}$. If multiple shock nodes project to the same child representative state, their probabilities are simply aggregated in $P_k(\mathbf{x}_k, \mathbf{x}')$, so the number of distinct child functions can only be smaller.

Let

$$D_k(\mathbf{x}_k) := \{\Pi_k(\mathbf{x}_k, \epsilon_{kq}) : q = 1, \dots, Q_k\} \subseteq \mathcal{X}_{k+1}$$

denote the set of distinct projected child states, and let

$$d_k(\mathbf{x}_k) := |D_k(\mathbf{x}_k)| \leq Q_k.$$

Then

$$\tilde{J}_{k+1}(\cdot; \mathbf{x}_k) = \sum_{\mathbf{x}' \in D_k(\mathbf{x}_k)} P_k(\mathbf{x}_k, \mathbf{x}') \tilde{J}_{k+1}(\cdot, \mathbf{x}').$$

For each child state $\mathbf{x}' \in D_k(\mathbf{x}_k)$, the value function $\tilde{J}_{k+1}(\cdot, \mathbf{x}')$ has $n_{k+1}(\mathbf{x}')$ affine segments and therefore $n_{k+1}(\mathbf{x}') - 1$ interior breakpoints. By definition,

$$n_{k+1}(\mathbf{x}') \leq n_{k+1}^{\max}.$$

The weighted sum defining $\tilde{J}_{k+1}(\cdot; \mathbf{x}_k)$ can change affine formula only at a breakpoint of one of the child value functions. Therefore the common refinement of the child breakpoint lists has at most

$$1 + \sum_{\mathbf{x}' \in D_k(\mathbf{x}_k)} (n_{k+1}(\mathbf{x}') - 1) \leq 1 + d_k(\mathbf{x}_k)(n_{k+1}^{\max} - 1) \leq 1 + Q_k(n_{k+1}^{\max} - 1)$$

affine segments. Hence the continuation function $\tilde{J}_{k+1}(\cdot; \mathbf{x}_k)$ has

$$O(Q_k n_{k+1}^{\max})$$

affine segments.

Because the piecewise-affine value functions are stored by their sorted breakpoint lists and affine coefficients, the continuation function can be assembled by merging the breakpoint lists of the child value functions. Since Q_k is treated as a fixed shock-discretization parameter, merging at most Q_k sorted lists is linear in the total number of breakpoint entries. On the resulting common refinement, the affine coefficient of $\tilde{J}_{k+1}(\cdot; \mathbf{x}_k)$ is obtained by summing the corresponding child coefficients weighted by $P_k(\mathbf{x}_k, \mathbf{x}')$. Thus, for the fixed representative state $\mathbf{x}_k$, the continuation function can be assembled in

$$O(Q_k n_{k+1}^{\max})$$

time.

If Q_k were instead treated as a growing asymptotic parameter, a generic comparison-based multiway merge could introduce an additional $\log(Q_k + 1)$ factor. This factor is not included in the stated bound because Proposition 3.4 treats Q_k as a fixed discretization parameter.

Once

$$\tilde{J}_{k+1}(\cdot; \mathbf{x}_k)$$

has been assembled, Corollary 3.3 applies directly. Since this continuation function has

$$O(Q_k n_{k+1}^{\max})$$

affine segments, the exact one-dimensional SoC Bellman update for the fixed representative state $\mathbf{x}_k$ can be computed in

$$O(Q_k n_{k+1}^{\max})$$

time. The update produces $\tilde{J}_k(\cdot, \mathbf{x}_k)$ and an associated optimal policy $u_k^*(\cdot, \mathbf{x}_k)$.

Therefore, the total work for one representative state $\mathbf{x}_k \in \mathcal{X}_k$ is

$$O(Q_k n_{k+1}^{\max}).$$

Since

$$m_k := |\mathcal{X}_k|,$$

the total work at period k is

$$O(m_k Q_k n_{k+1}^{\max}).$$

Summing over $k = 0, \dots, N-1$ gives the stated backward-pass time complexity:

$$O\left(\sum_{k=0}^{N-1} m_k Q_k n_{k+1}^{\max}\right).$$

For memory, storing the piecewise-affine representation of $\tilde{J}_k(\cdot, \mathbf{x}_k)$ for a fixed representative state $\mathbf{x}_k \in \mathcal{X}_k$ requires

$$O(n_k(\mathbf{x}_k))$$

space. Storing the value functions for all representative states at period k therefore requires

$$O\left(\sum_{\mathbf{x}_k \in \mathcal{X}_k} n_k(\mathbf{x}_k)\right) \leq O(m_k n_k^{\max}).$$

If the associated piecewise-affine policy $u_k^*(\cdot, \mathbf{x}_k)$ is also stored, this changes only the constant factor. Summing over all periods gives

$$O\left(\sum_{k=0}^N m_k n_k^{\max}\right).$$

This memory count is for the value and policy representations used by the backward dynamic program. It excludes storage of the precomputed reduced exogenous graph, including the representative state sets, projection maps, and transition kernels, which are produced by the forward graph construction.

Therefore, conditional on the precomputed reduced graph and fixed shock discretization sizes Q_k , the backward pass can be executed in

$$O\left(\sum_{k=0}^{N-1} m_k Q_k n_{k+1}^{\max}\right)$$

time and

$$O\left(\sum_{k=0}^N m_k n_k^{\max}\right)$$

memory.

B.1.5 Proof of Proposition 3.5

Fix any bounded continuation function

$$J : [0, M^e] \times \mathbb{X}_{k+1} \rightarrow \mathbb{R}.$$

Define

$$G_k(b, u, \mathbf{x}) := -\max\{\eta_d u, u/\eta_c\} \hat{p}_k(\mathbf{x}) + \int J(b + u, h(\mathbf{x}, \epsilon)) dF_k(\epsilon),$$

$$G_k^Q(b, u, \mathbf{x}) := -\max\{\eta_d u, u/\eta_c\} \hat{p}_k(\mathbf{x}) + \sum_{q=1}^{Q_k} \pi_{kq} J(b + u, h(\mathbf{x}, \epsilon_{kq})),$$

and

$$\tilde{G}_k(b, u, \mathbf{x}) := -\max\{\eta_d u, u/\eta_c\} \hat{p}_k(\mathbf{x}) + \sum_{q=1}^{Q_k} \pi_{kq} J(b + u, A_{k+1}(h(\mathbf{x}, \epsilon_{kq}))).$$

Then

$$(T_k J)(b, \mathbf{x}) = \max_{u \in U(b)} G_k(b, u, \mathbf{x}), \quad (\tilde{T}_k J)(b, \mathbf{x}) = \max_{u \in U(b)} \tilde{G}_k(b, u, \mathbf{x}).$$

Using the elementary inequality

$$\left| \max_{u \in U(b)} a_u - \max_{u \in U(b)} b_u \right| \leq \sup_{u \in U(b)} |a_u - b_u|,$$

we obtain

$$\left| (T_k J)(b, \mathbf{x}) - (\tilde{T}_k J)(b, \mathbf{x}) \right| \leq \sup_{u \in U(b)} \left| G_k(b, u, \mathbf{x}) - \tilde{G}_k(b, u, \mathbf{x}) \right|.$$

Now add and subtract $G_k^Q(b, u, \mathbf{x})$:

$$\left| G_k(b, u, \mathbf{x}) - \tilde{G}_k(b, u, \mathbf{x}) \right| \leq \left| G_k(b, u, \mathbf{x}) - G_k^Q(b, u, \mathbf{x}) \right| + \left| G_k^Q(b, u, \mathbf{x}) - \tilde{G}_k(b, u, \mathbf{x}) \right|.$$

The first term is exactly the error from replacing the integral under F_k with the discrete support $\{(\epsilon_{kq}, \pi_{kq})\}_{q=1}^{Q_k}$. Hence

$$\left| G_k(b, u, \mathbf{x}) - G_k^Q(b, u, \mathbf{x}) \right| \leq \Delta_k^{\text{disc}}(J).$$

For the second term, the common immediate-reward term cancels, and the triangle inequality gives

$$\begin{aligned} |G_k^Q(b, u, \mathbf{x}) - \tilde{G}_k(b, u, \mathbf{x})| &= \left| \sum_{q=1}^{Q_k} \pi_{kq} \left[J(b + u, h(\mathbf{x}, \epsilon_{kq})) - J(b + u, A_{k+1}(h(\mathbf{x}, \epsilon_{kq}))) \right] \right| \\ &\leq \sum_{q=1}^{Q_k} \pi_{kq} |J(b + u, h(\mathbf{x}, \epsilon_{kq})) - J(b + u, A_{k+1}(h(\mathbf{x}, \epsilon_{kq})))| \\ &\leq \Delta_k^{\text{proj}}(J). \end{aligned}$$

Therefore, for every $(b, \mathbf{x}) \in [0, M^e] \times \mathbb{X}_k$,

$$|(T_k J)(b, \mathbf{x}) - (\tilde{T}_k J)(b, \mathbf{x})| \leq \Delta_k^{\text{disc}}(J) + \Delta_k^{\text{proj}}(J).$$

Taking the supremum over $b \in [0, M^e]$ and $\mathbf{x} \in \mathbb{X}_k$ yields the result.

B.2 Holt-Winters model embedding

This appendix makes the Holt-Winters (HW) embedding explicit in the notation of Section 3.3. Throughout, $\mathbf{x}_k$ denotes the exogenous state observed at the *beginning* of period k , before p_k is realized.

B.2.1 Baseline additive HW embedding

Let m denote the seasonal period. To avoid ambiguity from zero-based versus one-based indexing, define the seasonal-position map

$$\rho(k) := 1 + ((k - 1) \bmod m) \in \{1, \dots, m\}.$$

Store the seasonal factors in chronological order as

$$\mathbf{s}_{k-1} := (s_{k-m}, s_{k-m+1}, \dots, s_{k-1})^\top \in \mathbb{R}^m.$$

Hence the factor relevant for forecasting period k is the first component,

$$[\mathbf{s}_{k-1}]_1 = s_{k-m}.$$

The beginning-of-period HW state is

$$\mathbf{x}_k^{\text{HW}} := (l_{k-1}, t_{k-1}, \mathbf{s}_{k-1}).$$

The additive HW observation equation is

$$p_k = l_{k-1} + t_{k-1} + [\mathbf{s}_{k-1}]_1 + \varepsilon_k, \quad (\text{B.1})$$

so the conditional mean forecast is

$$\hat{p}_k(\mathbf{x}_k^{\text{HW}}) = l_{k-1} + t_{k-1} + [\mathbf{s}_{k-1}]_1.$$

In the experiments, the one-step shock ε_k has mean zero and an hour-specific distribution $F_{\rho(k)}$, with variance $\sigma_{\rho(k)}^2$ estimated from the corresponding training residuals.

The conventional additive HW updating equations are

$$l_k = \alpha(p_k - s_{k-m}) + (1 - \alpha)(l_{k-1} + t_{k-1}), \quad (\text{B.2})$$

$$t_k = \beta(l_k - l_{k-1}) + (1 - \beta)t_{k-1}, \quad (\text{B.3})$$

$$s_k = \gamma(p_k - l_k) + (1 - \gamma)s_{k-m}, \quad (\text{B.4})$$

with smoothing parameters $\alpha, \beta, \gamma \in (0, 1)$.

Substituting (B.1) into (B.2)–(B.4) yields

$$l_k = l_{k-1} + t_{k-1} + \alpha\varepsilon_k, \quad (\text{B.5})$$

$$t_k = t_{k-1} + \alpha\beta\varepsilon_k, \quad (\text{B.6})$$

$$s_k = s_{k-m} + (1 - \alpha)\gamma\varepsilon_k. \quad (\text{B.7})$$

Now define

$$\mathbf{s}_k = (s_{k-m+1}, s_{k-m+2}, \dots, s_k)^\top.$$

Equivalently, if $C \in \mathbb{R}^{m \times m}$ denotes the cyclic-shift matrix

$$C = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 1 \\ 1 & 0 & \cdots & 0 & 0 \end{bmatrix}, \quad \mathbf{e}^{(m)} := (0, \dots, 0, 1)^\top,$$

then (B.7) implies

$$\mathbf{s}_k = C\mathbf{s}_{k-1} + (1 - \alpha)\gamma\varepsilon_k \mathbf{e}^{(m)}.$$

Therefore the next-period state $\mathbf{x}_{k+1}^{\text{HW}} = (l_k, t_k, \mathbf{s}_k)$ satisfies the linear transition

$$h_{\text{HW}}(\mathbf{x}_k^{\text{HW}}, \varepsilon_k) = \left(l_{k-1} + t_{k-1} + \alpha \varepsilon_k, t_{k-1} + \alpha \beta \varepsilon_k, C \mathbf{s}_{k-1} + (1 - \alpha) \gamma \varepsilon_k \mathbf{e}^{(m)} \right). \quad (\text{B.8})$$

Thus additive HW fits the generic structure in Section 3.3: a Markov forecasting state, a conditional mean forecast, and a one-step transition driven by a scalar shock.

B.2.2 Dependence-aware extension used in the real-data comparison

If the real-data comparison retains the dependence-aware specification in Section 3.4.3, the HW embedding above extends directly.

Let

$$e_k := p_k - \hat{p}_k(\mathbf{x}_k^{\text{HW}})$$

denote the period- k forecast error, and suppose the residual dependence is modeled by

$$e_k = \phi e_{k-1} + \nu_k,$$

where ν_k is a zero-mean innovation with an hour-specific distribution $F_{\rho(k)}^\nu$. Then the beginning-of-period augmented state is

$$\mathbf{x}_k^{\text{dep}} := (\mathbf{x}_k^{\text{HW}}, e_{k-1}).$$

Given $\mathbf{x}_k^{\text{dep}}$, the realized period- k forecast error is

$$e_k = \phi e_{k-1} + \nu_k,$$

so the observation equation becomes

$$p_k = \hat{p}_k(\mathbf{x}_k^{\text{HW}}) + e_k.$$

The next augmented state therefore satisfies

$$\mathbf{x}_{k+1}^{\text{dep}} = h_{\text{dep}}(\mathbf{x}_k^{\text{dep}}, \nu_k) := (h_{\text{HW}}(\mathbf{x}_k^{\text{HW}}, \phi e_{k-1} + \nu_k), \phi e_{k-1} + \nu_k).$$

Hence the dependence-aware specification remains Markov once the lagged forecast error is included in the state, and the stagewise-independent shock in the DP is the innovation ν_k , not the full forecast error e_k .

B.3 Deterministic LP benchmark (certainty-equivalent and perfect foresight).

Given a price path $\{\tilde{p}_k\}_{k=0}^{N-1}$, per-period charge/discharge limits $\{M_k^c, M_k^d\}$, energy capacity M^e , round-trip efficiency parameter $\eta \in (0, 1]$, and big- M constants $M_u, M_g > 0$, solve:

$$\max_{\{b_k, u_k, g_k, y_k\}} - \sum_{k=0}^{N-1} g_k \tilde{p}_k \quad (\text{B.9})$$

$$\text{s.t. } b_{k+1} = b_k + u_k, \quad k = 0, \dots, N-1, \quad (\text{B.10})$$

$$0 \leq b_k \leq M^e, \quad k = 0, \dots, N, \quad (\text{B.11})$$

$$-M_k^d \leq u_k \leq M_k^c, \quad k = 0, \dots, N-1, \quad (\text{B.12})$$

$$-b_k \leq u_k \leq M^e - b_k, \quad k = 0, \dots, N-1, \quad (\text{B.13})$$

$$u_k \geq -M_u(1 - y_k), \quad u_k \leq M_u y_k, \quad k = 0, \dots, N-1, \quad (\text{B.14})$$

$$g_k \geq \frac{1}{\eta} u_k - M_g(1 - y_k), \quad g_k \leq \frac{1}{\eta} u_k + M_g(1 - y_k), \quad k = 0, \dots, N-1, \quad (\text{B.15})$$

$$g_k \geq \eta u_k - M_g y_k, \quad g_k \leq \eta u_k + M_g y_k, \quad k = 0, \dots, N-1, \quad (\text{B.16})$$

$$y_k \in \{0, 1\}, \quad u_k, g_k \in \mathbb{R}, \quad b_k \in [0, M^e]. \quad (\text{B.17})$$

Here $y_k = 1$ (charging mode) forces $u_k \geq 0$ and $g_k = \frac{1}{\eta} u_k$, while $y_k = 0$ (discharging mode) forces $u_k \leq 0$ and $g_k = \eta u_k$. Choose big- M constants that dominate feasible ranges, where

$$M_u \geq \max_k \{M_k^c, M_k^d, M^e\}, \quad M_g \geq \max_k \left\{ \frac{M_k^c}{\eta}, \eta M_k^d, \frac{M^e}{\eta}, \eta M^e \right\}.$$

Appendix C

SUPPLEMENTARY MATERIAL FOR CHAPTER CHAPTER 4: WORKER SCHEDULE QUALITY DISPARITIES

C.1 Robustness Analysis

In this section, we conduct several robustness checks to validate our main analysis.

C.1.1 Results Using Server Workers Only

Although our main specification controls for role composition via $Role_i$, it may not fully capture the heterogeneity across different role types. For example, bartenders and cocktail servers may follow distinct scheduling rules from servers, reflecting the differing demand patterns of alcohol service versus standard dining. To address this, Table C.1 reports the full-control regressions over the 9-month subsample restricted to workers who exclusively work as servers. For White and minority female workers, the estimated coefficients remain consistent with our baseline findings in terms of both direction and magnitude. For minority male workers, we no longer observe a statistically significant sufficiency gap in offered hours compared to White men, and they now have significantly more worked hours. Nevertheless, the directional pattern between offered and worked sufficiency gaps remains similar: minority male workers initially receive fewer offered hours than what they are available for and subsequently rely on last-minute additions to attain more hours.

C.1.2 Alternative Race Classification

To classify the race and gender of workers, two widely-used algorithms, NamSor and NamePrism, were used in combination with manual verification in cases of disagreement. NamSor is a machine-learning algorithm trained on global datasets, including the names of Olympic athletes. It analyzes both first and last names to infer race and gender, reporting an average accuracy of 75% for race classification across U.S. ethnicities and over 95% for gender

Table C.1: Schedule Disparities Among Server-Only Workers.

Dependent Variables:	$\log(\text{OfferedWeeklyHrs}_{ijt} + 1)$	$RTARatio_{ijt}$	$SNARatio_{ijt}$	$\log(\text{WorkedWeeklyHrs}_{ijt} + 1)$
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
<i>WhiteFemale_i</i>	-0.0913*** (0.0128)	-0.0002 (0.0057)	0.0085** (0.0041)	-0.0797*** (0.0136)
<i>MinorityMale_i</i>	0.0360 (0.0224)	-0.0051 (0.0089)	0.0106 (0.0064)	0.0880*** (0.0224)
<i>MinorityFemale_i</i>	-0.0493*** (0.0177)	0.0157** (0.0075)	0.0231*** (0.0057)	0.0197 (0.0183)
Controls and fixed-effects	Yes	Yes	Yes	Yes
Observations	16,130	15,669	15,552	16,130
R ²	0.08852	0.13797	0.22624	0.17029

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

identification. By using both names, NamSor provides a nuanced and context-sensitive prediction. NamePrism, in contrast, uses a database of 74 million names derived from contact lists spanning 118 countries. It generates probabilistic scores for different racial and ethnic groups based on first names alone, offering a broader but sometimes less specific classification.

In the main analysis, for gender classification, NamSor results are used.¹ For race classification, the process began with both algorithms independently assigning race categories to the workers based on their names. For context, NamSor is particularly effective for names with strong ethnic associations, such as "Nguyen" for Vietnamese or "Gonzalez" for Hispanic individuals. NamePrism, relying solely on first names, can struggle in cases where the first name is ethnically neutral or widely shared across groups, such as "John" or "Sarah". In contrast, NamSor's use of both first and last names allows it to detect patterns like

¹In some cultures, such as East Asian cultures, gender can be challenging to infer from first names alone. However, in our dataset, all East Asian last names are paired with English first names, making gender inference straightforward and accurate.

"Sarah Chang" being more likely to belong to an Asian individual and "Sarah Johnson" being more likely to belong to a White individual. However, both algorithms can face challenges when dealing with names that do not conform to typical patterns, necessitating the manual verification process.

When discrepancies between NamSor and NamePrism predictions occurred, manual verification was used by searching for each name in Google Images and reviewing publicly available photographs to infer the most likely race based on human judgment. For example, if NamSor predicted a worker as "Asian" and NamePrism classified the same individual as "Hispanic Latino", the specific name would be manually searched online to determine the most likely classification based on publicly available images. Following this approach, workers were classified into two genders—female and male—and four racial groups: White (Non-Hispanic), Black (Non-Hispanic), Asian, and Hispanic Latino. Approximately 12% of the names required manual adjudication, where the algorithms disagreed or provided low-confidence predictions. This additional verification step ensured that classifications were as accurate as possible, even in challenging cases. In our analyses, all of Black, Asian, and Hispanic Latino workers were combined into the larger minority worker group.

Since the main analysis relies on race classifications derived by combining predictions from two algorithms, we assess robustness by re-running the analysis using each algorithm's predictions independently. We present the corresponding results in Tables C.2 and C.3. When using race classifications generated by Namsor (Table C.2), the gender and racial disparities observed remain largely consistent with our main results in both direction and magnitude. The only exception is minority female workers, who now exhibit a significant sufficiency gap in worked hours relative to White men. Nevertheless, the difference between their offered sufficiency gap (10.7%) and worked sufficiency gap (6.5%) is comparable to the difference observed in our baseline results (6.9% versus 0). This suggests that although differences in availability may partly explain the sufficiency gap observed for minority female workers when using Namsor-based racial assignment, they do not fully account for this disparity.

When using NamPrism-based racial assignment, the results for both White and minority female workers remain consistent with those from the main model. For minority male

Table C.2: Schedule Disparities Using Namsor Algorithm.

Dependent Variables:	$\log(\text{OfferedWeeklyHrs}_{ijt} + 1)$	$RTARatio_{ijt}$	$SNARatio_{ijt}$	$\log(\text{WorkedWeeklyHrs}_{ijt} + 1)$
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
<i>WhiteFemale_i</i>	-0.0990*** (0.0099)	0.0036 (0.0044)	0.0120*** (0.0031)	-0.0808*** (0.0103)
<i>MinorityMale_i</i>	-0.0275** (0.0132)	-0.0018 (0.0056)	0.0078* (0.0040)	0.0005 (0.0138)
<i>MinorityFemale_i</i>	-0.1067*** (0.0110)	0.0141*** (0.0049)	0.0109*** (0.0035)	-0.0647*** (0.0116)
Controls and fixed-effects	Yes	Yes	Yes	Yes
Observations	29,767	29,014	28,922	29,767
R ²	0.12147	0.12545	0.20610	0.18124

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

workers, although the estimated disparities are no longer statistically significant, their magnitudes remain comparable to those in the main results. The reduced statistical significance may be attributable to the smaller sample size under this classification method; in the main analysis, minority males comprise only 7% of the workers, which limits statistical power.

C.1.3 A Three-Year Full Sample

Table C.4 presents the results with all controls except for worker intrinsic abilities using the 3-yr full sample.

C.2 Estimating Worker Skill Levels

We estimate workers' skill levels across three dimensions by extending the fixed-effects approach employed by Tan and Netessine [2019] and Kamalahmadi et al. [2021]. Specifically, we assume that a worker's skill remains constant within a given month but may vary from month to month, allowing for potential learning and forgetting effects over time [Shafer et al., 2001].

Table C.3: Schedule Disparities Using NamPrism Algorithm.

Dependent Variables:	$\log(OfferedWeeklyHrs_{ijt} + 1)$	$RTARatio_{ijt}$	$SNARatio_{ijt}$	$\log(WorkedWeeklyHrs_{ijt} + 1)$
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
$WhiteFemale_i$	-0.0980*** (0.0085)	0.0059 (0.0037)	0.0085*** (0.0026)	-0.0794*** (0.0088)
$MinorityMale_i$	-0.0333 (0.0211)	-0.0108 (0.0087)	0.0023 (0.0061)	-0.0060 (0.0223)
$MinorityFemale_i$	-0.0829*** (0.0153)	0.0153** (0.0068)	0.0128** (0.0051)	-0.0473*** (0.0166)
Controls and fixed-effects	Yes	Yes	Yes	Yes
Observations	29,767	29,014	28,922	29,767
R ²	0.12144	0.12539	0.20603	0.18130

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

To estimate skill levels in sales, service quality, and speed, we run the following worker-hour-level fixed-effects model separately for each month $m = 1, \dots, 9$:

$$\log(Y_{ihm}) = \beta_0 + \theta_{im} + \beta_1 \mathbf{X}_{ihm} + \epsilon_{ihm}, \quad (\text{C.1})$$

where θ_{im} captures the skill of worker i in month m , and $\mathbf{X}_{ihm}$ includes controls for hour of the day, day of the week, week of the year, and store fixed effect. To avoid confounding skill estimates with differences in job responsibilities, we restrict our analysis to transactions completed during server shifts only. Sales handled by bartenders, for instance, tend to be lower due to differing customer interaction patterns and cannot be adequately adjusted for via fixed effects or shift-type controls alone. We also note that customer characteristics will not bias the estimates of workers skills in our context. This is because customers are assigned to workers in a round round-robin fashion, which means customers served by different servers should be statistically similar.

As mentioned in Section 4.3.3, customers are assigned to servers in a round robin fash-

Table C.4: Schedule Disparities Using 3-yr Full Sample

Dependent Variables:	$\log(\text{OfferedWeeklyHrs}_{ijt} + 1)$	$RTARatio_{ijt}$	$SNARatio_{ijt}$	$\log(\text{WorkedWeeklyHrs}_{ijt} + 1)$
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
<i>WhiteFemale_i</i>	-0.0523*** (0.0046)	0.0104*** (0.0020)	0.0086*** (0.0014)	-0.0363*** (0.0048)
<i>MinorityMale_i</i>	-0.0148* (0.0080)	0.0108*** (0.0034)	0.0159*** (0.0024)	0.0219*** (0.0081)
<i>MinorityFemale_i</i>	-0.0449*** (0.0067)	0.0238*** (0.0030)	0.0187*** (0.0021)	0.0014 (0.0070)
<i>MidTenure_{ijt}</i>	-0.0065 (0.0050)	-0.0119*** (0.0021)	-0.0170*** (0.0015)	-0.0238*** (0.0052)
<i>LongTenure_{ijt}</i>	-0.0497*** (0.0062)	-0.0190*** (0.0027)	-0.0384*** (0.0018)	-0.1011*** (0.0064)
<i>FullAbsenteesim_{ijt}</i>	-1.017*** (0.0233)	-0.4536*** (0.0088)	0.1366*** (0.0097)	-2.014*** (0.0268)
<i>PartialAbsenteesim_{ijt}</i>	-0.4479*** (0.0313)	-0.6296*** (0.0102)	0.0701*** (0.0099)	-0.7090*** (0.0239)
<i>OfferedWeeklyHrs_{ijt}</i>		-0.0050*** (0.0002)	-0.0121*** (0.0001)	
<i>OfferedLTWeeklyHours_{ij}</i>		0.0020*** (0.0002)	0.0091*** (0.0002)	
<i>Fixed-effects</i>				
<i>Store_j</i>	Yes	Yes	Yes	Yes
<i>Year_t</i>	Yes	Yes	Yes	Yes
<i>Week_t</i>	Yes	Yes	Yes	Yes
<i>Role_i</i>	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	112,524	109,926	109,599	112,524
R ²	0.11298	0.10011	0.18126	0.17501

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table C.5: Schedule Disparities Using 3-yr Full Sample.

	Sch.Hours	RTARatio	SNARatio	MgrRTCRatio	MgrSNCRatio	Act.Hours
<i>WhiteFemale_i</i>	-4.98%***	1.30 p.p***	0.83 p.p***	-	-	-3.16%***
<i>MinorityMale_i</i>	-1.50%*	1.18 p.p***	1.59 p.p***	-	-	2.14%***
<i>MinorityFemale_i</i>	-4.08%***	2.89 p.p**	1.73 p.p***	-0.94 p.p***	-	-
Ability Controls	No	No	No	No	No	No
Other Controls						
and fixed-effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	112,524	109,926	109,599	109,926	109,926	112,524

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

ion. As such, customers served by different servers should be statistically similar and thus customer characteristics will not bias the three measurements above for workers' ability.

Each regression is estimated using a different performance metric Y_{ihm} , corresponding to one of the three skill dimensions:

- **Sales Ability.** We measure this using *per person average sales* (PPA), defined as the total sales divided by the total number of guests served by worker i in hour h of month m :

$$Y_{ihm} = \frac{\sum_{k \in ihm} \text{Sales}_k}{\sum_{k \in ihm} \text{Guests}_k}. \quad (\text{C.2})$$

PPA is widely used in the restaurant industry as a benchmark for worker performance [Mill, 2007].

- **Score Ability.** To capture service quality, we compute the average customer recommendation score across all transactions handled by the worker during each hour:

$$Y_{ihm} = \frac{1}{|ihm|} \sum_{k \in ihm} \text{RecommendationScore}_k. \quad (\text{C.3})$$

- **Speed Ability.** We assess service speed using the inverse of the average transaction duration in the focal hour. A higher value corresponds to faster service. Unlike the model used to evaluate sales and score ability, the speed model includes several additional controls: average party size per check, average sales per check, average number of items sold per check, and average worker load per check. These adjustments are essential because transaction duration is closely linked to sales effort—faster service may reflect either genuine efficiency or reduced sales effort (e.g., less upselling). Controlling for these factors allows us to interpret sales performance and service speed as distinct dimensions of worker productivity. Equation (C.4) presents the model used to estimate service speed.

$$\begin{aligned} \log \left(\frac{1}{\frac{1}{|ihm|} \sum_{k \in ihm} \text{Duration}_k} \right) = & \beta_0 + \theta_{im} + \beta_1 \mathbf{X}_{ihm} + \text{AvgNumItems}_{ihm} \\ & + \text{AvgSales}_{ihm} + \text{AvgWorkerLoad}_{ihm} + \text{AvgGuestCount}_{ihm} + \epsilon_{ihm}. \end{aligned} \quad (\text{C.4})$$

In total, we estimate 27 monthly fixed-effects models (nine months $\times$ three skill dimensions), generating skill estimates for each worker across all three dimensions over time. These estimates are then normalized to facilitate comparability and used as control variables in our main regression analyses.

C.3 Earning Potentials By Week-Hour

Hourly earning potentials are presented in Figure C.1. Peak lunch (12:00 PM - 1:00 PM) and dinner hours (5:00 PM - 7:00 PM) offer the highest hourly earning potential per worker. Weekends are slightly more profitable than weekdays with extended lunch and dinner windows.




Figure C.1: Expected Hourly Tip Earning Potential (\$/Worker) for Day-of-Week and Hour-of-Day Combinations.

Note: Values shown are the averages across all locations.

C.4 Sales Performance Analysis Procedure

Recall that in order to estimate the causal impact of level of advance notice on workers' sales performance during identical shifts, in Section 4.5 we proposed model (4.3), which is written as

$$\log(\text{Sale}_k) = \alpha_0 + \alpha_1 \text{SNA}_k + \alpha_2 \text{RTA}_k + \alpha_3 \mathbf{W}_k + \eta_k.$$

where

$$\mathbf{W}_k = (\text{GuestCount}_k, \log(\text{Tenure}_k), \text{StoreLoad}_k, \text{WorkerLoad}_k, \text{WorkerLoad}_k^2, \log(\text{Fatigue}_k), \text{EndOfShift}_k, \text{Hour}_k, \text{DayWeek}_k, \text{WeekYear}_k, \text{Month}_k, \text{Store}_k, \text{Worker}_k).$$

To properly effectively account for potential endogeneity due to the correlation between last-minute shift assignments and check size, we employ the Probit-2SLS method described in chapter 18.4 of Woodridge [2010] and used by Kamalahmadi et al. [2021]. To begin with, we construct instrumental variables, $\text{Lagged}\% \text{SNA}_k$ and $\text{Lagged}\% \text{RTA}_k$, corresponding to the key variables SNA_k and RTA_k , respectively. Specifically, $\text{Lagged}\% \text{SNA}_k$ ($\text{Lagged}\% \text{RTA}_k$) represents the percentage of workers who were on a short-notice (real-time) addition shift during the same hour as the focal transaction one week prior. The

estimation procedure is as follows.

Probit stage. Because the endogenous explanatory variables of interest, SNA_k and RTA_k , are binary indicators of shift assignment, we begin the two-stage procedure with a probit model. Unlike a linear first stage, the probit specification accounts for the nonlinear relationship between the instruments and these binary outcomes, yielding more appropriate fitted values. Specifically, we estimate the probability that each shift assignment occurs as a function of the instrumental variables $\mathbf{Z}_k$ and additional controls $\mathbf{W}_k$:

$$P(SNA_k = 1) = \Phi(\beta_0^{SN} + \beta_1^{SN}\mathbf{Z}_k + \beta_2^{SN}\mathbf{W}_k + \eta_k^{SN}), \quad (\text{C.5})$$

$$P(RTA_k = 1) = \Phi(\beta_0^{RT} + \beta_1^{RT}\mathbf{Z}_k + \beta_2^{RT}\mathbf{W}_k + \eta_k^{RT}), \quad (\text{C.6})$$

where Φ denotes the cumulative distribution function of the standard normal distribution. The predicted probabilities from this stage, denoted $\hat{P}(SNA_k = 1)$ and $\hat{P}(RTA_k = 1)$, are then used in the second-stage regression as fitted values for the endogenous shift assignment variables.

2SLS stage. Having obtained the probit-predicted probabilities for the binary shift indicators, we proceed with a standard 2SLS regression to estimate their causal effect on transaction-level sales. For each endogenous indicator $V_{end,k} \in \{SNA_k, RTA_k\}$, we regress

$$V_{end,k} = \beta_0^{V_{end}} + \beta_1^{V_{end}}\hat{P}(SNA_k = 1) + \beta_2^{V_{end}}\hat{P}(RTA_k = 1) + \beta_3^{V_{end}}\mathbf{W}_k + \eta_k^{V_{end}}, \quad (\text{C.7})$$

where $\mathbf{W}_k$ is the vector of exogenous controls. This yields fitted values $\widehat{SNA_k}$ and $\widehat{RTA_k}$. We then estimate the main sales performance model by replacing each endogenous variable with its fitted counterpart:

$$\log(Sale_k) = \alpha_0 + \alpha_1\widehat{SNA_k} + \alpha_2\widehat{RTA_k} + \alpha_3\mathbf{W}_k + \eta_k \quad (\text{C.8})$$

To examine whether last-minute notices affect sales performance differently across worker groups, we extend the model to include interaction terms between the shift notice indicators (RTA_k and SNA_k) and demographic dummies ($WhiteFemale_k$, $MinorityMale_k$, and $MinorityFemale_k$). The coefficients on these interaction terms capture how the impact of additional or short-notice shifts varies by demographic group, relative to the baseline group

of White males. The probit stage remains unchanged, while the second stage of the 2SLS procedure includes additional endogenous regressors corresponding to the interaction terms.

To simplify the notation and avoid repeatedly listing each interaction term explicitly, we use vector notation. Let $\mathbf{G}_k = [WhiteFemale_k, MinorityMale_k, MinorityFemale_k]$ represent the vector of demographic indicators, and let $\mathbf{N}_k = [SNA_k, RTA_k]$ denote the vector of shift-notice indicators. The interaction terms are represented compactly by $\mathbf{N}_k \otimes \mathbf{G}_k$, the Kronecker product of notice indicators and demographic dummies.

Specifically, in the first stage of the 2SLS procedure, for each endogenous variable $V_{end,k}$ in the set $\{\mathbf{N}_k, \mathbf{N}_k \otimes \mathbf{G}_k\}$, we estimate:

$$V_{end,k} = \beta_0^{V_{end}} + \beta_1^{V_{end}} \hat{\mathbf{P}}(\mathbf{N}_k) + \beta_2^{V_{end}} [\hat{\mathbf{P}}(\mathbf{N}_k) \otimes \mathbf{G}_k] + \beta_3^{V_{end}} \mathbf{W}_k + \eta_k^{V_{end}}, \quad (\text{C.9})$$

where $\hat{\mathbf{P}}(\mathbf{N}_k)$ denotes the vector of predicted probabilities from the probit stage. The fitted values from this stage are represented as $\hat{V}_{end,k}$.

In the second stage, we estimate the impact on log sales using these predicted endogenous regressors:

$$\log(Sale_k) = \alpha_0 + \alpha_1' \hat{\mathbf{N}}_k + \alpha_2' [\widehat{\mathbf{N}_k \otimes \mathbf{G}_k}] + \alpha_3 \mathbf{W}_k + \eta_k. \quad (\text{C.10})$$

C.5 *Discordance with Manager Background*

Numerous studies have shown that the discordance in race or gender between managers and workers can result in worse job outcomes for workers (e.g., Cohen et al. 1998, Cohen and Huffman 2007, Giuliano et al. 2009, Grissom and Keiser 2011). In the context of schedule quality, Storer et al. [2020] found that workers experience more precarious schedules when managed by a racially discordant supervisor. Furthermore, because the majority of managers are White, non-White workers are disproportionately exposed to such discordance, contributing to substantial racial disparities in precarious scheduling practices. Inspired by these findings, we investigate the moderating effects of race and gender discordance with managers in our dataset. Out of the 25 locations in our dataset, manager names are available for 24 of them, from which managers' gender and race are inferred following the same approach for workers as discussed in Section 4.3.3. Among these managers, 18 are White

males, plus 2 White females and 4 minority males. There have been no changes in managers during the 9-month period. Given the relatively small and unbalanced sample of managers, the findings presented here should be interpreted with caution and not overgeneralized to broader contexts. Nevertheless, the results reveal systemic patterns within our dataset that are roughly consistent with broader trends reported in the literature.

The moderating effects of discordance with managers are presented in Table C.6. We introduce two binary variables, $MngrSameGender_{ij}$ and $MngrSameRace_{ij}$, which are set to 1 when the manager and the focal worker share the same gender or race², respectively, and 0 otherwise. As shown in Columns 1 and 3, minority workers face smaller RTA ratio gaps relative to average White males when they have managers of the same gender, while White females experience a smaller SNA gap. This aligns with literature suggesting that workers tend to receive more favorable treatment from managers who share a similar demographic background.

Racial discordance appears to play a mixed role in shaping worker-manager interactions. As shown in Columns 2 and 4, only minority male workers receive less frequent SNA than White male workers under managers of the same race, while female workers appear to face less favorable treatment when their managers share the same race. Notably, in our dataset, all minority managers are male, meaning that when minority male workers have managers of the same race, they are also of the same gender. This suggests that minority male managers may treat workers who share both their race and gender more favorably. In contrast, for minority female workers, male managers of the same race appear to assign them more demanding schedules. This again echoes intersectionality theory, which highlights how minority women face complex challenges that go beyond the additive effects of race- and gender-based inequalities [Crenshaw, 2013]. In the case of White females, the fact that they receive a larger RTA gap under White managers suggests that they face fewer disadvantages under minority managers, potentially because their gender-related challenges are offset by an implicit racial status advantage in interactions with the manager.

²For minority workers and managers, race alignment is based on the more granular race classification, namely Black, Asian, or Hispanic Latino.

Table C.6: Moderating Effects of Gender and Race Discordance with Manager.

Dependent Variables:	$RTARatio_{ijt}$		$SNARatio_{ijt}$	
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
$WhiteFemale_i$	0.0051 (0.0042)	-0.0118 (0.0089)	0.0092*** (0.0029)	0.0044 (0.0060)
$MinorityMale_i$	0.0414* (0.0218)	0.0034 (0.0072)	0.0086 (0.0145)	0.0118** (0.0052)
$MinorityFemale_i$	0.0328*** (0.0061)	0.0305*** (0.0060)	0.0181*** (0.0045)	0.0152*** (0.0044)
$WhiteFemale_i \times MngrSameGender_{ij}$	-0.0018 (0.0162)		-0.0194* (0.0113)	
$MinorityMale_i \times MngrSameGender_{ij}$	-0.0445* (0.0230)		-0.0012 (0.0154)	
$MinorityFemale_i \times MngrSameGender_{ij}$	-0.0462** (0.0189)		-0.0065 (0.0141)	
$WhiteFemale_i \times MngrSameRace_{ij}$		0.0210** (0.0097)		0.0044 (0.0067)
$MinorityMale_i \times MngrSameRace_{ij}$		-0.0111 (0.0214)		-0.0368*** (0.0142)
$MinorityFemale_i \times MngrSameRace_{ij}$		-0.0234 (0.0183)		0.0340** (0.0149)
Controls and fixed-effects	Yes	Yes	Yes	Yes
Observations	27,967	27,967	27,885	27,885
R ²	0.12473	0.12445	0.20636	0.20669

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

VITA

Xiaoxuan Hou received her Ph.D. in Business Administration from the University of Washington. Her research interests include sustainable operations management, renewable energy operations, electricity markets, and workforce equity.