

Essays on cost pass-through under changing economic conditions

by

Kinza Emad Khan

B.Sc., Lahore University of Management Sciences, 2016

M.S., Lahore University of Management Sciences, 2018

AN ABSTRACT OF A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Economics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2026

Abstract

This dissertation consists of three essays. The first essay examines one of the most important questions in the business cycle literature: how consumer prices adjust to equilibrium after a cost shock to the economy. We consider tight credit conditions as a state affecting firms' pricing decisions. As credit conditions change, so should the willingness of firms to leave their price unchanged in response to a large increase in the cost of production. This essay evaluates the hypothesis that cost shocks are passed on to consumers more quickly when credit conditions tighten. Using time series data from FRED on U.S. airfares, jet fuel prices, and credit conditions, we examine the effect of credit conditions on the pass-through of jet fuel price shocks to U.S. airfares over the period 1990–2019. We model the coefficients determining pass-through as a function of credit conditions at the time of the shock to jet fuel prices. We find evidence of much faster responses of airfares to jet fuel price shocks when credit conditions tighten. With normal credit conditions, cost pass-through is delayed by approximately eight months due to hedging practices, suggesting that hedging postpones, but does not eliminate, the transmission of cost shocks to airfares.

The second essay examines whether the pass-through of jet fuel cost shocks differs systematically across high and low travel intensity seasons in the U.S. airline industry. Since factors such as route competition, capacity utilization, and airline pricing power vary systematically between peak and off peak travel periods as travel demand fluctuates, cost pass through may also exhibit seasonal variation. Using aggregate data on airfares and jet fuel prices, we estimate a structural vector autoregressive (SVAR) model and compute impulse response functions that vary by the month in which the shock hits. The results show that seasonality primarily affects the timing of pass-through rather than its long run magnitude. Fuel cost shocks are transmitted more rapidly when fares are set during peak travel periods,

while adjustment is more gradual during off-peak periods. These findings highlight seasonality as an important determinant of fuel cost pass-through and suggest that ignoring seasonal demand conditions may obscure important patterns in airline pricing behavior.

The third essay contributes to the literature on wage setting. The gender pay gap literature emphasizes that wage differentials are no longer primarily driven by human capital differences, but by workplace structures that reward long, inflexible work schedules while penalizing workers who value flexibility. Since high-skilled women, particularly those raising children, are more likely to demand flexibility, they are more exposed to these wage penalties. Increases in low-skilled immigration may relax their time constraints by reducing the cost of childcare through the expansion of low-cost household services. This paper examines whether low-skilled immigration affects the occupational inflexibility of high-skilled native women through this channel. We construct an index of occupational inflexibility using five occupation-level characteristics from O*NET and merge it with pooled individual-level data from the ACS of 2009 and 2019. To address endogenous immigrant location choices, we employ a shift-share instrumental variable strategy. We find suggestive evidence that low-skilled immigration reduces childcare costs at the PUMA level. However, its effect on occupational inflexibility among high-skilled native women is negative but statistically insignificant. Heterogeneity analysis shows that the effect remains negative for women without children but turns positive for mothers, with the largest effects among those with at least one child under age five, though estimates are imprecise across all subgroups.

Essays on cost pass-through under changing economic conditions

by

Kinza Emad Khan

B.Sc., Lahore University of Management Sciences, 2016

M.S., Lahore University of Management Sciences, 2018

A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Economics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2026

Approved by:

Major Professor
Lance Bachmeier

Copyright

© Kinza Emad Khan 2026.

Abstract

This dissertation consists of three essays. The first essay examines one of the most important questions in the business cycle literature: how consumer prices adjust to equilibrium after a cost shock to the economy. We consider tight credit conditions as a state affecting firms' pricing decisions. As credit conditions change, so should the willingness of firms to leave their price unchanged in response to a large increase in the cost of production. This essay evaluates the hypothesis that cost shocks are passed on to consumers more quickly when credit conditions tighten. Using time series data from FRED on U.S. airfares, jet fuel prices, and credit conditions, we examine the effect of credit conditions on the pass-through of jet fuel price shocks to U.S. airfares over the period 1990–2019. We model the coefficients determining pass-through as a function of credit conditions at the time of the shock to jet fuel prices. We find evidence of much faster responses of airfares to jet fuel price shocks when credit conditions tighten. With normal credit conditions, cost pass-through is delayed by approximately eight months due to hedging practices, suggesting that hedging postpones, but does not eliminate, the transmission of cost shocks to airfares.

The second essay examines whether the pass-through of jet fuel cost shocks differs systematically across high and low travel intensity seasons in the U.S. airline industry. Since factors such as route competition, capacity utilization, and airline pricing power vary systematically between peak and off peak travel periods as travel demand fluctuates, cost pass through may also exhibit seasonal variation. Using aggregate data on airfares and jet fuel prices, we estimate a structural vector autoregressive (SVAR) model and compute impulse response functions that vary by the month in which the shock hits. The results show that seasonality primarily affects the timing of pass-through rather than its long run magnitude. Fuel cost shocks are transmitted more rapidly when fares are set during peak travel periods,

while adjustment is more gradual during off-peak periods. These findings highlight seasonality as an important determinant of fuel cost pass-through and suggest that ignoring seasonal demand conditions may obscure important patterns in airline pricing behavior.

The third essay contributes to the literature on wage setting. The gender pay gap literature emphasizes that wage differentials are no longer primarily driven by human capital differences, but by workplace structures that reward long, inflexible work schedules while penalizing workers who value flexibility. Since high-skilled women, particularly those raising children, are more likely to demand flexibility, they are more exposed to these wage penalties. Increases in low-skilled immigration may relax their time constraints by reducing the cost of childcare through the expansion of low-cost household services. This paper examines whether low-skilled immigration affects the occupational inflexibility of high-skilled native women through this channel. We construct an index of occupational inflexibility using five occupation-level characteristics from O*NET and merge it with pooled individual-level data from the ACS of 2009 and 2019. To address endogenous immigrant location choices, we employ a shift-share instrumental variable strategy. We find suggestive evidence that low-skilled immigration reduces childcare costs at the PUMA level. However, its effect on occupational inflexibility among high-skilled native women is negative but statistically insignificant. Heterogeneity analysis shows that the effect remains negative for women without children but turns positive for mothers, with the largest effects among those with at least one child under age five, though estimates are imprecise across all subgroups.

Table of Contents

List of Figures	x
List of Tables	xi
Acknowledgements	xii
Dedication	xvi
1 When Credit Tightens, Prices Take Off: Evidence from Jet Fuel Price Shocks in U.S. Airfares	1
1.1 Introduction	1
1.2 Literature Review	4
1.2.1 Time-Dependent versus State-Dependent Pricing	4
1.2.2 Our Contribution	5
1.3 Fuel Hedging	6
1.4 Related Work on Airfares	8
1.5 Methodology	10
1.5.1 Identification of Impulse Response Functions	14
1.5.2 Impulse Response Functions (IRFs)	15
1.5.3 Confidence Intervals for Impulse Response Functions	16
1.6 Empirical Results	17
1.6.1 Robustness to Lag Length Selection	19
1.6.2 Robustness to the Measurement of Credit Conditions	21
1.7 Conclusion & Implications	23

2	Is Cost Pass-Through in the Airline Industry Seasonal?	26
2.1	Introduction	26
2.2	Literature Review	30
2.3	Methodology	34
2.3.1	Identification of Impulse Response Functions	36
2.3.2	Impulse Response Functions (IRFs)	37
2.3.3	Confidence Intervals for Impulse Response Functions	38
2.4	Results	39
2.4.1	Do Shocks in Different Months Affect the Same Future Month Differ- ently?	44
2.4.2	Airfare Responses to Extreme Jet Fuel Price Shocks	47
2.4.3	The Iran War Shock: Counterfactual Seasonal Pass Through	53
2.5	Conclusion & Implications	56
3	Low Skilled Immigration and the Hidden Link to the Gender Pay Gap	58
3.1	Introduction	58
3.2	Literature Review	63
3.3	Data Sources and Variable Construction	67
3.3.1	Descriptive Statistics	72
3.4	Mechanism - Analysis of Childcare Market	75
3.4.1	Identification Strategy: Endogeneity and Instrumental Variables	76
3.4.2	Evidence from the Childcare Market Mechanism	78
3.5	Low-Skilled Immigration and Occupational Flexibility	82
3.5.1	Heterogeneous Effects by Household Childcare Needs	86
3.6	Conclusion	90
	Bibliography	93

List of Figures

1.1	Debt-to-Asset Ratio by Airline (1990-2019)	3
1.2	Transition Weights under Alternative Values of the Smoothing Parameter . .	11
1.3	Logistic Transition Functions with respect to Credit Conditions	12
1.4	Domestic Revenue Passenger Miles (RPMs)	13
1.5	Response of Airfares to Jet Fuel Shock	18
1.6	Response of Airfares to Jet Fuel Shock - VAR(3)	20
1.7	Response of Airfares to Jet Fuel Shock (6 months Credit Change)	22
1.8	Response of Airfares to Jet Fuel Shock (3 months Credit Change)	23
2.1	Domestic Revenue Passenger Miles (RPMs)	35
2.2	Response of Airfares to Jet Fuel Shock	40
2.3	Response of Airfares to Jet Fuel Shock by Month	42
2.4	Airfares Responses in a Given Month to Jet Fuel Price Shocks in Previous Months	45
2.5	Response of Airfares to Jet Fuel Shocks by Month	49
2.6	Airfares Responses in a Given Month to Jet Fuel Price Shocks in Previous Months	52

List of Tables

2.1	Impulse Response of Airfares to Jet Fuel Shocks	54
3.1	Occupational Inflexibility by Occupation Category	73
3.2	Summary Statistics for Explanatory Variables and Labor Market Outcomes .	74
3.3	Effects of Low Skilled Immigration on Childcare Labor Supply and Wages .	80
3.4	Impact of LSIs on Occupational Inflexibility among High Skilled Native Women	84
3.5	Impact of LSIs on Occupational Inflexibility by Family Structure	87

Acknowledgments

In the name of Allah, the Most Gracious, the Most Merciful.

“For indeed, with hardship [will be] ease.”

— Quran - Surah Ash-Sharh (The Opening-Up of the Heart) (94:5)

All praise and gratitude belong to Allah alone. This dissertation is, above all else, a testament to His infinite mercy, wisdom, and blessings. Every step of this journey was written by Him long before I could understand it. He alone witnessed every moment that no one else did: the silent prayers that no one else saw, the uncertainty, the setbacks, the tears, the exhaustion, and the hope that often seemed beyond reach. There were countless times when I questioned my own ability to continue, yet Allah always knew the strength He had placed within me, even when I could not see it myself. Every opportunity I received, every person who entered my life, every challenge that shaped me, and every success I have been blessed with came only through His grace. I pray that this work is accepted as a humble expression of gratitude for the countless blessings He has bestowed upon me.

I owe my deepest gratitude to my advisor, Dr. Lance Bachmeier. It is no exaggeration to say that this dissertation would not exist without his mentorship. Throughout my doctoral studies, he provided not only invaluable academic guidance but also unwavering encouragement and confidence in my abilities. His remarkable insight, creativity, and ability to transform broad ideas into meaningful research fundamentally shaped every chapter of this dissertation. More importantly, he believed in me during moments when I struggled to believe in myself. As I worried about completing my dissertation in time for the job market and questioned whether I could meet the expectations I had set for myself, he remained patient, supportive, and reassuring. His confidence gave me confidence. His thoughtful advice, constructive criticism, and constant availability made even the most difficult research

challenges feel manageable. I consider myself incredibly fortunate to have learned from him and worked under his supervision.

I would also like to express my sincere appreciation to Dr. Valerie Bostwick for her generosity, encouragement, and willingness to collaborate on the third chapter of this dissertation. Her outstanding teaching in labor economics not only deepened my understanding of the field but also inspired many of the ideas that eventually became part of my research. I am grateful for her thoughtful guidance and her confidence in my work.

My sincere thanks also go to Dr. Steven P. Cassou, Dr. William Blankenau, and Dr. Jennifer Ifft. Each of them has contributed significantly to my education through the excellent courses they taught and the valuable knowledge they shared. I am especially thankful for their willingness to serve on my dissertation committee, for their thoughtful feedback, and for the time they invested in helping me complete this important milestone.

I am grateful to the Department of Economics at Kansas State University for providing financial support throughout my doctoral studies. I also sincerely appreciate the department's flexibility and understanding in accommodating my teaching responsibilities, which allowed me the time needed to focus on my research and successfully complete this dissertation.

No words can adequately express my gratitude to my parents. Everything I have achieved began with the values they instilled in me from a very young age. They taught me that education is the strongest foundation a person can build—one that can never be taken away regardless of life's circumstances. They encouraged me to dream beyond my limits, to work with sincerity, and to pursue excellence with humility. My father always dreamed that at least one of his children would earn a Ph.D., and I feel deeply honored to have fulfilled that dream. Although he is no longer with us to witness this milestone, every page of this dissertation carries his dreams, sacrifices, and prayers. One of the greatest sorrows of my life is that he was not here to see this achievement become a reality. I often wonder how proud he would have been, and I hope that, in a way beyond my understanding, he knows that his dream has finally been fulfilled. I miss him every day, and I dedicate this accomplishment to his memory with love, gratitude, and prayers that Allah grants him the highest place in Jannah. Every accomplishment in my academic journey belongs as much to them as it does

to me.

Words cannot express how grateful I am for my husband's unwavering love and support. Throughout this journey, Arsalan stood beside me through every high and every low. He celebrated every success with me, carried me through every disappointment, and constantly reminded me not to give up, even when I felt exhausted and overwhelmed. He understood my frustrations, fears, pain, and the emotional weight that often accompanies a Ph.D. Through every difficult moment, he held my hand and never allowed me to lose faith in myself. This achievement is as much his as it is mine. Thank you for being my greatest source of strength, my constant companion, and my best friend.

My precious daughter, Emaan, Allah's most beautiful gift to me, brought immeasurable joy, love, and strength into my life. She entered my life during this doctoral journey and transformed it in ways I could never have imagined. Before her, I measured life through accomplishments, deadlines, and goals. She taught me that life's greatest joys are found in the simplest moments—a smile, a laugh, and a tiny hand reaching for mine. She brought light into days that often felt overwhelming and reminded me why perseverance matters. Every time I looked at her, I found renewed strength to keep moving forward. One day, when she is old enough to read these words, I hope she will know that her presence made me stronger than I ever thought possible. She gave me not only happiness, but also a completely new understanding of love, gratitude, and the true meaning of success. For all that she has brought into my life, I will always be profoundly grateful.

I would also like to express my deepest gratitude to my eldest sister. She gave up many of her dreams so that her younger siblings could pursue theirs. She took responsibility for our family and generously supported my undergraduate education at one of the finest universities in Pakistan. As the youngest sibling, I received the privilege of pursuing higher education and eventually fulfilling our father's dream. I will always be indebted to my sister for her selflessness, generosity, and unconditional love.

I am profoundly grateful to Dr. Abid Aman Burki and Dr. Mushtaq Ahmed Khan, whose guidance and encouragement first inspired me to pursue doctoral studies. Their mentorship shaped my academic aspirations and gave me the confidence to embark on this journey. I

will always appreciate the faith they placed in me and the direction they provided during the early stages of my academic career.

Finally, I extend my sincere gratitude to every teacher, colleague, friend, family member, and well-wisher who supported me throughout this journey. Whether through advice, encouragement, friendship, prayers, or simple acts of kindness, each of you has contributed to this accomplishment in ways that words cannot fully express. If I have unintentionally omitted anyone by name, please know that your support has not been forgotten. I carry your kindness with me and will always remain grateful.

This dissertation is more than the completion of a doctoral degree. It is the culmination of years of faith, sacrifice, perseverance, and the unwavering love of those who stood beside me. I did not walk this path alone, and I will carry the lessons, generosity, and prayers of everyone who helped me throughout the rest of my life.

Alhamdulillah Rabbil-'Alamin.

Dedication

To my father, Emaddudin Muhammad Khan, whose dream became my purpose, and whose memory continues to inspire me every day.

To my mother, Salma Emad, for your endless sacrifices, unconditional love, and unwavering prayers.

To my husband, Arsalan Hussain, for believing in me, standing by my side, and never letting me give up.

To my daughter, Emaan, for bringing immeasurable joy, love, and meaning into my life.

Chapter 1

When Credit Tightens, Prices Take Off: Evidence from Jet Fuel Price Shocks in U.S. Airfares

1.1 Introduction

A standard feature of New Keynesian macroeconomic models is the assumption of a friction that slows the pass-through of shocks to consumer prices. [Aruoba et al. \(2026\)](#), referencing the overwhelming evidence that monetary policy shocks have effects on the real economy, write, “...real effects arise because prices do not fully adjust to nominal disturbances. How much they adjust depends on two distinct forces: nominal frictions that prevent firms from changing prices, and structural features of demand and costs that limit how much firms want to adjust prices even when they can.” The ability of modern macroeconomic models to match the data is dependent on assumptions about pass-through of shocks to consumer prices.

Most commonly, this pass-through is modeled by an explicit assumption of sticky prices. Time-dependent pricing, usually in the form of Calvo pricing, restricts the number of firms that can change price in a particular time period ([Calvo, 1983](#)). In contrast, state-dependent

pricing does not restrict the ability of firms to change prices, but it changes the firm’s profit-maximization problem so the decision of whether or not to change prices is dependent on the state of the economy. One way to derive state-dependent pricing is to impose a cost to changing prices, often referred to as a “menu cost” (Ball and Mankiw, 1995), so that it is only optimal for firms to change price if the current price deviates far enough from the optimal price conditional on changing price. It has been shown that the assumption of time-dependent versus state-dependent price rigidity is not innocuous (Alvarez et al., 2017).


This essay highlights the potential role of credit conditions in the pass-through of shocks to consumer prices. Specifically, we estimate a model that allows for asymmetry in the pass-through of jet fuel price shocks into airfares, where the impulse response function for airfares depends on national credit conditions. Credit conditions should matter for the pass-through of cost shocks to consumer prices, because the ability of a firm to temporarily absorb higher input costs is dependent on them having sufficient cash on hand or access to credit to make up for the unexpected increase in cost.

This effect should be particularly strong in the airline industry. U.S. airlines rely heavily on debt (Figure 1.1), while engaging in fuel price hedging (Morrell and Swan, 2006), which requires access to cash.¹ With normal credit conditions, airlines will have access to the cash needed for hedging, so an increase in the price of jet fuel will have little immediate effect on their marginal cost. If credit conditions are tight, the lack of access to credit may cause airlines to stop fuel price hedging. When that happens, a higher jet fuel price translates immediately into higher marginal cost for the airlines. That will in turn cause airlines to pass fuel cost shocks on to consumers in the form of higher ticket prices, because the tightening of credit conditions that reduced fuel price hedging also make it harder for airlines to absorb a cost shock. Similar reasoning should apply to other sectors of the economy. If the cost of changing prices is the fear of making customers upset (Blinder et al., 1998), the willingness

¹Air Carrier Financial data from [Schedule B1](#) of the Bureau of Transportation Statistics from of the Bureau of Transportation Statistics provide quarterly airline level financial information beginning in 1990. Four airlines with continuous coverage over the full sample period are retained. For each airline and quarter, total debt is constructed as the sum of current maturities of long term debt and long term debt. Quarterly debt and total assets are then aggregated to the annual level by summing across the four quarters, and the annual debt to asset ratio is calculated as annual total debt divided by annual total assets.

of a firm to absorb an increase in costs will be contingent on their cash flow situation. Firms that do not have easy access to credit will prefer to anger customers rather than take a chance of running out of cash.

Figure 1.1: *Debt-to-Asset Ratio by Airline (1990-2019)*



Consistent with a priori expectations, our empirical work finds that airlines increase fares immediately in response to a jet fuel cost shock when credit conditions are tight. Our preferred specification finds no pass-through of jet fuel price shocks into airfares in the first eight months after a shock when credit conditions are normal. It takes more than a year for airfares to fully adjust in that scenario, as expected, because fuel price hedging means the shock takes time to translate into a change in costs for the airlines. In contrast, when credit conditions are tight, a jet fuel price shock leads to an immediate increase in airfares, with the full adjustment happening in less than a year.

The remainder of this essay is organized as follows. Section 1.2 reviews the relevant literature and highlights the contribution of this study. Section 1.3 discusses fuel hedging practices in the airline industry and their role in mitigating exposure to jet fuel price volatility. Section 1.4 reviews the related literature on airline airfares and cost pass-through.

Section 1.5 describes the data and presents the empirical methodology used to estimate the effects of jet fuel price shocks on airfares under different credit conditions. Section 1.6 reports the empirical results and discusses their implications. Finally, Section 1.7 concludes with a summary of the main findings and discusses their policy implications.

1.2 Literature Review

To the best of our knowledge, this is the first empirical research to model the pass-through of cost shocks to consumer prices as dependent on credit conditions. We contribute to the literature on state-dependent pricing and time-varying cost pass-through, which we review below. We propose a new channel for state-dependent pricing and estimate an empirical model that captures it. There is a large empirical literature on nonlinear pass-through of cost shocks, but aside from the use of similar econometric techniques, there is limited overlap between this essay and that literature.

1.2.1 Time-Dependent versus State-Dependent Pricing

There are two methodological approaches to price setting in New Keynesian macroeconomic models. Models with time-dependent pricing assume firms set prices at fixed intervals, normally by imposing Taylor (1979) staggered price setting, or that there is a constant probability the firm updates prices each time period, most often by imposing Calvo (1983) price setting. Models with state-dependent pricing commonly impose a cost to the firm for changing prices, as in Ball and Mankiw (1994, 1995). Alvarez et al. (2017) show that the difference in modeling approaches matters primarily for large shocks. With state-dependent pricing, there is little price rigidity, since firms are better off to change prices quickly even if there is a cost to doing so. In contrast, by construction, price rigidity is time invariant with time-dependent pricing.

Recent empirical evidence provides direct support for the distinct implications of state and time dependent pricing for cost pass-through. Using survey data from UK firms, Bunn

[et al. \(2026\)](#) show that both pricing strategies are prevalent, but they generate systematically different responses to shocks. State dependent firms adjust prices more rapidly in response to cost shocks, while time dependent firms exhibit more delayed and smoother adjustment patterns. Importantly, these differences are not uniform: the gap in pass-through between the two groups is small for minor shocks but becomes substantially larger for large shocks, consistent with nonlinear state dependent pricing models. Complementary evidence based on randomized survey experiments, forecast errors, and local projections further confirms that state dependent firms exhibit faster and stronger pass-through of cost shocks, whereas time dependent firms adjust more gradually, with convergence occurring only after several quarters. Overall, these findings highlight that the nature of price setting behavior is a key determinant of the speed and nonlinearity of cost pass-through.

An implication of state-dependent pricing is that pass-through of cost shocks to consumer inflation is time-varying. Time variation in the pass-through of cost shocks has been studied in the case of exchange rates, motivated by the high volatility of the exchange rate. [Dainauskas \(2023\)](#) shows that the pass-through of exchange rates to the U.S. terms of trade fluctuated widely over the 1990-2018 time period.

1.2.2 Our Contribution

We estimate a model in which the pass-through of jet fuel price shocks to airfares depends on credit conditions. Two well-known characteristics of the airline industry are a reliance on debt and hedging against jet fuel price shocks, which are the main source of changes in variable costs for an airline. These two characteristics are related. Hedging requires access to cash, and given the heavy reliance of airlines on debt, their ability to participate in hedging depends on their access to credit. In tight credit conditions, airlines will reduce their hedging behavior and pass-through of jet fuel shocks to airfares will be faster. The importance of credit conditions for firm willingness to absorb cost shocks in the short run is likely to be important for firms in other sectors as well. We therefore propose to model prices as state-dependent, where the state is determined by credit conditions.

1.3 Fuel Hedging

Airlines are particularly vulnerable to jet fuel price shocks because fuel represents one of their largest operating expenses.² According to the Federal Aviation Administration (FAA), fuel accounts for approximately 53.5 percent of variable costs per block hour and 47.6 percent of total costs per block hour for U.S. airlines.³ Because fuel expenditures are so large, even modest increases in jet fuel prices can substantially raise operating costs. [Hornung \(2020\)](#) notes that Southwest Airlines estimated a one cent increase in the price of jet fuel would increase its annual fuel bill by roughly \$15 million, while American Airlines estimated that a 10 percent increase in fuel prices would raise fuel expenses by \$169 million in 2002 and \$649 million in 2008. These examples illustrate that relatively small changes in fuel prices can have economically significant effects on airline costs and profitability.

Given the importance of fuel costs and their sensitivity to changes in jet fuel prices, many airlines engage in fuel hedging to reduce their exposure to cost volatility. By locking in fuel prices through financial contracts, airlines seek to stabilize operating costs and cash flows and mitigate the financial risks associated with fluctuations in fuel prices ([Lim and Hong, 2014](#); [Merkert and Swidan, 2019](#); [Morrell and Swan, 2006](#)).

Airlines use a variety of derivative instruments to hedge fuel price risk, including forwards, futures, options, collars, and swaps ([Morrell and Swan, 2006](#)). Although these instruments differ in their contractual structure, they all can generate financing needs at different points in time. For example, futures require an initial margin and may involve additional cash payments through daily margin calls, options require an upfront premium, swaps typically involve periodic cash settlements, and forwards or collars can create payment obligations

²Estimates from the broader literature similarly indicate that jet fuel is one of the largest components of airline operating costs. [Koopmans and Lieshout \(2016\)](#) estimate that fuel accounts for 20 to 50 percent of total airline costs depending on the carrier, [Swidan and Merkert \(2019\)](#) report a range of 30 to 40 percent of operating costs, and the International Air Transport Association ([IATA, 2022](#)) estimates that aircraft fuel and oil account for approximately 28.7 percent of global airline operating costs.

³[FAA](#) offers detailed data in its Benefit-Cost Analysis Guidance for Airport Improvement Program Projects (2023), particularly in Section 4, which covers aircraft operating costs. Table 4-7 of this document categorizes costs into variable and fixed costs per block hour, with fuel and oil listed at \$2,760 per block hour. The total variable and total overall costs per block hour are \$5,161 and \$5,799, respectively. These figures imply that, as of June 2023, fuel costs accounted for approximately 53.5% of total variable costs and 47.6% of total costs per block hour.

at contract maturity or under specified market conditions. Consequently, maintaining a hedging program often requires access to liquidity or external financing, suggesting that an airline’s ability to hedge fuel price risk may depend on prevailing credit conditions.

How large are the liquidity requirements associated with fuel hedging in practice? [Swidan et al. \(2019\)](#) show that fuel hedging can generate large and time varying cash flow obligations when fuel prices decline and hedging positions become financially unfavorable. For example, they document that Southwest Airlines posted \$835 million in cash collateral and \$250 million in aircraft collateral in 2015, followed by \$301 million in cash collateral in 2016. After the decline in jet fuel prices following 2014, the gross fair value of Southwest’s outstanding fuel derivatives turned into a net liability of \$1.5 billion in 2015 and \$326 million in 2016, generating cash payments of \$577 million in 2015 and \$1.0 billion in 2016 to hedging counterparties. This highlights the strain that hedging can impose on airline balance sheets. Similarly, Delta Air Lines incurred net hedging losses of \$119 million in 2015 and \$38 million in 2016, while United Airlines and American Airlines stopped hedging altogether, citing collateral obligations that limited their ability to benefit from potential future declines in jet fuel prices.

Given these substantial liquidity requirements, what hedging intensity is optimal and how much do airlines actually hedge in practice? [Swidan et al. \(2019\)](#) show that the trade off between risk reduction and liquidity cost is non linear: modest hedge ratios provide substantial risk reduction with relatively limited collateral needs, while higher hedge ratios generate sharply increasing funding requirements that can outweigh additional hedging benefits. They further estimate that full hedging would require about \$1.36 billion in cash collateral, making it impractical even for large carriers, with risk reduction gains flattening beyond approximately a 70 percent hedge ratio. In practice, airlines typically hedge only a fraction of their expected fuel consumption, generally between one-third and two-thirds ([Morrell and Swan, 2006](#)), and evidence from airline annual reports reviewed by ([Swidan et al., 2019](#)) similarly shows that most carriers hedge around 50 to 60 percent of their fuel requirements, with full hedging never observed in practice.

Studies document several benefits of fuel hedging for airlines, primarily through its role

in reducing exposure to jet fuel price shocks and stabilizing financial performance. Evidence shows that hedging is associated with lower volatility in airline income and operating performance (Rao, 1999; Swidan and Merkert, 2019). A second strand of the literature links hedging to firm valuation and capital market outcomes, documenting higher Tobin’s Q, stronger firm value, higher stock returns, and lower total and idiosyncratic volatility among hedging airlines (Berkowitz et al., 2025; Carter et al., 2006; Samunderu et al., 2023). However, Morrell and Swan (2006) emphasize that the effect of hedging on profit volatility is not guaranteed, as its impact depends on the correlation between fuel prices and airline revenues; in some environments hedging reduces volatility when fuel prices and demand move in opposite directions, while in others its stabilizing effect may be limited when both move together.

1.4 Related Work on Airfares

Recent research suggests that the pass-through of jet fuel cost shocks to airfares is not uniform but instead depends on the economic environment in which airlines operate. In other words, airlines’ pricing responses appear to be state dependent, with the magnitude of pass through varying across different conditions. The notion of the relevant state differs across the literature. Some studies examine how pass through changes with market structure. Others focus on asymmetric responses to fuel cost increases and decreases, treating the direction of the cost shock itself as a distinct pricing state.

A prominent strand of the literature argues that the degree of competition shapes the extent of jet fuel cost pass-through to airfares. Airlines operating in less competitive markets have greater discretion over pricing decisions and may therefore respond differently to cost shocks than airlines facing intense competitive pressure (Koopmans and Lieshout, 2016). Consistent with this view, Jiang et al. (2022) show that market concentration, common ownership, and tacit collusion slow fare reductions following declines in fuel costs. Scotti and Volta (2018) similarly find that fare adjustments are stronger on routes with lower competition and reduced capacity, while Cannon and Watanabe (2020) report that airlines are

more likely to retain fuel cost savings when they possess greater pricing power. Collectively, these findings highlight competition as an important state variable governing fuel cost pass-through, as airlines with stronger market power have greater discretion over when and to what extent cost changes are reflected in airfares.

Another strand of the literature examines whether pass-through depends on the direction of the fuel cost shock. While [Kaufmann \(2017\)](#) and [Pal and Mitra \(2022\)](#) find evidence consistent with symmetric pass-through, [Jiang et al. \(2022\)](#), [Chuang \(2020\)](#), [Scotti and Volta \(2018\)](#), and [Cannon and Watanabe \(2020\)](#) report asymmetric pricing, with fuel cost increases generally being passed through more readily than fuel cost decreases. Overall, the evidence on whether the direction of the fuel cost shock affects pass-through remains inconclusive.

Taken together, the existing literature shows that the pass-through of jet fuel costs to airfares is not constant and may vary with market structure or the direction of the fuel cost shock. Related work also suggests that fuel hedging can smooth the transmission of fuel cost shocks to airfares. For example, [Gayle and Lin \(2021\)](#) show that airlines with greater hedging capacity exhibit smaller immediate pass-through of fuel cost shocks and document that widespread hedging contributed to unexpectedly high airfares during the 2014 oil price collapse. However, despite the widespread use of fuel hedging and the industry's heavy reliance on external financing, existing studies have not examined whether credit conditions influence airlines' ability or willingness to absorb fuel cost shocks before passing them through to consumers. This essay addresses that gap by studying credit conditions as a source of state dependent cost pass-through.

1.5 Methodology

We use a bivariate VAR model with airfares and jet fuel prices to model airfare dynamics:

$$A_t = \alpha_A + \sum_{i=1}^{13} \phi_{i,AA} A_{t-i} + \sum_{i=1}^{13} \phi_{i,AF} F_{t-i} + \psi \omega_t + \sum_{i=1}^{13} \theta_{i,AF} (\omega_{t-i} \cdot F_{t-i}) + \mu_{At} \quad (1)$$

$$F_t = \alpha_F + \sum_{i=1}^{13} \phi_{i,FA} A_{t-i} + \sum_{i=1}^{13} \phi_{i,FF} F_{t-i} + \mu_{Ft} \quad (2)$$

The airfare equation captures a smooth transition process for the coefficients on the lagged jet fuel terms, where the transition weight, ω_t , takes a value between 0 and 1, and is a function of credit conditions at time t . We model the transition weight as a logistic function of credit conditions

$$\omega_t = \frac{1}{1 + \exp(-\gamma(c_t - \bar{d}))} \quad (3)$$


where c_t is a measure of credit conditions at time t that serves as the transition variable, γ controls the speed of transition between regimes, and $\bar{d}$ is the location parameter.⁴

The main specification question is the choice of c_t . Our baseline specification using the year-over-year change in the Chicago Fed National Financial Conditions Credit Subindex (NFCI) as reported by FRED series [NFCICREDIT](#). The use of the year-over-year change in NFCI is motivated by a desire to capture a recent change in credit conditions sufficient to affect airline pricing behavior. Our robustness analysis checks the effect of changing the horizon of the change in NFCI as well as monotonic transformations of the level of NFCI.


The ability of our model specification to capture nonlinearity in this process depends on the ability of the transition function applied to cover the full range of probabilities (i.e., ω taking values over the full interval (0,1) rather than just a small portion of it) and many transition speeds. While in principle the logistic function covers both cases, the realization of the logistic function for a particular dataset might fail in one of the dimensions when

⁴Both parameters were chosen using a grid search over different values of $\bar{d}$ and γ in equation (1), selecting the pair that minimized the sum of squared residuals of the airfare equation. The optimal values found were $\bar{d} = -0.142$ and $\gamma = 1.64$.

Figure 1.2: *Transition Weights under Alternative Values of the Smoothing Parameter*



doing estimation, so we plot four examples calculated for our dataset showing the transition function coverage. Figures 1.2 and 1.3 show that this is not a concern for our dataset.⁵



The lag length of 13 was chosen by the AIC. To reduce the number of estimated parameters, our robustness analysis considers shorter lag lengths on the nonlinear term $\omega_{t-i}F_{t-i}$. It is plausible that credit conditions are only important for the response to recent jet fuel price changes.

The airfare variable is the year-over-year percentage change in the non-seasonally adjusted monthly Consumer Price Index for All Urban Consumers: Airline Fares in U.S. City Average (Series [CUUR0000SETG01](#)). The jet fuel variable is the year-over-year change in the non-seasonally adjusted monthly data for Kerosene-Type Jet Fuel Prices: U.S. Gulf Coast (Series [MJFUELUSGULF](#)). We take year-over-year percentage changes of both variables to remove seasonality. In addition, when computing impulse response functions, the point estimates of the response are equal to the cumulative responses for horizons up to 12 months, simplifying inference. Both series are deflated by the CPI as measured by the Consumer Price Index for

⁵For Figure 1.3, $\bar{d}$ is held fixed at -0.142 while only γ is varied, ensuring a common point of inflection across all curves.

All Urban Consumers: All Items in U.S. City Average ([CPIAUCNS](#)).

The starting dates of the data series are not the same. Although data on airfares and credit conditions is available from the early 1970s, data on jet fuel prices only begins in April 1990. As a result, this study uses a sample period starting from April 1990. The data series end in December 2019, yielding a total of 357 monthly observations. This cut-off is specified because the beginning of the COVID-19 pandemic had an unprecedented effect on the airline industry, where passenger travel has fallen almost instantaneously.

Figure 1.4: *Domestic Revenue Passenger Miles (RPMs)*




Figure 2.1 shows the effect of the COVID-19 pandemic on the airline industry by using the U.S. Domestic Revenue Passenger Miles ([RPMD](#)) from the FRED database, where the periods of the U.S. recessions are shaded in light grey based on the National Bureau of Economic Research (NBER) definitions. RPM, a typical indicator for the traffic in the airline industry, is defined as the number of paying passengers multiplied by the distance traveled. The index indicates a sharp decline in April 2020, where the RPM index dropped to about 96% below its December 2019 value. It took almost one and a half years for the travel demand to recover, where the RPM index reached its pre-pandemic value by July

2021. In order to avoid the effects of the extraordinary COVID-19 pandemic, the data is cut-off at December 2019.

1.5.1 Identification of Impulse Response Functions

The reduced form residuals μ_{At} and μ_{Ft} are assumed to be driven by the structural shocks to both variables, ε_{At} and ε_{Ft} , and a measure of US aggregate demand, $\varepsilon_{AD,t}$. Accounting for aggregate demand is important for the purposes of this study, because changes in the state of the economy will influence airfares and jet fuel prices in the same direction, potentially biasing the estimated response of airfares to jet fuel prices. We use a dummy variable that captures NBER recession dates for R_t . Our goal is to capture US-specific aggregate demand, not world aggregate demand.

$$\mu_{At} = \lambda_A \mu_{Ft} + \lambda_B R_t + \varepsilon_{At} \quad (4)$$

$$\mu_{Ft} = \lambda_C \varepsilon_{At} + \lambda_D R_t + \varepsilon_{Ft} \quad (5)$$

Our proxy SVAR framework requires an external instrument that is correlated with ε_{Ft} but not with ε_{At} . We use the precautionary oil demand shock of [Kilian \(2009\)](#). These shocks capture changes in jet fuel prices driven by unexpected shifts in demand for global crude oil, typically in response to news about future supply shortages.

Following [Ramey \(2016\)](#), We estimate equation (4) using the 2SLS method. In the first stage, the reduced-form jet fuel price residuals (μ_{Ft}) are regressed on the precautionary demand shocks and aggregate demand shocks. The fitted values from this stage are then used as instruments in the second stage, where the reduced-form airfare residuals (μ_{At}) are regressed on the predicted jet fuel prices along with aggregate demand shocks. We use the residuals from equation (4) as a proxy for the reduced-form airfare residuals (μ_{At}) when estimating equation (5) with the OLS.

For the purpose of simulating the impulse responses of airfares to jet fuel price shocks,

the main focus is on the coefficient λ_A , as it reflects the immediate effect of jet fuel prices on airfares. The Proxy SVAR estimation produces a value of $\lambda_A = 0.14$, indicating that a one-unit shock in jet fuel prices leads to an immediate increase of 0.14 in airfares. This can be properly interpreted as: a one percentage-point increase in the year-over-year change in jet fuel prices leads to an immediate 0.14 percentage points increase in the year-over-year change in airfares, within the same month.

1.5.2 Impulse Response Functions (IRFs)

After identifying the impact effect, the next step is to forecast the future path of airfares. To this end, we employ the local projection approach proposed by (Jordà, 2005). Specifically, we consider a 12 month horizon and estimate equation (6) using OLS separately for each horizon ($h = 1, \dots, 12$).

$$A_{T+h} = \alpha_0 + \sum_{i=0}^2 \phi_{i,A} A_{T-i} + \sum_{i=0}^2 \phi_{i,F} F_{T-i} + \psi \omega_T + \theta_{\omega F} (\omega_T \cdot F_T) \quad (6)$$

We include two lags of both A_T and F_T in the local projection regressions. This allows us to use a wild bootstrap procedure for inference, which is robust to heteroskedasticity and conditional heteroskedasticity but not to serial correlation; serial correlation is addressed by including lagged dependent and explanatory variables in the specification.

For each horizon, we forecast airfares h periods ahead under two scenarios: one with a jet fuel price shock at time T (denoted S_T), and one without (denoted N_T), each evaluated under tighter credit (H) and normal credit (L) conditions. In the shock scenario, a one-unit increase in the jet fuel price F_T leads to a 0.14-unit increase in airfares A_T . In the no-shock scenario, the change in F_T and hence in A_T is zero. For both sets of forecasts, the lagged variables A_{T-1} , A_{T-2} , F_{T-1} , and F_{T-2} are set to zero.

We define tighter credit (H) as credit conditions one standard deviation above the mean ($\bar{\omega}^H = \bar{\omega} + \sigma_\omega$) and normal credit (L) as one standard deviation below the mean ($\bar{\omega}^L = \bar{\omega} - \sigma_\omega$). Using these definitions, the forecasts under both shock and no-shock scenarios for tight and normal credit conditions can be expressed as follows:

$$\begin{aligned}
\hat{A}_{T+h}^H | S_T &= \hat{\alpha}_0^H + 0.14 \times \hat{\phi}_A^H + \hat{\phi}_F^H + \psi^H \bar{\omega}^H + \hat{\theta}_{\omega_F}^H(\bar{\omega}^H) \\
\hat{A}_{T+h}^H | N_T &= \hat{\alpha}_0^H + \psi^H \bar{\omega}^H \\
\hat{A}_{T+h}^L | S_T &= \hat{\alpha}_0^L + 0.14 \times \hat{\phi}_A^L + \hat{\phi}_F^L + \psi^L \bar{\omega}^L + \hat{\theta}_{\omega_F}^L(\bar{\omega}^L) \\
\hat{A}_{T+h}^L | N_T &= \hat{\alpha}_0^L + \psi^L \bar{\omega}^L
\end{aligned} \tag{7}$$

The impulse response functions (IRFs) for tight and normal credit conditions are then computed for each horizon as:

$$\begin{aligned}
IRF_{T+h}^H &= \hat{A}_{T+h}^H | S_T - \hat{A}_{T+h}^H | N_T \\
IRF_{T+h}^L &= \hat{A}_{T+h}^L | S_T - \hat{A}_{T+h}^L | N_T
\end{aligned} \tag{8}$$

We then compute the cumulative sum of the IRFs over all forecast horizons.

1.5.3 Confidence Intervals for Impulse Response Functions

To obtain confidence intervals for the IRFs, we apply the wild bootstrap with 10,000 replications. For each bootstrap replication we generate a new dataset by combining the model-fitted values for airfares ($\hat{A}_t$) from equations (1) with its estimated residuals ($\hat{\mu}_{At}$), using:

$$\tilde{A}_t = \hat{A}_t + z \hat{\mu}_{At} \tag{9}$$

Here, z is a random variable that takes values from $\{-1, 1\}$. In each replication, equations (6) is re-estimated using the simulated dataset generated for that replication. The forecasts, IRFs, and cumulative IRFs are then computed over the entire forecast horizon. The median cumulative IRFs are presented in the results, along with the 2.5th and 97.5th percentile

values to represent the 95% confidence intervals.

1.6 Empirical Results

Figure 1.5 depicts the response of airfares to jet fuel price shocks over 12 months under varying credit conditions, produced by estimating the local projection model specification in equation (6). The solid red line represents the cumulative median response under tight credit, while the blue line corresponds to normal credit conditions. Shaded bands around each line are 95% confidence intervals. This figure illustrates how the impact of fuel price shocks builds up over time in each credit environment.

The IRFs show that airfares respond immediately to jet fuel price shocks when credit conditions are tight, indicating that airlines begin passing higher fuel costs on to consumers within the same month as the shock. The pass-through effect increases steadily over the following nine months and then continues to rise at a slower pace, reaching a cumulative peak of approximately 0.67 percentage points after 12 months. This implies that a 1 percentage point increase in jet fuel prices results in a cumulative 0.67 percentage point increase in airfares one year after the shock. The estimated response remains statistically significant throughout the forecast horizon.


The most striking result is the delayed response of airfares to a jet fuel price shock under normal credit conditions. During the first eight months after the shock, the estimated response is not statistically different from zero, suggesting that airlines largely absorb the increase in fuel costs rather than immediately passing it on to consumers. By the ninth month, the cumulative response reaches approximately 0.117 percentage points and becomes statistically significant, implying that a 1 percentage point increase in jet fuel prices leads to a cumulative 0.117 percentage point increase in airfares.

One possible explanation for the delayed pass-through under normal credit conditions is the widespread use of fuel hedging by airlines. Fuel hedging contracts allow airlines to lock in fuel prices for several months, insulating operating costs from contemporaneous increases in spot jet fuel prices. Since these contracts typically have maturities ranging from six months

to one year (Morrell and Swan, 2006), airlines need not immediately adjust ticket prices following a fuel price shock. The finding that airfare responses remain statistically insignificant for approximately the first eight months after the shock is therefore consistent with the duration of typical hedging contracts. As existing hedge positions expire, airlines become increasingly exposed to higher fuel costs and begin passing these costs on to consumers, suggesting that hedging delays rather than eliminates cost pass-through.

Credit conditions likely influence the extent to which airlines can rely on this hedging mechanism. As discussed earlier, maintaining fuel hedging programs often requires substantial liquidity to meet margin calls, collateral requirements, option premiums, and other cash flow obligations. When credit conditions are normal, airlines have greater access to external financing and are better able to meet these funding needs while absorbing temporary increases in fuel costs. In contrast, when credit conditions tighten, financing becomes less available or more expensive, reducing airlines' ability to sustain hedge positions and smooth cash flow fluctuations. Consequently, airlines face stronger incentives to pass higher fuel costs on to consumers immediately, resulting in faster and larger cost pass-through.

Figure 1.5: *Response of Airfares to Jet Fuel Shock*




Overall, the analysis highlights the crucial role that credit conditions play in shaping how airlines respond to fuel price shocks. Tighter credit constraints lead to a faster and higher pass-through of costs to consumers, reflecting the airlines' reduced ability to absorb shocks due to limited financing options. In contrast, when credit is more accessible and affordable, airlines are better positioned to absorb fuel cost shocks without passing them aggressively on to consumers. This indicates that tighter financial conditions amplify the transmission of input cost shocks to consumer prices, suggesting that credit constraints play a key role in shaping inflation dynamics.

1.6.1 Robustness to Lag Length Selection

The baseline results presented above are based on the smooth transition vector autoregressive (STVAR) model specified in Equations (1) (2) and (3), estimated with 13 lags as selected by the AIC. An alternative lag selection criterion, the SIC, instead recommends a more parsimonious specification with 3 lags. To assess the sensitivity of our findings to the choice of lag length, we reestimate the STVAR model using 3 lags while continuing to use the year-over-year change in credit conditions as the transition variable. Comparing the results from the two specifications allows us to evaluate whether the estimated effects are robust to the choice of lag length.

Overall, the results are consistent with the baseline specification in showing that jet fuel cost pass-through is larger under credit tightening than under normal credit conditions. In the first three months following the shock, the responses under the two regimes are statistically indistinguishable. After this initial period, pass-through increases more rapidly under tight credit conditions, while it rises more gradually under normal credit conditions. As a result, the difference between the two responses widens over time and reaches its maximum approximately nine months after the shock. Beyond this point, pass-through continues to increase under tight credit conditions, but at a slower rate, whereas it continues to rise more noticeably under normal credit conditions, leading to a partial narrowing of the gap. This dynamic pattern is consistent with the baseline results, highlighting stronger and faster cost

Figure 1.6: *Response of Airfares to Jet Fuel Shock - VAR(3)*



pass-through under credit tightening and more delayed transmission when credit conditions are loose.

Quantitatively, under tight credit conditions, airfares increase by approximately 0.67 percentage points one year after a 1 percentage point increase in jet fuel prices. Under normal credit conditions, the corresponding response is substantially higher at around 0.37 percentage points. While both responses are statistically different from zero over the full horizon, they become statistically different from each other after approximately three months following the shock.

When the lag length is reduced to three, the responses under normal credit conditions increase more rapidly relative to the baseline specification, reflecting the model's reduced ability to capture delayed adjustment dynamics, as a shorter lag structure cannot fully accommodate the gradual, multi-month transmission of fuel cost shocks into airfares. In contrast, the tight credit responses remain largely unchanged, consistent with the more immediate pass-through of cost shocks in this regime. Overall, these results suggest that the baseline specification with a longer lag structure is particularly important for capturing

the gradual adjustment under normal credit conditions, whereas the dynamics under tight credit conditions are less sensitive to the choice of lag length.

1.6.2 Robustness to the Measurement of Credit Conditions

In the baseline specification, credit conditions are measured using a year-over-year change. This subsection examines the sensitivity of the results to alternative ways of constructing this variable. In particular, we re estimate the model using a six month over six month change, three month over three month change and a month over month change in credit conditions, while keeping all other aspects of the specification unchanged, including the use of 13 lags in the STVAR. This exercise allows us to assess whether the estimated relationship between credit conditions and cost pass-through is driven by the choice of measurement window or remains robust across alternative definitions of credit dynamics.


Figure 1.7 presents the results from the STVAR specification in which the transition weights are constructed using the six month over six month change in credit conditions. Overall, the findings are consistent with the baseline specification and continue to support stronger pass-through of jet fuel cost shocks to airfares when credit conditions tighten. Under tight credit conditions, a one percentage point increase in jet fuel prices leads to a cumulative increase of approximately 0.76 percentage points in airfares after twelve months.

The estimated response under normal credit conditions also remains consistent with the baseline results and indicates substantially lower cost pass-through. For roughly the first six months after the shock, the responses under the two credit regimes are not statistically different from one another. Beyond this point, however, the responses begin to diverge, with pass-through becoming considerably stronger under tight credit conditions. After twelve months, a one percentage point increase in jet fuel prices raises airfares by only about 0.36 percentage points under normal credit conditions, compared with 0.76 percentage points under tight credit conditions.

Figure 1.8 presents the results from the STVAR specification in which the transition weights are constructed using the three month over three month change in credit conditions.

Overall, the findings are consistent with the baseline specification and continue to support stronger pass through of jet fuel cost shocks to airfares when credit conditions tighten. Under tight credit conditions, a one percentage point increase in jet fuel prices leads to a cumulative increase of approximately 0.71 percentage points in airfares after twelve months.


Figure 1.7: *Response of Airfares to Jet Fuel Shock (6 months Credit Change)*



Although the estimated response under normal credit conditions is smaller than under tight credit conditions, the two responses are not statistically different from each other until approximately nine months after the shock. Relative to the baseline specification, the divergence in pass-through across regimes is less pronounced. After twelve months, a one percentage point increase in jet fuel prices raises airfares by about 0.56 percentage points under normal credit conditions, compared with 0.71 percentage points under tight credit conditions.

We also estimated the model using month over month changes in credit conditions to construct the transition weights. Under this specification, the estimated transition weights were essentially constant, ranging from approximately 0.99996 to 0.99999, implying virtually no regime switching. As a result, the model effectively collapses to a linear specification. This

Figure 1.8: *Response of Airfares to Jet Fuel Shock (3 months Credit Change)*



outcome is consistent with the notion that very short run fluctuations in credit conditions are unlikely to generate economically meaningful changes in firms' pricing behavior. Instead, the nonlinear effects documented in the baseline and alternative specifications emerge only when credit conditions are measured over longer horizons that better capture persistent changes in financing constraints. The absence of regime variation under the month over month definition therefore reinforces the interpretation that the estimated non-linearities are driven by sustained changes in credit conditions rather than by high frequency noise in the transition variable.

1.7 Conclusion & Implications

This essay proposes credit conditions as a new source of state dependent cost pass-through. We argue that a firm's ability to absorb temporary increases in input costs depends not only on the size of the shock but also on its access to external finance. Using the airline industry as a case study, we examine whether the pass-through of jet fuel price shocks to airfares

varies with national credit conditions.

Our proposed mechanism operates through airlines' reliance on debt financing and fuel hedging. When credit conditions are favorable, airlines have greater access to the liquidity needed to hedge fuel costs and temporarily absorb input price shocks, reducing the need for immediate fare adjustments. When credit conditions tighten, financing constraints limit this ability, causing increases in jet fuel costs to translate more quickly into higher airfares.

Consistent with this mechanism, our empirical results reveal markedly different pricing responses across credit regimes. During periods of tight credit, airlines raise fares almost immediately following a jet fuel price shock. In contrast, under normal credit conditions, the response remains negligible for roughly eight months before becoming positive and statistically significant. We interpret this delayed adjustment as evidence that, when financing conditions are normal, fuel hedging temporarily insulates airlines from changes in jet fuel costs. As hedging contracts expire, however, the increase in fuel costs is gradually reflected in ticket prices. These findings suggest that fuel hedging delays, rather than eliminates, the pass-through of jet fuel price shocks to airfares. Our findings are also robust to alternative empirical specifications. Similar patterns emerge under different lag selections and when credit conditions are measured using alternative definitions, indicating that the estimated relationship between credit conditions and fuel cost pass-through is not driven by a particular modeling choice.

Our results also have broader implications beyond the airline industry. In many sectors, firms do not adjust prices immediately in response to cost shocks because they rely on contracts, hedging arrangements, or temporary absorption of higher costs to preserve customer relationships. The ability to pursue these strategies depends on access to financing. Our findings suggest that credit conditions can therefore influence the speed with which cost shocks are transmitted to consumer prices, making effective price rigidity state dependent rather than constant over time. Incorporating this channel into macroeconomic models may improve our understanding of inflation dynamics and the transmission of monetary policy, particularly during periods of financial stress.

Furthermore, our findings have direct policy implications. If cost shocks are passed

through to prices more rapidly when credit conditions are tight, then the short run aggregate supply curve becomes steeper during such periods. As a result, expansions in aggregate demand, whether driven by economic shocks or policy interventions, are more likely to generate inflation and less likely to increase real output when credit conditions are tight than when credit conditions are favorable.

Chapter 2

Is Cost Pass-Through in the Airline Industry Seasonal?

2.1 Introduction

Jet fuel costs constitute one of the largest and most volatile components of airline operating expenses and represent a substantial share of total costs in the U.S. airline industry. Although estimates of the exact cost share vary across sources, existing evidence consistently highlights the central importance of jet fuel expenditures for airlines. For example, [Koopmans and Lieshout \(2016\)](#) estimate that jet fuel accounts for between 20 and 50 percent of total airline costs, depending on the carrier. At the global level, aircraft fuel and oil account for 28.7 percent of total airline operating expenses, ranking among the largest cost categories alongside depreciation and crew salaries (IATA).¹ In the United States, Federal Aviation Administration (FAA) data indicates that fuel costs account for approximately 53.5 percent of total variable costs per block hour and 47.6 percent of total costs per block hour². The economic significance of fuel price movements is illustrated by the case of Southwest Airlines,

¹[IATA, 2022](#) – Unveiling the biggest airline costs.

²[FAA](#) offers detailed data in its Benefit-Cost Analysis Guidance for Airport Improvement Program Projects (2023), particularly in Section 4, which covers aircraft operating costs. According to the FAA (2023, Table 4-7), fuel and oil account for approximately 53.5% of variable aircraft operating costs and 47.6% of total operating costs per block hour.

which consumed about 50 million barrels of jet fuel in 2019; a \$10 per barrel increase in fuel prices, if not passed through to fares, would reduce the airline’s net income by roughly \$500 million (Wolter et al., 2021). As a result, fluctuations in jet fuel prices play a central role in shaping airline pricing decisions.

How airlines respond to changes in fuel costs is important for several reasons. (Pal and Mitra, 2022) highlight three key considerations. First, fuel expenditures directly affect the financial viability of airlines, given their large share in total operating costs. Second, airline pricing behavior is closely monitored by regulators to ensure that airlines do not coordinate prices or exercise market power in ways that harm competition. Third, policymakers are concerned with consumer welfare and pay close attention to whether airlines pass fuel cost savings on to passengers when fuel prices fall. Interest in this issue increased sharply following the large decline in oil prices between June 2014 and June 2016, when crude oil prices fell by about 57 percent while average airfares declined by only around 9 percent, raising concerns about asymmetric pricing behavior (Pal and Mitra, 2022). These concerns led the U.S. Department of Justice to investigate airlines in 2015 for their alleged failure to pass fuel cost reductions on to consumers (Wolter et al., 2021).

Only a few studies examine why the sharp decline in fuel costs between 2014 and 2016 was not fully passed through to airfares. Kaufmann (2017) investigates whether asymmetric pricing behavior can explain this pattern but finds no evidence that asymmetries in the relationship between jet fuel prices and airfares account for the muted fare response. In contrast, Cannon and Watanabe (2020) document asymmetric pass-through and attribute the weak pass-through to airline pricing power, arguing that carriers passed fuel cost increases on to consumers while retaining much of the gains from fuel cost declines. This debate has motivated a growing body of research on the asymmetric transmission of jet fuel price shocks to airfares using different datasets, levels of aggregation, data frequencies, and empirical strategies. The resulting evidence is mixed, suggesting that fuel cost pass-through depends on the underlying economic and market environment. This naturally raises a related question that has received little attention in the literature: does the transmission of fuel cost shocks to airfares vary systematically across seasons?

While several studies account for seasonality, they typically remove seasonal variation either by using seasonally adjusted data (Cannon and Watanabe, 2020; Kaufmann, 2017; Scotti and Volta, 2018) or by including seasonal fixed effects (Chuang, 2020; Jiang et al., 2022), rather than directly examining whether fuel cost pass-through itself differs across seasons. These approaches allow average fares to vary over the year but impose a common pass-through relationship between fuel costs and airfares throughout the sample period. However, the mechanisms emphasized in the literature as drivers of pass-through, including pricing power, demand conditions, and capacity constraints, are likely to vary systematically between peak and off-peak travel seasons. If so, estimates that restrict pass-through to be uniform may conceal meaningful seasonal heterogeneity in how fuel cost shocks are transmitted to airfares.

Moreover, seasonality also has a spatial dimension in addition to its temporal one. At the route level, seasonal patterns differ across markets and across time of year, implying that airlines operating on different routes may face distinct demand and competitive environments when the same fuel cost shock occurs. As a result, the pass-through of identical fuel cost shocks may vary not only across seasons but also across routes, depending on whether the shock coincides with peak or off-peak travel conditions in that specific market. These considerations raises an important question: does the asymmetry documented in the literature persist once fuel cost pass-through is allowed to vary across seasons?

In this essay, we address the first question of whether cost pass-through differs systematically across seasons. We estimate a two-variable structural vector autoregressive (SVAR) that includes airfares and jet fuel prices. The airfare equation incorporates seasonal indicators to capture systematic differences in average fare levels across the year and interacts jet fuel price changes with these indicators to allow pass-through to vary between peak and off-peak travel periods. Identification of fuel price shocks follows a proxy SVAR approach in the spirit of Ramey (2016), where jet fuel prices are instrumented using Kilian’s (2009) precautionary oil demand shocks and indicators for NBER recessions. The estimated impact effects from the SVAR are then used in a lag-augmented local projection framework to trace the response of airfares to a unit jet fuel price shock over time.

The baseline results indicate that pass-through of jet fuel cost shocks to airfares in the peak season, defined as July, is immediate, reaching its maximum within approximately three months after the shock and then declining gradually over time. In contrast, in the off-peak season, defined as February, pass-through is more muted in the initial months following the shock and increases more gradually, reaching its peak only after nine months. Notably, the long-run magnitude of pass-through is similar across July and February; the key difference lies in the timing of adjustment rather than the eventual level of pass-through.

To complement the baseline analysis, we adopt an alternative timing based approach that examines whether the transmission of fuel cost shocks depends on the interaction between the season in which the shock occurs and the season in which airfares are observed. The motivation is that airline tickets are typically purchased in advance, so prices observed in a given month may reflect cost shocks originating several months earlier. We therefore compare a high travel intensity pricing month to a low season shock with a low travel intensity pricing month to a high season shock. In particular, we compare the response of airfares in August, a peak travel month, to a jet fuel price shock in February with the response in January to a shock in July. This design is chosen so that the lag between shock and pricing month is held constant at seven months in both cases. The results show a clear asymmetry in pass-through across seasonal states. The response of airfares in August to a February shock is 0.10, compared to 0.07 for the response in January to a July shock, indicating stronger transmission when the pricing month occurs in a high demand period.

We also investigate whether our empirical results depend on the environment for oil prices. In other words, we ask if the pass-through depends on what is happening in oil markets. We use the net oil price increase measure of Hamilton (1996) to distinguish between times with normal oil price behavior and times of extreme oil price spikes. The results show that pass-through does not meaningfully change when the economy is experiencing a large oil price spike.

Finally, to show the importance of allowing for asymmetric pass-through in practice, we calculate the estimated responses of airfares to the jet fuel price shocks caused by the recent military conflict between Iran and the United States. We show how the behavior of airfares

would differ if the jet fuel price shocks had occurred in other months of the year, as opposed to March and April. The results show that if these shocks had occurred during peak travel periods, there would have been substantially faster pass-through of fuel costs to airfares, whereas if those shocks had hit during low demand periods, there would have been a much slower but more persistent adjustment of airfares. This case study highlights the importance of accounting for the timing of jet fuel price shocks.

The remainder of this essay is organized as follows. Section 2.2 reviews the relevant literature and identifies the research gap. Section 2.3 describes the methodology and data. Section 2.4 presents the empirical results, and Section 2.5 concludes.

2.2 Literature Review

Several studies have noted that the sharp decline in oil prices between 2014 and 2016 was not accompanied by a comparable decline in U.S. airfares. This outcome is surprising because fuel is one of the largest operating costs for airlines. [Kaufmann \(2017\)](#) reports that the price of crude oil fell by approximately 57 percent between June 2014 and May 2016, while airline fares, measured using the CPI for airline fares, declined by only about 12 percent over the same period. Using a different measure of airfares, [Pal and Mitra \(2022\)](#) similarly find that a 57 percent decline in crude oil prices was associated with only a 9 percent reduction in airfares. [Gayle and Lin \(2021\)](#) report that crude oil prices fell from \$111.8 per barrel in June 2014 to \$38.01 per barrel in December 2015, a decline of about 66 percent. Financial reports show that the four major U.S. airlines experienced more than a 30 percent reduction in fuel expenses from 2014 to 2015, yet average airfares declined by less than 5 percent over the same period.³ Taken together, these studies suggest that airlines retained a substantial share of the cost savings generated by the collapse in oil prices. Why were lower fuel costs not more fully reflected in airfares during this period, despite the substantial cost savings enjoyed by airlines?

³For example, American and Delta reported more than 40% jet fuel expense saving, while United and Southwest had over 30% reduction in fuel costs, according to these airlines' 10 K reports.

Among the studies on airline fuel cost pass-through, [Kaufmann \(2017\)](#), [Cannon and Watanabe \(2020\)](#), and [Gayle and Lin \(2021\)](#) directly examine why changes in fuel costs are not fully reflected in airfares. [Kaufmann \(2017\)](#) investigates several explanations for the weak fare response during the 2014 to 2016 oil price collapse. Using monthly industry level data, he finds no evidence that the modest decline in airfares can be explained by slow adjustment of fares to fuel costs, asymmetric fare responses to oil price changes, or asymmetric price transmission within oil markets. Instead, his results suggest that neither conditions within the oil market nor the pricing relationship between fuel costs and airfares can fully explain the limited fare response. In contrast, [Cannon and Watanabe \(2020\)](#) argue that the period from 2009 to mid-2016 was characterized by strong airline pricing power. Their results indicate that airlines passed fuel cost increases on to consumers while retaining much of the savings generated by fuel cost declines. [Gayle and Lin \(2021\)](#) suggest that one reason lower oil prices may not translate into lower airfares is fuel hedging. Their theoretical and empirical results indicate that higher fuel hedging reduces the pass-through of crude oil price changes to fares, implying that airlines with more extensive hedging are less likely to pass fuel cost savings on to consumers. They further show that market competition and route characteristics also influence the extent of pass-through. Together, these studies highlight that the modest decline in airfares during the 2014 to 2016 fuel price collapse remains only partially understood.

More generally, the literature provides mixed evidence on whether fuel cost changes are transmitted symmetrically to airfares. Using a multiple threshold nonlinear autoregressive distributed lag model, [Pal and Mitra \(2022\)](#) find a long-run cointegrating relationship between airfares and both crude oil and jet fuel prices, but report no evidence that airfares respond more rapidly to fuel price increases than to fuel price decreases. In contrast, several studies document asymmetric pass-through. Using U.S. quarterly fare data from 2001 to 2016, [Jiang et al. \(2022\)](#) find evidence of asymmetric pricing behavior. Similarly, [Chuang \(2020\)](#), using firm-level quarterly airline fare data from 2006 to 2016, finds that positive fuel cost shocks are transmitted to airfares more rapidly than negative shocks, with most adjustments occurring within the same quarter. Evidence of asymmetric fare adjustment is

also reported outside the United States. Using route-level European airline data from 2010 to 2015, [Scotti and Volta \(2018\)](#) find that fare responses to fuel cost changes are asymmetric. Taken together, these studies suggest that the nature of fuel cost pass-through remains unsettled, with some evidence supporting symmetric adjustment and other evidence pointing to asymmetric transmission of fuel cost shocks to airfares. This mixed evidence raises the broader question of what factors determine the extent of fuel cost pass-through in the airline industry.

Several studies argue that fuel cost pass-through depends on market conditions rather than reflecting a fixed relationship between fuel costs and airfares. [Koopmans and Lieshout \(2014, 2016\)](#) emphasize that pass-through varies with market structure, demand conditions, and the competitive environment in which airlines operate. A recurring finding in the literature is that competition and pricing power play a central role in shaping fare responses to fuel cost shocks. [Jiang et al. \(2022\)](#) show that market concentration, common ownership, and tacit collusion slow fare reductions following fuel cost declines. Similarly, [Scotti and Volta \(2018\)](#) find that asymmetric fare adjustments are stronger on routes with lower competition and reduced capacity. [Cannon and Watanabe \(2020\)](#) find that airlines were more likely to retain fuel cost savings during periods of strong pricing power. Together, these studies suggest that airlines facing less competitive pressure have greater ability to pass fuel cost increases on to consumers while retaining the benefits of fuel cost reductions. However, market power alone may not fully explain asymmetric pass-through. [Tappata \(2006\)](#) shows that asymmetric pricing can arise even in competitive markets when consumers face search costs and imperfect information, providing an alternative explanation for why fares may respond more quickly to fuel cost increases than to fuel cost decreases.

The factors discussed above suggest that the extent of fuel cost pass-through may vary across different economic environments. This raises a related question that has received little attention in the literature: does the transmission of fuel cost shocks to airfares differ systematically across seasons? A potential reason to expect seasonal variation comes from the role of capacity constraints in airline pricing. [Escobari \(2013\)](#) argues that, unlike many industries where observed input prices are the primary source of cost variation, an important

cost component in the airline industry is the opportunity cost of selling limited seat capacity. Because airlines set capacity in advance and sell seats under uncertain demand, the cost of selling an additional seat depends on the likelihood that future demand will exceed available capacity.

To capture this cost component, [Escobari \(2013\)](#) measures costs using relative inventory levels, defined as the ratio of seats already sold to total aircraft capacity. Holding the number of days before departure fixed, higher realized demand implies fewer remaining seats and therefore a higher shadow cost of capacity for the next seat sold. The study finds that the response of prices to these cost shifts varies across the selling season. In particular, positive cost shocks are more likely to increase prices than negative cost shocks are to reduce them, and the effect of cost shocks is stronger during the summer travel season. These findings suggest that the pass-through of costs to airfares may itself vary across seasons. However, while [Escobari \(2013\)](#) demonstrates seasonal variation in the pass-through of capacity-related costs, no study has directly examined whether the pass-through of fuel cost shocks to airfares also differs systematically between high-demand and low-demand seasons.

While several studies account for seasonality, they generally do so by removing seasonal variation rather than examining whether fuel cost pass-through itself differs across seasons. [Kaufmann \(2017\)](#), [Cannon and Watanabe \(2020\)](#), and [Scotti and Volta \(2018\)](#) use seasonally adjusted data, while [Jiang et al. \(2022\)](#) and [Chuang \(2020\)](#) include quarter-specific fixed effects. These approaches allow average fares to vary over the year but maintain a common pass-through relationship between fuel costs and airfares throughout the sample period. Yet the factors identified in the literature as shaping pass-through, including pricing power, demand conditions, and capacity constraints, are likely to vary systematically between peak and off-peak travel seasons. If so, estimates based on a common pass-through relationship may mask important seasonal differences in how fuel cost shocks are transmitted to airfares. This raises an important question: does the asymmetry documented in the literature persist once fuel cost pass-through is allowed to vary across seasons?

2.3 Methodology

We use a bivariate VAR model with airfares and jet fuel prices to model airfare dynamics:

$$\begin{aligned}
A_t = & \alpha_A + \pi_1 P + \pi_2 O + \sum_{i=1}^{13} \beta_{AA,i} A_{t-i} + \sum_{i=1}^{13} \beta_{AF,i} F_{t-i} \\
& + \sum_{i=1}^{13} \gamma_{AF,i} (P \cdot F_{t-i}) + \sum_{i=1}^{13} \delta_{AF,i} (O \cdot F_{t-i}) + \mu_{A,t}
\end{aligned} \tag{1}$$

$$F_t = \alpha_f + \sum_{i=1}^{13} \beta_{FA,i} A_{t-i} + \sum_{i=1}^{13} \beta_{FF,i} F_{t-i} + \mu_{F,t} \tag{2}$$

The airfare equation includes seasonal indicators to account for differences in average fare levels across the year and interacts jet fuel price changes with these indicators to allow fuel cost pass-through to vary between peak and off peak travel periods. The lag length of the model is set to thirteen, based on the Akaike Information Criterion (AIC).

We use passenger data from the [TranStats](#) database (T-100 Domestic Market, U.S. Carriers) to identify seasonal variation in air travel demand and to classify peak and off-peak periods based on observed passenger volumes. Because the data are reported separately for each airline and month, we aggregate passenger counts across all carriers within each month and year to construct a monthly time series of total passenger traffic. The peak season (P) is defined as the two consecutive months with the highest passenger volumes, measured by the average monthly share of annual passenger traffic. This share is calculated as the number of passengers in each month relative to total annual passenger traffic in the corresponding year. The peak months are consistently July and August. The off-peak season (O) is defined analogously as the two consecutive months with the lowest passenger volumes, which consistently occur in January and February.

The airfare variable is the year-over-year percentage change in the non-seasonally adjusted monthly Consumer Price Index for All Urban Consumers: Airline Fares in U.S. City Average (Series [CUUR0000SETG01](#)). The jet fuel variable is the year-over-year change in the non-

seasonally adjusted monthly data for Kerosene-Type Jet Fuel Prices: U.S. Gulf Coast (Series [MJFUELUSGULF](#)). We take year-over-year percentage changes of both variables to remove seasonality. In addition, when computing impulse response functions, the point estimates of the response are equal to the cumulative responses for horizons up to 12 months, simplifying inference. Both series are deflated by the CPI as measured by the Consumer Price Index for All Urban Consumers: All Items in U.S. City Average ([CPIAUCNS](#)).

The starting dates of the data series differ across variables. Although airfare data are available from the early 1970s, passenger and jet fuel price data are only available beginning in 1990. In particular, the jet fuel price series begins in April 1990. Consequently, the analysis uses a sample period from April 1990 through December 2019, yielding 357 monthly observations. The sample ends in December 2019 to exclude the onset of the COVID 19 pandemic, which caused an unprecedented disruption to the airline industry and an abrupt collapse in passenger travel.

Figure 2.1: *Domestic Revenue Passenger Miles (RPMs)*




Figure 2.1 shows the effect of the COVID-19 pandemic on the airline industry by using the U.S. Domestic Revenue Passenger Miles ([RPMD](#)) from the FRED database, where the

periods of the U.S. recessions are shaded in light grey based on the National Bureau of Economic Research (NBER) definitions. RPM, a typical indicator for the traffic in the airline industry, is defined as the number of paying passengers multiplied by the distance traveled. The index indicates a sharp decline in April 2020, where the RPM index dropped to about 96% below its December 2019 value. It took almost one and a half years for the travel demand to recover, where the RPM index reached its pre-pandemic value by July 2021. In order to avoid the effects of the extraordinary COVID-19 pandemic, the data is cut-off at December 2019.

2.3.1 Identification of Impulse Response Functions

The reduced form residuals μ_{At} and μ_{Ft} are assumed to be driven by the structural shocks to both variables, ε_{At} and ε_{Ft} , and a measure of US aggregate demand, $\varepsilon_{AD,t}$. Accounting for US aggregate demand is important for the purposes of this study, because changes in the state of the economy will influence airfares and jet fuel prices in the same direction, potentially biasing the estimated response of airfares to jet fuel prices. We use a dummy variable that captures NBER recession dates for R_t . Our goal is to capture US-specific aggregate demand, not world aggregate demand.

$$\mu_{At} = \lambda_A \mu_{Ft} + \lambda_B R_t + \varepsilon_{At} \quad (3)$$

$$\mu_{Ft} = \lambda_C \varepsilon_{At} + \lambda_D R_t + \varepsilon_{Ft} \quad (4)$$

Our proxy SVAR framework requires an external instrument that is correlated with ε_{Ft} but not with ε_{At} . We use the precautionary oil demand shock of Kilian (2009). These shocks capture changes in jet fuel prices driven by unexpected shifts in demand for global crude oil, typically in response to news about future supply shortages.

Following [Ramey \(2016\)](#), We estimate equation (3) using the 2SLS method. In the first stage, the reduced-form jet fuel price residuals (μ_{Ft}) are regressed on the precautionary

demand shocks and aggregate demand shocks. The fitted values from this stage are then used as instruments in the second stage, where the reduced-form airfare residuals (μ_{At}) are regressed on the predicted jet fuel prices along with aggregate demand shocks. We use the residuals from equation (3) as a proxy for the reduced-form airfare residuals (μ_{At}) when estimating equation (4) with the OLS.

For the purpose of simulating the impulse responses of airfares to jet fuel price shocks, the main focus is on the coefficient λ_A , as it reflects the immediate effect of jet fuel prices on airfares. The Proxy SVAR estimation produces a value of $\lambda_A = 0.016$, indicating that a one-unit shock in jet fuel prices leads to an immediate increase of 0.016 in airfares. This can be properly interpreted as: a one percentage-point increase in the year-over-year change in jet fuel prices leads to an immediate 0.016 percentage points increase in the year-over-year change in airfares, within the same month.

2.3.2 Impulse Response Functions (IRFs)

After identifying the impact effect, the next step is to forecast the future path of airfares. To this end, we employ the local projection approach proposed by [Jordà \(2005\)](#). Specifically, we consider a 12 month horizon and estimate equation (5) using OLS separately for each horizon ($h = 1, \dots, 12$).

$$A_{T+h} = \beta_0 + \beta_{1,T}Travel_T + \sum_{i=0}^2 \beta_{2,i}A_{T-i} + \sum_{i=0}^2 \beta_{3,i}F_{T-i} + \sum_{i=0}^2 \beta_{4,i}(Travel_{T-i} \cdot F_{T-i}) \quad (5)$$

We include two lags of both A_T and F_T in the local projection regressions. This allows us to use a wild bootstrap procedure for inference, which is robust to heteroskedasticity and conditional heteroskedasticity but not to serial correlation; serial correlation is addressed by including lagged dependent and explanatory variables in the specification. $Travel_T$ denotes the monthly share of annual passenger traffic, calculated as the number of passengers in month m of year y relative to total passenger traffic in year y .

For each horizon, we construct forecasts of airfares h periods ahead under two scenarios: a jet fuel price shock at time T (denoted S_T) and a no shock counterfactual (denoted N_T), evaluated separately for each month (M). In the shock scenario, a one unit increase in the jet fuel price F_T increases airfares A_T by 0.016 units, while in the no shock scenario the change in F_T and hence A_T is set to zero. For both sets of forecasts, the lagged variables A_{T-1} , A_{T-2} , F_{T-1} , and F_{T-2} are set to zero.

To allow the response to vary across months, we substitute the mean monthly share of annual passenger travel ($Travel_T$) into the forecasting equation, which captures seasonal differences in travel demand. Using these definitions, the forecasts under both shock and no-shock scenarios across months can be expressed as follows:

$$\hat{A}_{T+h}^M | S_T = \hat{\beta}_0^M + (\hat{\beta}_{1,T}^M \times \delta_A) + (0.016 \times \hat{\beta}_{2,0}^M) + \hat{\beta}_{3,0}^M + (\hat{\beta}_{4,0}^M \times \delta_A) \quad (6)$$

$$\hat{A}_{T+h}^M | N_T = \hat{\beta}_0^M + (\hat{\beta}_{1,T}^M \times \delta_A)$$

The impulse response functions (IRFs) for month m are then computed for each horizon as:

$$IRF_{T+h}^M = \hat{A}_{T+h}^M | S_T - \hat{A}_{T+h}^M | N_T \quad (7)$$

We then compute the the IRFs over all forecast horizons.

2.3.3 Confidence Intervals for Impulse Response Functions

To obtain confidence intervals for the IRFs, we apply the wild bootstrap with 10,000 replications. For each bootstrap replication we generate a new dataset by combining the model-fitted values for airfares ($\hat{A}_t$) from equations (1) with their estimated residuals ($\hat{\mu}_{At}$), using the following:

$$\tilde{A}_t = \hat{A}_t + z \hat{\mu}_{At} \quad (8)$$


Here, z is a random variable that takes values from $\{-1, 1\}$. In each replication, the models specified in equation (5) is re-estimated using the simulated dataset generated for that replication. The forecasts and IRFs are then computed over the entire forecast horizon. The median IRFs are presented in the results, along with the 2.5th and 97.5th percentile values to represent the 95% confidence intervals.

2.4 Results

We first compare the pass-through of jet fuel cost shocks during the highest and lowest travel months. Based on the average monthly share of annual passenger traffic, July has the highest share, whereas February has the lowest. Figure 2.2 presents the impulse responses of airfares over a 12 month horizon to a jet fuel price shock occurring in July and February, based on the local projection model in Equation (5). The solid red line represents the median response for July, while the blue line corresponds to February. The shaded areas indicate 95 percent confidence intervals. The figure illustrates how the extent of cost pass-through varies with the month of travel.

The impulse response functions show that airfares respond immediately to a jet fuel price shock occurring in July, indicating that airlines begin passing higher fuel costs on to consumers in the same month. Airfares continue to increase over subsequent months, reaching their peak approximately three months after the shock at about 0.10, before returning toward their pre-shock level. Thus, a one percentage point increase in jet fuel prices in July raises airfares by about 0.10 percentage points at the peak response. In contrast, a jet fuel price shock occurring in February is not passed through immediately. Instead, airfares begin to increase only after about one month and continue increasing gradually over the forecast horizon, reaching a maximum response of about 0.11 around nine months after the shock, before declining slightly thereafter. Notice that the peak airfare response is of similar mag-

Figure 2.2: *Response of Airfares to Jet Fuel Shock*



nitude when the shock occurs in July or February; however, the timing of the adjustment differs substantially. In July, airfares reach their maximum response of about 0.10 within approximately three months after the shock, whereas in February a comparable peak is only reached after about nine months. This suggests that the main difference across travel months lies in the speed of pass-through rather than the eventual magnitude of adjustment.

While Figure 2.2 highlights differences between the highest and lowest travel months, it does not reveal the full seasonal pattern of cost pass through across the year. pass-through may vary meaningfully with travel intensity, and focusing only on two extreme months may mask important variation across intermediate months. To address this, Figure 2.3 presents the impulse responses of airfares to a jet fuel price shock occurring in each month of the year, traced over a 12 month horizon. The subplots are arranged by average monthly share of annual passenger traffic, ordered from highest to lowest travel intensity. The first row reports the two most travel intensive months, followed by the next two in the second row, and so on. Within each row, travel intensity decreases from left to right, with the most travel intensive month on the left.

Figure 2.3 reveals that airfares respond immediately and strongly to jet fuel price shocks occurring during the highest travel intensity months, namely July, August, and June. In these months, airlines begin passing higher fuel costs on to consumers in the same month, with the response reaching its peak approximately three months after the shock before declining rapidly and returning close to its pre-shock level by the end of the forecast horizon.


The next three months in terms of travel intensity, March, May, and October, exhibit a similar initial pattern. Fuel cost shocks are again passed through immediately and the peak response occurs around three months after the shock. However, unlike the highest travel intensity months, the subsequent decline is much more gradual, with the responses remaining elevated between approximately five and nine months after the shock. This persistence is particularly pronounced in October.

For the next group of months, April, December, and November, cost shocks are still passed through immediately, but the timing of the peak shifts toward approximately six months after the shock. Moreover, the responses remain relatively stable from about three to nine months after the shock before gradually declining, indicating a more prolonged adjustment process.

Finally, the three lowest travel intensity months, September, January, and February, display noticeably delayed pass-through dynamics. In September, airfares begin to respond earlier than in January and February. However, in all three cases, pass-through occurs gradually, with the peak response shifting further into the future and occurring around nine months after the shock.

Overall, Figure 2.3 reveals a clear seasonal gradient in cost pass through. As travel intensity declines, airlines adjust airfares more slowly following jet fuel price shocks, with the peak response occurring progressively later and the adjustment becoming more gradual and persistent. These findings suggest that seasonal differences in pass through are driven primarily by the timing and speed of price adjustment rather than by large differences in the eventual magnitude of the response.

Figure 2.3: *Response of Airfares to Jet Fuel Shock by Month*



The intuition behind the greater pass-through in high-demand periods is that in the short run, approximately within six months of a fuel cost shock, airlines face substantial operational rigidities. Flight schedules, aircraft assignments, crew rotations, and airport slot allocations are largely predetermined, which limits airlines’ ability to adjust capacity or reoptimize their networks ([Dobruszkes et al., 2019](#)). As a result, pricing becomes the primary margin of adjustment. During peak demand seasons, air travel demand is relatively inelastic due to congestion, limited substitution across travel dates, and a higher share of time-sensitive travel. This allows airlines to pass-through a larger fraction of fuel cost increases into airfares without losing many passengers. In contrast, demand during off-peak periods is more elastic, as travelers have greater flexibility in travel timing, destination choice, and mode of transportation. To avoid losing passengers, airlines absorb a larger share of cost increases, resulting in lower short run pass-through during low-demand seasons. In addition, airlines may rely on fuel hedging strategies that smooth fuel costs over horizons of up to a year ([Morrell and Swan, 2006](#)), which further limits immediate fare adjustments and contributes to lower observed pass-through in off-peak periods.

Over longer horizons, airlines gain access to additional margins of adjustment that fundamentally change pass-through behavior. In particular, airlines can reoptimize their networks by reducing flight frequencies, reallocating aircraft, downsizing capacity, or exiting marginal routes ([Aviation at Lazarski University, 2025](#); [Reynolds, 2008](#); [Sri, 2023](#)). These adjustments are disproportionately concentrated in off-peak periods, where load factors are lower and routes are more likely to become unprofitable following a fuel cost shock. Capacity reductions in off-peak periods increase load factors on the remaining flights and weaken competitive pressure as marginal competitors reduce service or exit the market, effectively increasing market power despite more elastic demand. As a result, airlines are able to raise fares more aggressively in off-peak periods over the long run. In contrast, persistent capacity reductions during peak seasons are more difficult, as flights are already highly utilized and competition for profitable peak routes remains strong. Moreover, in the long run, consumers adjust their travel behavior and firms anticipate these responses, which further constrains pricing in peak periods. Taken together, these mechanisms imply that while pass-through

is higher during peak seasons in the short run, long run pass-through becomes relatively larger in off-peak seasons as capacity reductions and network restructuring tighten supply and enhance pricing power.


2.4.1 Do Shocks in Different Months Affect the Same Future Month Differently?

So far, the analysis has examined seasonal variation in cost pass through using impulse responses that trace how airfares respond over the twelve months following a jet fuel price shock occurring in a given month. For example, we study how a shock in August affects airfares in August and in the subsequent months.

Another way to think about seasonality in cost pass-through is to examine how airlines adjust fares for peak or off-peak travel months when the fuel cost shock occurs prior to the travel month. This is important because travelers often purchase tickets in low travel intensity months for travel that takes place in high travel intensity months. For example, a passenger may purchase a ticket in February for travel in August. Therefore, an ideal setting would provide airfare series for different travel months observed at the time the ticket is purchased: fares for April travel observed in February, fares for May travel observed in February, fares for August travel observed in February, and so on. This would allow us to examine how the same February fuel price shock affects fares for different future travel periods. Specifically, we could compare the immediate response of April and August fares to the February shock and trace how each response evolves over the subsequent months. This would allow us to determine whether airlines adjust fares differently for peak versus off peak travel periods following the same fuel cost shock.

However, our data only allows us to observe the month when the ticket was purchased, not the actual month of travel. We do not know whether a ticket purchased in a given month is for travel in that same month, the following month, or several months later. As a result, we cannot directly link a fare adjustment to the actual travel season in which the passenger flies.

Figure 2.4: *Airfares Responses in a Given Month to Jet Fuel Price Shocks in Previous Months*



What we can do is treat the month in which airfare is observed as the travel month and examine how airfares respond to fuel price shocks that occurred in earlier months. For example, we can fix the airfare observed in August and examine how an earlier fuel price shock in February affects fares for the August travel period, which corresponds to a peak travel month. We can then compare this response with the response of fares observed in April to the same February shock, where April represents a lower travel intensity month. This approach provides suggestive evidence on whether airlines adjust fares differently across peak and off peak travel periods following fuel cost shocks that occurred in previous months. An important limitation of this approach is that the timing between the shock and the observed airfare response differs across travel months. Specifically, the response of April fares to a February fuel shock is measured two months after the shock, whereas the response of August fares to the same February shock is measured six months after the shock. Therefore, differences in the estimated responses may reflect differences in the time elapsed since the shock rather than differences in airlines' pricing behavior across peak and off peak travel periods.

Given this limitation, the best comparison possible in our setting is to examine whether the response of peak season fares to an off peak fuel price shock differs from the response of off peak season fares to a peak season shock while holding the time elapsed since the shock constant. Specifically, we compare how airfares observed in August respond to a jet fuel price shock occurring in February with how airfares observed in January respond to a shock occurring in July. In both cases, the response is measured six months after the shock, allowing us to reduce the influence of differences in pass through horizons. This comparison provides suggestive evidence on whether airlines adjust fares differently across peak and off peak travel periods when exposed to fuel cost shocks occurring at different points in the travel cycle.

Figure 2.4 presents the results for this approach. For each response month, it shows how airfares in that month react to jet fuel price shocks occurring in the same month and in earlier months. In each panel, we highlight two specific cases. The blue dots show the response in a given month to a jet fuel price shock in February, illustrating how a low season

shock carries over into different pricing periods. The red dots show the response in a given month to a jet fuel price shock in July, capturing how a high season shock is reflected across the year.

We find that the response of August airfares to a jet fuel price shock occurring in February is 0.10, while the response of January airfares to a jet fuel price shock occurring in July is 0.07. The similar magnitude of these responses suggests that, in our setting, fare adjustments are not substantially different when an off peak fuel price shock affects a peak travel period compared with when a peak season shock affects an off peak travel period.

A similar pattern emerges when comparing groups of months. High travel intensity months such as June, July, and August exhibit pass through responses of similar magnitude to low travel intensity months such as January, February, and September. Comparing shocks originating in low travel intensity months (e.g., February) with shocks originating in high travel intensity months (e.g., July) provides little evidence that airlines adjust fares differently across peak and off peak travel periods. Overall, these results suggest that the timing of the travel period does not substantially alter the pass through of fuel cost shocks.

2.4.2 Airfare Responses to Extreme Jet Fuel Price Shocks

Next, we investigate whether the pass-through of jet fuel cost shocks to airfares across seasons depends on the strength of the underlying oil price shock. As argued by Hamilton (1996), the macroeconomy responds to oil prices in an asymmetric and nonlinear way. His central idea is that consumers and firms mainly adjust their behavior when oil prices reach historically high and unfamiliar levels. In contrast, routine price changes that only reverse recent declines are less likely to have strong economic effects because they remain within normal expectations.

Motivated by this idea, we examine whether airlines respond differently to large or “extreme” oil price increases than to normal price changes. If airlines adjust fares differently following extreme oil price shocks, then the seasonal pattern of cost pass-through may also differ between periods of unusually large price increases and periods of more typical fuel price movements.


To test this hypothesis, we modify our baseline model of local projections in equation (5) by introducing a Net Oil Price Increase indicator and interacting it with jet fuel prices and monthly travel intensity. Hamilton defines a Net Oil Price Increase as the amount by which the current oil price exceeds its maximum level over the previous 12 quarters, assigning a value of zero whenever the current price does not exceed that historical peak. We convert this measure into a binary indicator that equals one when the Net Oil Price Increase is positive and zero otherwise. This specification allows us to examine whether unusually large oil price shocks are passed through to airfares differently than normal price shocks and whether any such differences vary across periods of high and low travel intensity. The modified specification is given as:

$$\begin{aligned}
A_{T+h} = & \beta_0 + \beta_{1,T}Travel_T + \gamma_{1,T}NOPI_T + \sum_{i=0}^2 \beta_{2,i}A_{T-i} + \sum_{i=0}^2 \beta_{3,i}F_{T-i} \\
& + \sum_{i=0}^2 \beta_{4,i}(Travel_{T-i} \cdot F_{T-i}) + \sum_{i=0}^2 \gamma_{2,i}(Travel_{T-i} \cdot F_{T-i} \cdot NOPI_{T-i}) \quad (9)
\end{aligned}$$

We first implement this extension by examining how jet fuel price shocks occurring in a given month affect airfares observed in the same month as well as in the 11 subsequent months. This framework allows us to assess whether unusually large fuel price shocks have heterogeneous effects on airfares depending on the timing relative of the shock.

Figure 2.5 compares the IRFs airfares to normal price shocks and extreme oil price shocks across months ranked by travel intensity. Overall, the results indicate that extreme oil price shocks do not lead to a stronger immediate pass-through into airfares. Instead, they tend to generate a more gradual adjustment path, particularly for shocks occurring during high travel intensity months. For example, following a July shock, the response under a normal price shock rises to 0.09 after two months and peaks at approximately 0.10 after three months. In contrast, under an extreme oil price shock, the corresponding responses are only 0.06 and 0.10, with the response remaining elevated for longer and exhibiting a secondary increase around the sixth month. Similar patterns are observed for August and June, where normal

Figure 2.5: *Response of Airfares to Jet Fuel Shocks by Month*



price shocks generate larger responses during the initial months following the shock, while extreme oil price shocks produce a slower buildup.

As travel intensity declines across the calendar, the differences between the two regimes become less pronounced. For lower travel intensity months such as January and February, both normal and extreme oil price shocks eventually generate responses of similar magnitude, although the adjustment under extreme shocks remains more gradual. For instance, the January response reaches approximately 0.11 under both normal price shock and extreme oil price shock around the seventh month after the shock, despite the latter exhibiting substantially smaller responses at shorter horizons. Taken together, these findings suggest that the primary effect of extreme oil price shocks is to alter the timing of cost pass-through rather than to increase its overall magnitude. Airlines appear to incorporate unusually large fuel cost increases into fares more gradually, especially when shocks occur during periods of high travel intensity.

We next examine whether cost pass-through differs between high and low travel intensity months under normal and extreme oil price shocks when the shock occurs before the month in which the airfare is observed. As in the previous subsection, we fix the response month and evaluate how shocks originating in earlier months affect airfares in that specific month under the two oil price regimes. We then plot the responses to normal and extreme oil price shocks on the same figure for comparison.

Figure 2.6 presents the results. Circles denote responses to normal oil price shocks, while triangles denote responses to extreme oil price shocks. To facilitate interpretation, shocks originating in February, the lowest travel intensity month, are highlighted in blue, allowing us to track how a shock during a low demand period affects airfares in each response month. Similarly, shocks originating in July, the highest travel intensity month, are highlighted in red, illustrating how shocks occurring during a peak travel period transmit to airfares across different response months.

Overall, we observe that, for nearly all response months, the triangles lie below the corresponding circles for most shock months. This indicates that airfare responses to extreme oil price shocks are generally smaller than responses to normal oil price shocks. The finding is


consistent with the evidence presented in Figure 2.5, which suggests that airlines do not immediately pass-through unusually large increases in fuel costs. Instead, they appear to adjust fares more gradually following extreme oil price shocks, resulting in lower contemporaneous pass-through relative to normal shocks.

To examine whether extreme oil price shocks occurring during low travel intensity months affect airfare responses in high travel intensity months differently from normal oil price shocks, we compare the response of airfares in August, a high travel intensity month, to shocks originating in February, the lowest travel intensity month. Notice that airfare responses in August are largest for shocks occurring in February relative to shocks originating in any other month. However, there is virtually no difference between the responses to extreme and normal oil price shocks, as illustrated by the overlap of the blue triangle and the blue circle in Figure 2.6. This suggests that when a shock originates during the lowest travel intensity month and affects pricing in a peak travel month, the magnitude of pass-through is largely independent of whether the underlying oil price shock is extreme or normal.

To analyze whether this relationship differs when airfares are observed in a low travel intensity month and the shock originates in a high travel intensity month, we consider the response of airfares in January to a shock occurring in July. As in the August February comparison, there is little evidence of differential pass-through between extreme and normal oil price shocks. The red triangle and red circle in Figure 2.6 almost completely overlap, indicating that the magnitude of the response is essentially the same under the two oil price regimes.

These findings support the broader conclusion of this subsection that pass-through is generally stronger for normal oil price shocks than for extreme oil price shocks. Even during periods of high travel intensity, airlines appear reluctant to fully and immediately pass-through extreme increases in fuel costs, instead delaying or smoothing these adjustments relative to the response observed following normal oil price shocks.

Figure 2.6: *Airfares Responses in a Given Month to Jet Fuel Price Shocks in Previous Months*



2.4.3 The Iran War Shock: Counterfactual Seasonal Pass Through

To illustrate the importance of accounting for seasonal variation in the pass-through of jet fuel costs to airfares, we do a counterfactual analysis of the recent conflict involving Iran. Following the onset of the conflict, we assume that the observed jet fuel price changes in March and April were driven by the Iran conflict and therefore treat them as jet fuel price shocks. The corresponding increases in jet fuel prices were 1.43 percent in March and 0.23 percent in April. Based on the estimated impact coefficient (0.016), these shocks increase airfares by 0.024 units (1.43×0.016) in March and by 0.004 units (0.23×0.016) in April on impact.

Using these impact responses, we compute counterfactual responses conditional on the conflict starting each of the 12 months of the year. In other words, we compute the impulse response for each month of the calendar year. We separately estimate the response to the first and second shocks. To obtain the joint effect of two consecutive shocks, we combine the impulse responses by adding the response at horizon (h) to the first shock with the response at horizon ($h-1$) to the second shock, for ($h=1, \dots, 11$). At horizon 0, only the contemporaneous response to the first shock is observed because the second shock has not yet occurred. This procedure yields the joint response of airfares to consecutive jet fuel price shocks for each two month combination.

To facilitate comparison across seasons, we rank the two month combinations according to their average passenger traffic. Specifically, we first compute the average passenger share for each calendar month across all years in the sample. We then calculate two month moving averages of these monthly shares and use these values to rank the month combinations from the highest to the lowest passenger demand.

Table 2.1 reports the joint response of airfares following two consecutive jet fuel price shocks occurring in consecutive months. We begin by examining the March-April combination because it corresponds to the sequence of shocks observed following the conflict. Under this observed sequence, the model predicts an increase in airfares of approximately 0.024 in the month of the initial shock. The response then rises steadily over the follow-

ing months, reaching about 0.146 after three months and peaking at approximately 0.158 after four months. Although the response gradually declines thereafter, it remains elevated throughout the one year horizon, indicating that the effect of consecutive fuel price increases persists well beyond the initial shock.

Table 2.1: *Impulse Response of Airfares to Jet Fuel Shocks*

Months of Shock	Months after shock											
	0	1	2	3	4	5	6	7	8	9	10	11
Jul-Aug	0.024	0.048	0.106	0.166	0.170	0.137	0.136	0.125	0.104	0.108	0.101	0.064
Jun-Jul	0.024	0.044	0.104	0.162	0.174	0.140	0.139	0.132	0.109	0.113	0.103	0.069
May-Jun	0.024	0.041	0.096	0.151	0.169	0.145	0.147	0.145	0.127	0.131	0.116	0.082
Mar-Apr	0.024	0.042	0.090	0.146	0.158	0.147	0.152	0.150	0.140	0.144	0.130	0.088
Apr-May	0.024	0.038	0.089	0.143	0.164	0.149	0.154	0.156	0.142	0.146	0.128	0.092
Aug-Sep	0.024	0.047	0.090	0.147	0.147	0.144	0.152	0.145	0.142	0.148	0.137	0.087
Oct-Nov	0.024	0.039	0.085	0.138	0.154	0.150	0.158	0.159	0.152	0.157	0.139	0.096
Nov-Dec	0.024	0.035	0.081	0.133	0.158	0.154	0.161	0.168	0.159	0.163	0.141	0.103
Sep-Oct	0.024	0.031	0.081	0.131	0.164	0.156	0.162	0.172	0.160	0.163	0.139	0.105
Feb-Mar	0.024	0.025	0.076	0.123	0.170	0.161	0.167	0.184	0.170	0.172	0.142	0.114
Dec-Jan	0.024	0.036	0.073	0.125	0.142	0.155	0.169	0.175	0.178	0.182	0.161	0.112
Jan-Feb	0.024	0.027	0.061	0.108	0.143	0.165	0.181	0.199	0.205	0.207	0.177	0.132

Based on our ranking, the busiest two month period is July-August, followed closely by June-July and May-June, and these combinations appear at the top of Table 2.1. We next consider a counterfactual scenario in which the same sequence of jet fuel price shocks occurred during these peak travel periods instead of March and April. Specifically, we ask how airfares would have evolved if the conflict had begun in July, generating consecutive shocks in July and August, if it had begun in June, generating shocks in June and July, or if it had started in May, resulting in shocks in May and June. The results indicate substantially faster pass-through during these high demand periods. Following a July-August shock sequence, the joint response increases from 0.024 on impact to 0.106 after two months and reaches approximately 0.170 after four months. The June-July combination exhibits a very similar pattern, with the response reaching approximately 0.174 after four months. Likewise, the May-June combination follows nearly the same trajectory, with the response increasing to approximately 0.169 after four months. Compared with the March-April scenario, these counterfactuals show that airfares rise more rapidly when identical jet fuel price shocks occur during the busiest travel months, suggesting that airlines are able to pass higher fuel

costs through to consumers more quickly when demand is strongest.

At the other end of the ranking are the January-February, December-January, and February-March combinations, which correspond to the periods with the lowest passenger demand. We next consider the counterfactual in which the same sequence of jet fuel price shocks occurred during these low travel periods. In contrast to the peak season results, pass-through is considerably slower. Following a January-February shock sequence, the joint response increases only marginally from 0.024 on impact to 0.027 after one month and reaches 0.061 after two months. The response continues to build gradually over time, eventually reaching approximately 0.207 after nine months, the largest cumulative response among all month combinations. The December-January and February-March combinations exhibit similar delayed adjustment patterns. Under the December-January scenario, the response reaches only 0.125 after three months before continuing to increase to approximately 0.182 after nine months. Likewise, the February-March combination rises to only 0.123 after three months but continues to increase, reaching approximately 0.172 after nine months. These findings indicate that when consecutive fuel price shocks occur during periods of weak travel demand, airlines initially absorb a larger share of the increase in fuel costs and delay passing them through to consumers. As demand gradually strengthens in subsequent months, airlines progressively incorporate these higher costs into ticket prices, resulting in a much slower but more persistent adjustment process than that observed during the peak travel season.

An important feature of the results is that seasonality primarily affects the timing rather than the eventual magnitude of pass-through. Consecutive fuel price shocks occurring during peak travel months generate a rapid increase in airfares, whereas similar shocks during low demand periods lead to substantially delayed adjustments. Thus, the economic consequences of consecutive oil price increases depend not only on their size but also on the time of year in which they occur.

2.5 Conclusion & Implications

This essay examines whether fuel cost pass-through in the U.S. airline industry varies systematically across demand seasons. Using aggregate industry-level data, we document that airfare responses to jet fuel price shocks vary depending on the month in which they occur. The extent of pass-through depends on whether the shock occurs during peak or off-peak travel periods.

Our results show that seasonal conditions primarily affect the timing of jet fuel cost pass-through to airfares rather than its long run magnitude. Pass-through is more immediate during peak travel periods, with fares in July responding quickly and reaching their maximum within a few months, whereas responses in February build up more gradually despite converging to a similar long run magnitude. Using an alternative timing framework that accounts for advance ticket purchases, we also find that fuel cost shocks originating in low travel intensity months are transmitted more strongly when fares are set during peak travel periods than vice versa. We further show that these seasonal patterns depend on the size of the underlying fuel cost shock. Normal oil price shocks are passed through more rapidly, while extreme oil price shocks lead to delayed fare adjustments, suggesting that airlines smooth the transmission of unusually large cost increases rather than immediately passing them on to consumers. Finally, using the recent Iran conflict as a motivating example, counterfactual analysis shows that the observed jet fuel price shocks would have generated substantially faster airfare adjustments when they occur during peak travel periods, whereas the same shocks during low demand periods produce a slower but more persistent pass-through process.

The intuition behind our findings is that airlines face different pricing and operational constraints across seasons and over time. In the short run, limited flexibility to adjust capacity makes fare changes the primary mechanism for responding to fuel cost shocks, and relatively inelastic demand during peak travel periods allows airlines to pass through a larger share of these costs to consumers. As the adjustment horizon lengthens, airlines gain greater scope to modify flight frequencies and network capacity, particularly during

off-peak periods where such changes are more feasible. These longer run supply adjustments strengthen pricing power and increase pass-through in low demand seasons, implying that seasonal differences arise mainly through the timing of adjustment rather than the eventual magnitude of fuel cost pass-through.

These findings have important implications for the broader literature on asymmetric pass-through, as estimates that ignore seasonal heterogeneity may attribute differences in pricing behavior to sign asymmetries or other nonlinearities when they instead reflect differences in demand conditions across the year.

Chapter 3

Low Skilled Immigration and the Hidden Link to the Gender Pay Gap

3.1 Introduction

Women are stepping up in their role in the modern economy through receiving more education and actively participating in the labor force. Over the past five decades female labor supply has grown across the globe and more so in the western economies.¹ Although the increased role of women in the labor market has improved their living standards, it has come with a tradeoff between fertility and employment opportunities. This tradeoff can be particularly high for mothers who are highly educated and work in high paying jobs. The provision of low-cost childcare, whether private or public, can ease this tradeoff.

Research has shown that an increase in low skilled immigrant inflows increases the supply of labor in the market for domestic services, suppressing the wages and hence reducing the cost of childcare (Furtado, 2016). Metropolitan regions which receive proportionately higher immigrant inflows may witness larger decrease in the wages of childcare workers relative to the regions which receive lower immigrant inflows. An increase in the population of low skilled immigration can elicit varied responses from native working women: increased

¹[Our World in Data](#)

fertility because of the low cost of childcare or, alternatively increased labor hours ([Blau and Robins](#); [Cortés and Tessada, 2011](#); [Furtado, 2016](#)) leading to decreases in fertility. Given this evidence, is it plausible to assume that women may also respond to increase in low skilled immigration by taking up more time inflexible roles within an occupation? Or switch towards more time inflexible occupations?

Working in time inflexible occupations is important for women. ([Goldin, 2014](#)) demonstrates that in time inflexible occupations the relationship between earnings and work hours is nonlinear in a convex fashion. This implies more continuous work hours and work hours provided at particular moments are worth more for the firm and hence remunerated proportionately. However, women desire more time flexibility due to arrival of children ([Reynolds and Johnson, 2012](#)). Fewer working hours and more flexible schedule means significant wage penalties in a nonlinear fashion ([Mas and Pallais, 2017](#); [Weeden et al., 2016](#)). Thus, low skilled immigrants, by providing low-cost childcare and household services, may encourage high skilled women to provide more and continuous work hours to avoid wage penalties. Moreover, high skilled women may develop preferences for time inflexible occupations if their childrearing duties are eased due to provision of low-cost childcare. Native women who rely more on family support for childcare might not respond to the lowering of childcare costs. However, it is expected that many of the highly educated women working in high paying jobs might live away from families and hence rely more heavily on market based childcare services ([Molloy et al., 2011](#)).

The literature has extensively studied how low-skilled immigration affects native women’s fertility decisions ([Bauernschuster et al., 2016](#); [Furtado, 2016](#)) and labor market outcomes ([Cortés and Tessada, 2011](#); [Forlani et al., 2020](#)). However, to our knowledge, no study has examined whether low-skilled immigration affects native women’s preference for time-inflexible jobs. We address this gap by constructing an occupation-level index of time inflexibility and exploiting variation across Public Use Microdata Areas (PUMAs) and over time in the share of predicted working-age population composed of recent low-skilled immigrants to explain changes in this index.

We construct the job inflexibility index using occupational characteristics from the Work

Activities and Work Context modules of O*NET. Following [Goldin \(2014\)](#), the index is based on five dimensions: time pressure, contact with others, structured versus unstructured work, freedom to make decisions, and establishing and maintaining interpersonal relationships. To construct the immigration and other relevant variables for our analyses, we use two five-year American Community Survey (ACS) waves from IPUMS: 2005–2009 (2009 ACS) and 2015–2019 (2019 ACS).

Our main variable of interest, low skilled immigration, is measured using two related definitions. The baseline measure is the share of a PUMA’s predicted working age population (ages 20–64) composed of recent low skilled immigrants, while the female measure is the corresponding share composed of recent low skilled female immigrants. We consider the female specific measure because domestic and childcare work is performed predominantly by women, making female immigrants more likely to substitute for unpaid household labor and influence the market for household services ([Cortés, 2023](#)).

The immigration literature has well established that immigrants’ location choices are not exogenous. Endogeneity may arise, for example, if immigrants select into areas experiencing improving economic conditions. In such cases, rising economic conditions may increase childcare worker wages and simultaneously attract additional immigrants, leading to a spurious correlation between immigration and outcomes of interest. To address this concern, we use a shift-share instrumental variable approach following [Tabellini \(2020\)](#), which predicts immigrant inflows to each PUMA based on historical settlement patterns of individuals from the same country of birth. We construct these historical settlement shares using Census 2000 data on the distribution of immigrants across PUMAs by country of birth.

We first examine whether inflows of recent low skilled immigrants reduce the cost of childcare services at the PUMA level by expanding childcare labor supply and thereby putting downward pressure on wages. For this purpose, we replicate the empirical strategy in [Furtado \(2016\)](#) and extend her framework by incorporating a female specific measure of low skilled immigration in addition to the baseline measure. We estimate four 2SLS specifications using two PUMA level outcomes, childcare labor supply and median childcare wages, combined with the two alternative measures of low skilled immigration and their corresponding instru-

ments.

Although the estimated effects are not statistically significant, the results consistently point in the expected direction. In the labor supply specifications, the coefficients on both immigration measures are positive, providing suggestive evidence that increases in low skilled immigration expand the local supply of childcare workers. In the median wage specifications, the coefficients are consistently negative, suggesting that this increase in labor supply places downward pressure on childcare wages and, consequently, reduces the cost of childcare services. Moreover, the female specific measure generally produces larger effects than the baseline measure, indicating that female low skilled immigrants play a disproportionately important role in shaping local childcare market outcomes. Taken together, these findings provide suggestive evidence that inflows of low skilled immigrants improve the availability and affordability of childcare at the local level.

The absence of statistically significant effects is consistent with the discussion in [Furtado \(2016\)](#), who argues that estimates based on ACS data are considerably noisier than those based on the decennial Census due to differences in data collection and the measurement of recent immigrant exposure. [Grieco et al. \(2018\)](#) further show that the ACS methodology can substantially understate the number of recent immigrants, which may attenuate estimated effects by introducing additional measurement error in the construction of immigrant shares.

We next examine whether reductions in childcare costs have encouraged high skilled native women to select into more time inflexible occupations. For this analysis, we construct a sample of high skilled native born women of working age (20–64), where high skilled individuals are defined as those with at least a bachelor’s degree. This restriction reduces the likelihood of direct labor market competition between low skilled immigrants and native women.

We estimate two 2SLS specifications using the baseline and female specific measures of recent low skilled immigration and their corresponding instruments. The estimated effects are small, negative and statistically insignificant. This suggests that, once we isolate plausibly exogenous variation in low skilled immigration, immigrant inflows are not associated with greater occupational inflexibility among high skilled native women on average.

Our main sample includes high-skilled women with heterogeneous family structures and childcare needs. This includes women without children, as well as mothers with children of different ages, whose demand for childcare services varies substantially. In particular, women with school-age children (ages 5–17) are likely to have lower childcare needs, while those with at least one child under age five face the highest demand for childcare and domestic services. These differences imply that the effect of low-skilled immigration on occupational choice may vary across groups. Accordingly, we consider three subsamples of high-skilled native women: women without children, women with school-age children, and women with at least one child under age five.

For women without children, the estimated effects are negative and statistically insignificant, consistent with limited reliance on childcare services and therefore weak exposure to changes in childcare costs. Among mothers, both those with school age children and those with children under the age of five, the estimates are positive but remain statistically insignificant. The effect for mothers with children under five is approximately 2.8 times larger than that for mothers with school age children, consistent with the idea that childcare needs are more intensive for this group. In addition, estimates based on female low skilled immigration are larger than those based on the baseline measure, supporting the view that female low skilled immigrants play a key role in substituting for unpaid household work. Overall, the lack of statistical significance implies that these results should be interpreted as suggestive rather than conclusive evidence.

The remaining essay is structured in the following way. Section 3.2 reviews the literature on the effects of immigrant inflows on native women’s labor market outcomes and fertility decisions. Section 3.3 describes the data sources and the construction of key variables. Section 3.4 presents the empirical strategy and results for the childcare market analysis and provides evidence on the proposed mechanism through which immigration affects occupational choices. Section 3.5 presents the empirical strategy and results for the main analysis. Section 3.6 concludes.

3.2 Literature Review

There is a large body of research examining how low-skilled immigration affects local labor markets in host countries. The arrival of low-skilled immigrants can influence native workers through both direct and indirect channels. Directly, immigration may affect native wages and employment opportunities by changing the supply of labor in local markets. Indirectly, low-skilled immigrants often provide relatively affordable household services, such as childcare, housekeeping, and elderly care. Access to these services can reduce the time native women spend on household responsibilities, allowing them to increase their labor force participation and work longer hours.

If low-skilled immigration reduces the cost of low-skilled labor, it can also lower the prices of goods and services that rely heavily on such labor, thereby increasing consumer welfare. Using confidential Consumer Price Index (CPI) microdata and variation in the share of low-skilled immigrants across U.S. cities between 1980 and 2000, [Cortés \(2008\)](#) finds that a 10 percent increase in the share of low-skilled immigrants in a city's labor force reduces the prices of immigrant-intensive services, such as housekeeping, babysitting, gardening, and dry cleaning, by approximately 2 percent. The study further estimates that immigration between 1980 and 2000 lowers the prices of these services by around 9–11 percent on average across cities. The main mechanism behind these price effects is a decline in wages for low-skilled workers, particularly low-skilled immigrants and natives who are more likely to compete directly with immigrants, such as Hispanic workers and individuals with limited English proficiency. Furthermore, using data from the Consumer Expenditure Survey, the study also shows that the gains from lower prices are larger for highly skilled natives, especially those with graduate degrees, because they spend a greater share of their income on immigrant-intensive household services. Thus, the overall effect of immigration on natives' purchasing power depends not only on wages, but also on changes in prices.

Because low-skilled immigration reduces the prices of services that substitute for household production, natives may replace time spent on household tasks with cheaper market-provided services. This suggests that low-skilled immigration can influence how natives

allocate their time between market work and household responsibilities. Building on this idea, [Cortés and Tessada \(2011\)](#) develop a simple model to identify which groups of women are most likely to adjust their time-use decisions in response to lower household service prices. Using U.S. Census data on labor supply, household work, and expenditures on housekeeping services, and an instrumental variable approach that exploits cross-city variation in low-skilled immigrant concentrations, they find that the effects of low-skilled immigration are concentrated among highly paid women. Specifically, low-skilled immigration significantly increases weekly hours worked for women in the top quartile of the female wage distribution, with smaller but still significant effects for women above the median wage, while no effects are found for women below the median wage. The study further shows that highly skilled women respond mainly through the intensive margin of labor supply rather than labor force participation. Since high-wage occupations often require long working hours and low-skilled immigrants provide relatively cheap and flexible household services, the authors argue that immigrant-provided services help women in demanding occupations work longer hours. Supporting this mechanism, they find that low-skilled immigration significantly increases the likelihood that women employed in occupations characterized by long working hours work more than 50 or 60 hours per week.

[Farré et al. \(2011\)](#) examine how immigration affects the size of the household services sector and the labor supply of natives in Spain. They pay particular attention to highly skilled native women, highly skilled men, and low-skilled natives, as well as whether women with greater family responsibilities respond more strongly to immigration. They primarily focus on the impact of female immigration, which in the case of Spain is more closely related to the household services sector. Their results are consistent with the broader literature, showing that immigration increases the size of the household services sector and leads to a significant rise in hours worked among women in high-earning occupations, by roughly two hours per week, while no comparable effect is found for similarly skilled men. They also find that the effects are stronger for women with greater family responsibilities, who are able to return to work more quickly after childbirth, remain in employment while caring for elderly relatives, and delay retirement when their spouses retire. Similar evidence from

other European contexts is provided by [Forlani et al. \(2020\)](#) and [Barone and Mocetti \(2011\)](#). [Forlani et al. \(2020\)](#) find that an increase in the share of female immigrants at the local level in Germany increases native women's working hours and raises the probability of having a child. [Barone and Mocetti \(2011\)](#) show that a higher supply of immigrant-provided household services increases native Italian women's working hours (intensive margin) without affecting labor force participation (extensive margin), with effects concentrated among highly skilled women whose time has a higher opportunity cost.

([Cortés and Pan, 2013](#)) examine how a policy change in Hong Kong that allowed the systematic importation of full-time live-in domestic workers (FDWs) affected labor supply and welfare of native women through increased access to affordable domestic services. They use two complementary empirical approaches. First, they compare employment rates of mothers with young children (aged 0–5) and those with older children (aged 6–17) in Hong Kong and Taiwan from 1978 to 2006, finding that mothers with a higher opportunity cost of time are more likely to purchase domestic services and increase their labor supply. In particular, the FDW program in Hong Kong increases female employment rates by 8–12 percentage points for mothers with young children relative to those with older children. Second, they develop a simple model of women's joint decisions to participate in the labor force and hire FDWs, and using pooled cross-sectional data from the 2001 and 2006 Hong Kong population censuses, they find strong complementarity between female labor force participation and the use of FDWs, such that reductions in the relative wage of FDWs significantly increase the probability of women entering the labor force. This complementarity is especially strong among mothers of very young children, suggesting substantial substitution between mothers' time and domestic worker time in childcare. Welfare analysis further shows that the gains from FDW availability are concentrated among mothers of very young children, highly educated women, and women with high unearned income.

In contrast to the literature focusing on immigration inflows, [East and Velasquez \(2019\)](#) exploit immigration enforcement as a source of exogenous variation that reduces the availability of immigrant labor, effectively generating a negative supply shock to immigrant-intensive household services. They study how enforcement policies affect the labor supply of high-

skilled native women by changing the cost of outsourcing household production. Because undocumented immigrants are disproportionately employed in household services, such as childcare, cooking, housekeeping, and gardening, stricter enforcement is expected to increase the price of these services. Using variation in the timing and location of enforcement policies, combined with American Community Survey (ACS) data from 2005–2012 and a staggered difference-in-differences design, they find that higher enforcement intensity reduces the labor supply of college-educated women with children. They provide supporting evidence that these effects operate through changes in the cost of household outsourcing, including increased time spent on household production among mothers in the American Time Use Survey, rising wages in household service occupations, and the absence of similar effects for high-skilled men or women without children.

These studies provide clear evidence that the arrival of low-skilled immigrants reduces the cost of household services in host countries. The main beneficiaries of these cheaper services are high-skilled native women, particularly those with young children. The findings suggest that low-skilled immigrants ease the time constraints associated with household and family responsibilities, thereby increasing high-skilled native women’s labor force participation and hours worked.

An important remaining question that has not been addressed in this literature is whether high-skilled native women also adjust their occupational choices by shifting toward more time-intensive or less flexible occupations. This question is particularly relevant to the broader gender pay gap literature. [Goldin \(2014\)](#) argues that much of the remaining gender wage gap in developed economies is not driven by differences in human capital, but by workplace structures that disproportionately reward long and inflexible work schedules while penalizing workers who demand flexibility. Although both men and women prefer to work fewer hours during the transition to parenthood ([Reynolds and Johnson, 2012](#)), high-skilled native women are more likely to demand workplace flexibility, especially after having children. However, demanding shorter or more flexible work schedules often carries substantial wage penalties ([Mas and Pallais, 2017](#); [Weeden et al., 2016](#)). In this context, it is important to examine whether reductions in the cost of household and childcare services induced

by low-skilled immigration relax high-skilled native women’s home production constraints and thereby influence their occupational choices, including potential shifts toward less flexible, higher-paying occupations. This essay addresses this gap in the literature and, to our knowledge, is the first to examine this relationship.

3.3 Data Sources and Variable Construction

For this study, we use two waves of the five year American Community Survey (ACS) data from IPUMS: ACS 2005 to 2009 and ACS 2015 to 2019. For simplicity, we refer to these datasets as the 2009 ACS and the 2019 ACS, respectively. We also use data from the 2000 Census to construct the instrumental variable. We exclude Alaska and Hawaii from the analysis and focus on the 48 contiguous U.S. states. We also drop individuals whose place of birth is in U.S. outlying territories, as well as immigrants who do not report a well defined country of birth.

Our study proceeds in two steps. We first examine whether the influx of low skilled immigrants reduces the cost of childcare services by putting downward pressure on the wages of childcare workers. Lower childcare costs increase the opportunity cost of not working for high skilled native women and may therefore allow them to choose occupations that are more time inflexible. We therefore examine whether low skilled immigration affects the occupational inflexibility of high skilled native women in the next step. This second step forms the main part of our analysis, while the first step examines a possible mechanism through which immigration may influence occupational choices and occupational switching. In the following discussion, we describe the samples used in each part of the analysis.

For our first step, we examine labor supply and wages in the childcare market. We use the full sample of working-age individuals to construct childcare labor market measures from individual-level ACS data, which are then collapsed to the PUMA level. Accordingly, the estimation sample consists of PUMAs, which serve as the unit of analysis.

We measure childcare labor supply as total childcare employment (including both native and immigrant workers) divided by the total labor force in each PUMA. To construct child-

care wages, we first calculate hourly wages using reported annual wage and salary income and annual hours worked. Annual earnings are converted to real 2019 dollars using the Consumer Price Index (CPI) from the [Bureau of Labor Statistics \(BLS\)](#). We then compute the median hourly wage at the PUMA level among native and pre-existing immigrant childcare workers.

Pre-existing immigrants are defined as immigrants who arrived approximately ten years before the census year in which they are observed. Accordingly, in the 2009 ACS we retain immigrants who arrived before 2000, while in the 2019 ACS we retain those who arrived before 2010. This restriction helps avoid a mechanical downward bias in median wages that may arise if recently arrived immigrants are disproportionately concentrated at the lower end of the wage distribution. We only consider those native and pre-existing immigrant childcare workers who report positive annual wage and salary income and positive usual weekly hours of work.

Our measure of low skilled immigration, which is the main variable of interest, is defined as the fraction of a PUMA’s predicted working age population (ages 20 - 64) made up of recent immigrants. Recent low skilled immigrants are defined as those who arrived after the year 2000 in the 2009 ACS and those who arrived after 2009 in the 2019 ACS. We define low skilled individuals as those with at most a high school degree and no college education.

Given the gender division of household work, female domestic workers are more likely to substitute for unpaid female labor within the household (Cortés, 2023). As a result, female low skilled immigration may have a stronger effect on the childcare market and, in turn, on the occupational choices of high skilled native women. Motivated by this, in addition to the baseline measure, we construct a female specific measure defined as the fraction of a PUMA’s predicted working age population (ages 20–64) made up of recent female low skilled immigrants.

Because a PUMA’s working age population is itself affected by immigration, we use a predicted measure of the working age population instead of the actual working age population. This is constructed by taking the PUMA’s working age population in the 2000 Census and scaling it using average growth in the total U.S. population, excluding growth in the

Census region where the PUMA is located (see the section 4.1 for details).²

For our main analysis, we restrict the sample to high skilled, native born, non Hispanic women of working age (20 - 64 years).³ High skilled individuals are defined as those who have completed at least a bachelor’s degree. This distinction between low skilled immigrants and high skilled native women reduces the likelihood of direct job competition that could affect female employment outcomes. We exclude high skilled women who are currently enrolled in college as well as those living in group quarters other than households.

We also exclude native Hispanic women. A large share of low skilled immigrants in the United States come from Latin America and tend to have higher fertility rates (Furtado, 2016). These immigrants may influence the fertility behavior of native Hispanic women through cultural transmission and peer effects. To avoid these potential spillovers, we focus only on non Hispanic native women.

To measure time inflexibility at the occupational level for high skilled native women, we construct a job inflexibility index using occupational characteristics from the Work Activities and Work Context modules of O*Net. Following Goldin (2014), we focus on five characteristics that capture time inflexibility in an occupation: time pressure, contact with others, structured versus unstructured work, freedom to make decisions, and establishing and maintaining interpersonal relationships. The first four measures come from the Work Context module, while the last is drawn from the Work Activities module. We provide details on each of these five characteristics below.

²Since our analysis requires the construction of variables at the PUMA level, a key challenge in earlier ACS releases is that PUMA boundaries and codes change after each census. As a result, PUMAs are not directly comparable over time. Recently, IPUMS introduced Consistent Public Use Microdata Areas (ConsPUMAs). These are defined as minimally aggregated groups of PUMAs that align well across different samples when combined. The goal is to match PUMAs across years using both consistent geographic boundaries and comparable population sizes. We use the CPUMA0010 version, which provides the smallest geographic units that can be consistently identified from the available geographic codes in the United States Census data from 2000 to 2021. More information is available on IPUMS’s website .

³As a robustness check, we re-estimate the analysis after including native Hispanic women in the sample. Although the magnitude of the estimated coefficients changes across subsamples (all women, women with school-aged children, and women with children under age five) and estimation methods (OLS and IV), the statistical significance of the results remains unchanged.

- Time pressure
 - *How often does this job require the worker to meet strict deadlines?*
 - Lower time pressure indicates that the worker does not need to be present at specific times to complete tasks.
- Contact with others - *How much does this job require the worker to be in contact with others (face-to-face, by telephone, or otherwise) in order to perform it?* - Less contact means greater flexibility.
- Freedom to make decisions
 - *How much decision-making freedom, without supervision, does the job offer.*
 - Generally, this means that the worker determines what each client should receive, rather than being given a specific project, and thus workers are poorer substitutes for each other the greater are these freedoms.
- Structured versus unstructured work
 - *To what extent is this job structured for the worker, rather than allowing the worker to determine tasks, priorities, and goals?*
 - If the job is highly structured to the worker, there would be a lower chance that the worker would have close substitutes.
- Establishing and maintaining interpersonal relationships
 - *Developing constructive and cooperative working relationships with others and maintaining them over time.*
 - The more working relationships, the more workers and clients the employee must be around.

The scale of measurement for each characteristic differs.⁴ Time pressure, contact with others, and freedom to make decisions are measured on a scale from 1 - 5, where 5 indicates the highest level. Structured versus unstructured work is measured on a scale from 1 - 3, where 3 indicates the highest level. Establishing and maintaining interpersonal relationships is recorded on two separate scales, with *importance* measured from 1 - 5 and *level* measured from 0 to 7. We use the *importance* scale for this characteristic.

We convert all characteristics to a common scale ranging from 0 to 1 using min - max normalization. The job inflexibility index is then constructed as the average of these normalized characteristics at the occupational level. A higher value of the index indicates a more inflexible occupation, while a lower value indicates a more flexible occupation.

To ensure consistency of occupations across the two rounds of ACS and O*NET, we use the six digit Standard Occupational Classification (SOC) crosswalks.⁵ Using these crosswalks, we harmonize occupational codes across ACS waves. In particular, we aggregate the 531 six digit occupations in the 2019 ACS to match the 471 six digit occupations in the 2009 ACS. After harmonization and aggregation, the pooled 2009 - 2019 ACS contains 432 occupations.

We apply the same occupational coding framework to the O*NET data. The O*NET dataset contains 894 occupations at the eight digit level, which we convert to six digit SOC codes using the crosswalks based on the 2009 ACS classification. We then aggregate O*NET occupations to the six digit level by taking weighted averages of characteristics, where the weights are the sample sizes reported for each occupation-characteristic cell at the eight digit level.⁶ This procedure yields 417 occupations at the six digit SOC level. We then merge the

⁴The details of these scales are available on the [O*NET website](#).

⁵Crosswalks are available on the [IPUMS website](#).

⁶For each characteristic and occupation, the data source is recorded as one of three categories: Occupational Expert, Incumbent, or Analyst Transition. For the first two categories, the information is based on surveys of individuals working in the occupation and the number of individuals surveyed is reported for each occupation - characteristic cell. For the Analyst Transition category, the information is provided by an analyst rather than through survey responses. We prioritize survey based data from incumbents and experts whenever it is available. We use analyst based data only when survey information is not available for an entire Standard Occupational Classification group. When the data for a given occupation characteristic is based only on an analyst source, we assign a weight of one in the weighted average.

O*NET data with the ACS data at the occupational level.⁷

3.3.1 Descriptive Statistics

We provide detailed summary statistics for the occupational inflexibility index in Table 3.1, which is constructed at the six-digit SOC level and aggregated to the two-digit occupation level for presentation. The table ranks occupations by mean inflexibility and reports the population-weighted distribution of high-skilled native women across these categories. This dual ordering allows us to jointly characterize (i) how restrictive different occupations are in terms of autonomy, scheduling, and interaction with others, and (ii) how much workers are exposed to these restrictions through the jobs they end up in.

At the aggregate level, the inflexibility index is relatively high, with a mean of 0.799 and a standard deviation of 0.051, indicating that, on average, occupations offer limited flexibility in how work can be adjusted. However, there is substantial heterogeneity across occupational groups. The most inflexible occupations are Community and Social Service Occupations, Healthcare Practitioners and Technical Occupations, and Management Occupations, each exhibiting mean values above 0.83. In contrast, Material Moving Occupations, Building and Grounds Cleaning and Maintenance Occupations, Food Preparation and Serving Occupations, and Farming, Fishing, and Forestry Occupations lie at the lower end of the distribution, with mean inflexibility below 0.70.

Importantly, the index also exhibits meaningful within-occupation dispersion. This is reflected in standard deviations within several broad occupational groups, particularly Production Occupations, Construction and Extraction Occupations, and Life, Physical, and Social Science Occupations. By contrast, Legal Occupations and Architecture and Engineering Occupations display relatively tight distributions around their means, suggesting greater internal homogeneity in job structure within these categories.

⁷15 occupations in the pooled ACS 2009 - 2019 data do not merge with the O*NET data. These include individuals under 16 years of age or not in the labor force, the unemployed and four military occupations. In addition, nine occupations are present in the ACS but are not available in O*NET, as no occupational characteristic data exist for these groups. We drop the military occupations and all occupations for which O*NET data are not available.

Table 3.1: *Occupational Inflexibility by Occupation Category*

Occupation	Mean	SD	Min	Percentiles					Max	Share
				10th	25th	50th	75th	90th		
Community and Social Service Occupations	0.844	0.019	0.800	0.818	0.836	0.836	0.863	0.863	0.863	0.049
Healthcare Practitioners and Technical Occupations	0.837	0.038	0.712	0.792	0.838	0.838	0.838	0.874	0.925	0.136
Management Occupations	0.831	0.026	0.750	0.810	0.810	0.824	0.851	0.873	0.878	0.141
Sales and Related Occupations	0.809	0.059	0.659	0.742	0.816	0.817	0.833	0.895	0.899	0.071
Legal Occupations	0.808	0.013	0.786	0.786	0.796	0.816	0.816	0.816	0.816	0.026
Business and Financial Operations Occupations	0.804	0.024	0.702	0.781	0.805	0.805	0.807	0.820	0.893	0.101
Education Instruction and Library Occupations	0.797	0.028	0.699	0.747	0.794	0.810	0.810	0.815	0.816	0.214
Protective Service Occupations	0.787	0.038	0.606	0.760	0.765	0.805	0.805	0.827	0.845	0.007
Arts, Design, Entertainment, Sports, and Media Occupations	0.772	0.039	0.717	0.717	0.744	0.771	0.806	0.815	0.875	0.038
Personal Care and Service Occupations	0.765	0.039	0.580	0.712	0.747	0.769	0.788	0.806	0.834	0.018
Construction and Extraction Occupations	0.763	0.064	0.597	0.682	0.725	0.750	0.849	0.849	0.855	0.001
Office and Administrative Support Occupations	0.761	0.043	0.595	0.702	0.726	0.771	0.799	0.799	0.877	0.101
Life, Physical, and Social Science Occupations	0.754	0.060	0.657	0.690	0.706	0.730	0.838	0.838	0.838	0.021
Installation, Maintenance, and Repair Occupations	0.753	0.033	0.666	0.723	0.738	0.748	0.783	0.792	0.819	0.001
Transportation Occupations	0.741	0.057	0.575	0.675	0.692	0.772	0.777	0.843	0.843	0.003
Computer and Mathematical Occupations	0.736	0.038	0.674	0.674	0.701	0.740	0.766	0.800	0.800	0.026
Healthcare Support Occupations	0.733	0.022	0.698	0.722	0.722	0.722	0.734	0.761	0.846	0.010
Architecture and Engineering Occupations	0.728	0.018	0.686	0.711	0.712	0.732	0.734	0.756	0.771	0.011
Production Occupations	0.716	0.069	0.542	0.632	0.685	0.716	0.751	0.825	0.825	0.008
Material Moving Occupations	0.673	0.045	0.580	0.606	0.606	0.701	0.701	0.701	0.754	0.003
Building and Grounds Cleaning and Maintenance Occupations	0.665	0.060	0.602	0.602	0.602	0.644	0.681	0.763	0.806	0.003
Food Preparation and Serving Related Occupations	0.656	0.042	0.538	0.631	0.631	0.631	0.673	0.717	0.771	0.011
Farming, Fishing, and Forestry Occupations	0.655	0.058	0.569	0.622	0.622	0.622	0.708	0.763	0.763	0.001
Total	0.799	0.051	0.538	0.725	0.781	0.810	0.828	0.851	0.925	1.000

Notes: Share reports the weighted proportion of high-skilled native women employed in each occupation category relative to the total sample of 1,720,363 observations.

The distribution of high-skilled native women across occupations reveals a key feature of the labor market: occupational sorting is concentrated in a subset of mid-to-high inflexibility occupations. Education Instruction and Library Occupations alone account for 21.4 percent of the sample, making it the single largest occupational group. Together with Management and Healthcare Practitioners, these three categories account for approximately half of all high-skilled native women. Notably, these occupations are not uniformly located at the extremes of the inflexibility distribution: Education occupations lie close to the sample mean, while Management and Healthcare Practitioners are among the most inflexible categories. This mismatch between employment concentration and the extreme tails of the inflexibility distribution is important, as it implies that aggregate exposure to occupational constraints is driven by a subset of large but not necessarily extreme occupations.

Table 3.2 presents summary statistics for the key variables used in the analysis. On average, recent low skilled immigrants account for about 1.6% of the predicted working age population, although this share varies substantially across PUMAs and census years, ranging from 0 to 12.1%. The corresponding measure for recent female low skilled immigrants is smaller on average, at approximately 0.7%, with values ranging from 0 to 5.6%. The instru-

mental variables constructed to predict these immigration measures exhibit similar variation across PUMAs and census years. Overall, the summary statistics indicate considerable geographic and temporal heterogeneity in immigrant inflows, providing the variation exploited in the empirical analysis.

Table 3.2: *Summary Statistics for Explanatory Variables and Labor Market Outcomes*

Variable	Mean	Standard Deviation	Minimum	Maximum	Observations
PUMA - Level Variables					
<i>Immigration Measures</i>					
Recent Low Skilled Immigrants	0.016	0.014	0	0.121	1066
Recent Female Low Skilled Immigrants	0.007	0.006	0	0.056	1066
<i>Instrumental Variables</i>					
IV for Recent Immigrants	0.014	0.015	0	0.119	1066
IV for Recent Female Immigrants	0.007	0.007	0	0.055	1066
<i>PUMA-Level Controls</i>					
Mean Income of Males with College (Dollars)	93940.017	23173.502	41181.453	214466.438	1066
Proportion Married (Age Group, PUMA, Year)	0.571	0.093	0.153	0.815	1066
Proportion Black (Age Group, PUMA, Year)	0.126	0.142	0.001	0.933	1066
Proportion Other Race (Age Group, PUMA, Year)	0.007	0.018	0	0.563	1066
<i>Labor Market Outcomes</i>					
Childcare Labor Supply	0.009	0.002	0.003	0.036	1066
Median Hourly Wage of Childcare Workers (Dollars)	10.496	1.935	5.567	25.087	1066
Individual - Level Variables					
<i>Education</i>					
Graduate Degree	0.363	0.481	0	1	1720363
<i>Marital Status</i>					
Currently Married	0.599	0.490	0	1	1720363
Were Married	0.143	0.350	0	1	1720363
Never Married	0.259	0.438	0	1	1720363
<i>Race</i>					
White	0.861	0.346	0	1	1720363
Black	0.097	0.295	0	1	1720363
Others	0.043	0.202	0	1	1720363

Notes: PUMA $\times$ Census Year level variables are aggregated at the PUMA $\times$ Census Year level and weighted by the PUMA population. The nine age groups used to construct the Age Group $\times$ PUMA $\times$ Census Year variables are 20–24, 25–29, . . . , 60–64. The childcare labor supply variable, defined as the share of childcare workers in the PUMA labor force, is weighted by the PUMA labor force size. The median wage of childcare workers is weighted by the number of childcare workers in each PUMA. All individual level variables are weighted using person weights.

The childcare labor supply measure, defined as total employment in childcare occupations among both natives and immigrants divided by the total PUMA labor force, has a mean of 0.009, indicating that childcare accounts for a relatively small share of overall local employment. Consequently, even economically meaningful changes in childcare employment correspond to relatively small changes in this measure. The median wage of childcare workers is calculated as the PUMA level median hourly wage of native and pre existing immigrant childcare workers and has a mean of 10.496 dollars per hour. The variation in this measure across PUMAs reflects meaningful differences in local childcare wage levels.

The PUMA level control variables reflect substantial heterogeneity in local economic and demographic environments. Mean income among college-educated men is about \$93,940, with considerable variation across areas, indicating large differences in local wage structures and economic conditions. The proportion married, measured within the relevant age group in each PUMA, averages 0.571 and captures variation in marital patterns among women across localities. Similarly, the racial composition variables are defined at the age-group level within PUMAs, with average shares of 0.126 for Black individuals and 0.007 for other race groups, reflecting differences in the demographic composition of women’s age cohorts across locations. Together, these controls capture systematic differences in local labor market conditions and demographic structure that may be correlated with both immigration patterns and labor supply decisions of high-skilled women.

The individual-level control variables show that 36.3% of the sample holds a graduate degree, while a substantial share is currently married (0.599), with smaller proportions previously married (0.143) or never married (0.259). The distribution of race indicates that most individuals identify as white (0.861), with smaller shares identifying as Black (0.097) or other races (0.043). These patterns suggest that the sample is highly educated and predominantly white and married, while still exhibiting meaningful variation in marital status and race across individuals.

3.4 Mechanism - Analysis of Childcare Market

An increase in the supply of low skilled immigrants is expected to raise the supply of childcare workers within the PUMA and thereby reduce wages in childcare occupations. A decrease in immigration is expected to have the opposite effect. To assess whether this relationship holds at the PUMA level, we pool data from the 2009 ACS and 2019 ACS and estimate the following regression model.

$$Y_{p,T} = \alpha_1 \text{LSI}_{p,T} + \alpha_2 \text{Income}_{p,T} + \mu_p + \mu_{r,T} + \varepsilon_{p,T} \quad (1)$$

We estimate two specifications of this equation. In the first specification, $Y_{p,T}$ is the proportion of childcare workers in the labor force of PUMA p in census year T . In the second specification, $Y_{p,T}$ is the log of median wages of childcare workers in PUMA p in census year T .

$LSI_{p,T}$ denotes recent low skilled immigrants in PUMA p in census year T , expressed as a fraction of PUMA p 's predicted working age population. Recent low skilled immigrants are defined as those who arrived after the year 2000 in the 2009 ACS and those who arrived after 2009 in the 2019 ACS. The variable $Income_{p,T}$ is the log of per capita income among working age male college graduates in PUMA p in census year T . We include income of male college graduates because they are more likely to demand household services, while their income is less directly linked to wages in low skilled occupations. We exclude female income because both earnings and labor supply of women may respond directly to wages in childcare occupations.

We include fixed effects for PUMA and region by census year. The term μ_p denotes PUMA fixed effects, and $\mu_{r,T}$ denotes census region by census year fixed effects. The error term is given by $\varepsilon_{p,T}$.

3.4.1 Identification Strategy: Endogeneity and Instrumental Variables

The parameter α_1 in the OLS regression may be endogenous because the location choices of immigrants are not exogenous. Endogeneity may arise if immigrants choose to settle in areas with improving economic conditions. In such cases, higher wages can act as a ‘pull factor’ for low skilled immigrants seeking employment. This can generate a positive correlation between $LSI_{p,T}$ and the proportion of child care workers in a PUMA, as well as between $LSI_{p,T}$ and wages in child care occupations.

The presence of endogeneity motivates the use of an instrumental variable approach. The instrument we use is based on the tendency of immigrants to settle in locations where individuals from the same country of birth are already concentrated, as shown by [Bartel](#)

(1989). We use a version of the shift share instrument described in [Tabellini \(2020\)](#), which can be defined mathematically as follows:⁸

$$IV_{p,T} = \frac{1}{\hat{P}_{p,T}} \sum_b \left(\frac{N_{p,2000}^b}{N_{2000}^b} \times NLS_T^{(b,-p)} \right) \quad (2)$$

where $\hat{P}_{p,T}$ is the predicted working age population of PUMA p in census year T , defined as follows:

$$\hat{P}_{p,T} = WAPOP_{p,2000} \times g_t^{-R} \quad (3)$$

Here, $WAPOP_{p,2000}$ denotes the working age population of PUMA p in the 2000 Census, and g_t^{-R} captures the growth rate of the U.S. working age population excluding Census region R , where PUMA p is located.⁹

The term $N_{p,2000}^b$ in equation (2) denotes the total number of immigrants from birthplace b who resided in PUMA p in 2000, while N_{2000}^b is the total number of immigrants from birthplace b in the United States in 2000. $NLS_T^{(b,-p)}$ represent the total number of recent low skilled immigrants from birthplace b in census year T , excluding those who eventually settled in PUMA p . This includes individuals who arrived after the year 2000 in the 2009 ACS and those who arrived after 2009 in the 2019 ACS.

As in [Cortés \(2008\)](#), the validity of the instrument relies on two key assumptions. First, the historical distribution of immigrants in census 2000 must be uncorrelated with subsequent differential changes in local economic conditions that affect both childcare market outcomes and occupational choices across PUMAs. Second, differential changes across PUMAs should not influence the aggregate inflow of low-skilled immigrants to the United States. Although these assumptions cannot be tested directly, several studies provide evidence supporting their plausibility.

Regarding the first condition, [Cortés and Tessada \(2011\)](#), who employ a similar shift-share instrument to study the impact of low-skilled immigration on the labor market outcomes of

⁸We again exclude Alaska and Hawaii and drop individuals whose place of birth is in U.S. outlying territories, as well as immigrants who do not report a well defined country of birth.

⁹The growth rate is defined as $g_t^{-R} = \frac{WAPOP_T^{-R} - WAPOP_{T-1}^{-R}}{WAPOP_{T-1}^{-R}}$.

high-skilled native women, show that their results are robust to controlling for a range of pre-existing local characteristics interacted with census year effects. Their findings suggest that the historical settlement patterns underlying the instrument are unlikely to be capturing differential local labor market trends that independently affect the outcomes of interest.

Regarding the second condition, a potential violation arises if low-skilled immigrants choose destinations in response to city-specific pull factors. [Boustan \(2010\)](#) assesses the quantitative importance of this concern by comparing instruments constructed using actual immigrant flows with those based on predicted flows driven by source-country push factors. She finds little difference between the resulting estimates, suggesting that destination-specific pull factors are not a primary driver of the instrument’s variation.

3.4.2 Evidence from the Childcare Market Mechanism

Previous research has already documented that low skilled immigration can expand the supply of household services and reduce their cost. For example, [Cortés and Tessada \(2011\)](#) find that increases in low skilled immigration raised the use of housekeeping services among highly skilled households, consistent with immigration lowering the price of these services. Similarly, [Furtado \(2016\)](#) shows that immigrant inflows increase the supply of childcare workers and reduce childcare wages, providing evidence that immigration can lower childcare costs. Motivated by these findings, we first examine whether low skilled immigration is associated with greater childcare labor supply and lower childcare wages in my sample. This exercise serves as a validation of the proposed childcare market mechanism before turning to its implications for occupational choice.

To estimate the causal impact of low skilled immigration on local childcare markets at the PUMA level, we estimate equation (1) using a two stage least squares (2SLS) framework, instrumenting for $LSI_{p,T}$ with the shift share instrument defined in equation (2). We consider two alternative measures of $LSI_{p,T}$: a baseline measure based on all recent low skilled immigrants and a female specific measure restricted to recent low skilled female immigrants. For the latter specification, the shift share instrument is constructed analogously using recent

female low skilled immigrants.

Table 3.3 reports the corresponding 2SLS estimates for equation (1) across two outcome variables, labor supply and median wages. For each outcome, the model is estimated using both definitions of $LSI_{p,T}$, with the baseline and female columns reporting the respective estimates for labor supply and median wages.

We find that the estimated coefficients on $LSI_{p,T}$ in the labor supply specifications are positive, although none are statistically significant. Nevertheless, they consistently suggest that increases in low skilled immigration at the PUMA level are associated with higher childcare labor supply. The coefficient based on the baseline definition of $LSI_{p,T}$ is 0.028, which is smaller than the corresponding estimate in Furtado (2016) (0.044), obtained using the same definitions of the dependent and independent variables in a comparable baseline specification but based on Census data from 1980–2000. In contrast, our coefficient based on the female measure is approximately 2.8 times larger than our baseline estimate and about 1.8 times larger than that reported in Furtado (2016), highlighting the more pronounced role of female immigrants in childcare provision.

The estimated coefficient for the baseline (female) specification is 0.028 (0.079), implying that a one percentage point increase in the share of a PUMA’s predicted population made up of all (female) recent immigrants increases the childcare employment share of the PUMA labor force by 0.028 (0.079) percentage points. In standard deviation terms, a one standard deviation increase 0.014 (0.006) in the all (female) immigrant share increases labor supply by 0.000392 (0.000474), which corresponds to approximately 19.6% (23.7%) of one standard deviation of labor supply (0.002; ref Table 3.2).¹⁰ Notably, although a one standard deviation increase in the female measure represents a substantially smaller increase in the immigrant share of the PUMA population than a one standard deviation increase in the baseline measure, it produces an even larger increase in childcare labor supply. This finding suggests that female low skilled immigrants play a disproportionately important role in driving childcare market dynamics.

For median wages, we consistently find negative coefficients for both the baseline and

¹⁰ $0.000392 \div 0.002 \approx 0.196$; $0.000474 \div 0.002 \approx 0.237$

Table 3.3: *Effects of Low Skilled Immigration on Childcare Labor Supply and Wages*

Variable	Labor Supply		Median Wages	
	Baseline	Female	Baseline	Female
Low Skilled Immigrants (PUMA, Year)	0.028 (0.023)	0.079 (0.057)	-0.668 (1.419)	-1.396 (3.367)
Log Mean Income of Males with College (PUMA, Year)	-0.000 (0.001)	-0.000 (0.001)	0.049 (0.078)	0.050 (0.077)
First-stage Coefficient	0.497*** (0.053)	0.505*** (0.053)	0.460*** (0.047)	0.482*** (0.060)
First-stage F	86.716	90.273	96.816	65.200
Mean Dep. Variable	0.009	0.009	2.303	2.303
Observations	2132	2132	2132	2132

* $p \leq 0.05$, ** $p \leq 0.01$, *** $p \leq 0.001$

Notes: Standard errors, clustered at the PUMA level, are reported in parentheses. All specifications include PUMA fixed effects, and census region $\times$ census year fixed effects. The labor supply regressions are weighted by the PUMA level labor force size, while the median wage regressions are weighted by the number of childcare workers in each PUMA.

female specifications of $LSI_{p,T}$. Although both estimates are statistically insignificant, they suggest that increases in the supply of childcare workers are associated with lower wages in the childcare market. The coefficients under the baseline and female specifications are -0.668 and -1.396, respectively, implying that the female estimate is approximately 2.1 times larger in magnitude than the baseline estimate. [Furtado \(2016\)](#) reports a coefficient of -4.281 using the same definitions of the dependent and independent variables as in our baseline specification. She notes that this estimate is considerably larger in magnitude than most existing estimates of the wage effects of low skilled immigration at the time.

The coefficient on the baseline (female) $LSI_{p,T}$ is -0.668 (-1.396), representing the percentage change in median wages of childcare workers associated with a one percentage point increase in the share of all (female) recent low skilled immigrants at the PUMA level. A one standard deviation increase in all (female) recent low skilled immigration, equal to 0.014 (0.006), which represents a typical change in exposure, is associated with a corresponding decrease in median wages of approximately 0.93% (0.84%). Although a one standard deviation increase in the female measure corresponds to a much smaller increase in the immigrant share of the PUMA population, it generates a reduction in childcare wages that is comparable in magnitude to that of the baseline measure. This pattern is consistent with female

low skilled immigrants having a particularly strong effect on childcare markets through their direct participation in childcare occupations.

The lack of statistically significant effects of recent low-skilled immigration on either childcare labor supply or median wages is somewhat surprising in light of earlier findings. [Furtado \(2016\)](#) reports statistically significant effects for both outcomes using Census data from 1980–2000. However, when she replicates the analysis using ACS data from 2007–2011, she finds that the estimates become considerably noisier, with substantially larger standard errors, preventing strong conclusions. She attributes this difference to the nature of ACS data collection. Unlike the decennial Census, which captures information over a relatively short reference period, the ACS is collected continuously over several years. As a result, measures of recent immigrant concentration constructed from the ACS may provide a noisier proxy for the immigrant environment faced by individuals when making labor market and family decisions.

This concern is reinforced by [Grieco et al. \(2018\)](#), who show that the ACS methodology can substantially understate the number of recent immigrants. They find that estimates of immigrants who report arriving in the survey year are underestimated by roughly 50 percent in single-year ACS files. Moreover, when annual ACS files are combined into multi-year datasets, estimates of immigrants who arrived within the previous five years can be underestimated by as much as 90 percent, depending on how recently the arrival occurred. Because our measure of low-skilled immigration is based on recent immigrant arrivals and is constructed using annual ACS files combined into multi-year datasets, this undercounting may weaken the measured relationship between immigration and childcare market outcomes, contributing to the statistically insignificant estimates reported in this essay.

Taken together, the results in Tables 3 suggest a coherent mechanism. When the share of low skilled immigrants increases in a PUMA, the supply of childcare workers rises. At the same time, wages in the childcare sector decline. This joint pattern indicates that the wage decline operates primarily through an expansion in labor supply rather than alternative channels. In other words, the increase in low skilled immigrants shifts the supply of childcare labor outward, and this supply response translates into lower equilibrium wages in the local

childcare market.

From the perspective of high skilled native women, the decline in childcare wages translates into lower childcare costs. This change increases the opportunity cost of supplying labor, in the sense that not working or working fewer hours becomes more costly when childcare services are cheaper. The existing literature has primarily focused on this channel, showing how lower childcare costs affect the opportunity cost of not working or reducing hours of work.

In this essay, we extend this idea to a different margin. We argue that lower childcare costs may also affect the opportunity cost of not working in a time inflexible occupation. When childcare becomes more affordable, the cost of committing to jobs that require fixed and less flexible hours becomes relatively less burdensome. As a result, high skilled native women may find it more attractive to select into occupations that are more time inflexible, since the effective cost of providing inflexible labor hours is reduced through cheaper childcare arrangements. In the next section, we examine whether these changes in incentives lead high skilled women to shift toward more time inflexible occupations.

3.5 Low-Skilled Immigration and Occupational Flexibility

For the empirical analysis of the impact of low skilled immigrants on the occupational choices of high skilled native women, we estimate the following model:

$$O_{i,p,t} = \alpha LSI_{p,T} + X'_{i,p,t}\beta + \Omega'_{p,T}\gamma + \theta_a + \theta_p + \theta_t + \theta_{r,T} + u_{i,p,t} \quad (4)$$

Here, $O_{i,p,t}$ denotes the job inflexibility index for high skilled woman i living in PUMA p in survey year t . The variable $LSI_{p,T}$ represents the share of recent low skilled immigrants in the predicted working age population of PUMA p in census year T .¹¹ The vector $X_{i,p,t}$

¹¹Census year T refers to the two ACS five year samples used in the analysis, namely ACS 2005 - 2009 and ACS 2015 - 2019. Survey year t refers to the individual calendar years within each ACS sample.

includes individual level controls such as an indicator for graduate degree, marital status indicators for currently married and previously married individuals (divorced, widowed, or separated), and race indicators for Black and other race individuals. The omitted categories are never married and White individuals.

The vector $\Omega_{p,T}$ includes PUMA-level controls measured in PUMA p in census year T that capture labor market and demographic characteristics relevant to the occupational choices of high-skilled women. These controls include the log of average income among college educated males in the PUMA, the proportion of women in the same age group who are married, the proportion who are Black, and the proportion belonging to other racial groups. These variables help account for differences in social norms and preferences related to fertility and labor supply decisions across local labor markets.

Finally, θ_a denotes age fixed effects (by individual year), θ_p denotes PUMA fixed effects, θ_t denotes survey year fixed effects, and $\theta_{r,T}$ denotes Census region by census year fixed effects.

Given the included fixed effects and individual and PUMA level controls, the coefficient α captures whether PUMAs that experienced larger increases in low skilled immigration also experienced larger increases in occupational inflexibility among high skilled native women relative to other PUMAs in the same Census region.

Although equation (4) includes a rich set of fixed effects, the OLS estimate of α may still suffer from endogeneity. As in equation (1), one concern is that low skilled immigrants may choose to settle in PUMAs with stronger economic conditions, which may independently affect the occupational choices of high skilled women. In addition, reverse causality may arise if an increase in high skilled women working in less flexible occupations raises the local demand for household services and, in turn, attracts more low skilled immigrants to the same PUMA.

OLS estimates may also be biased by unobserved changes in fertility preferences or household norms. For example, a local shift toward larger family size may simultaneously increase the demand for childcare services and encourage high skilled women to select into more flexible occupations. In this case, preferences related to fertility and family structure become a

confounding factor that generates a spurious relationship between low skilled immigration and occupational choice. Although we control for marital status and race, these variables are only imperfect proxies for preferences over fertility, motherhood, and work life balance. More importantly, such preferences and social norms may evolve over time, including between 2009 and 2019, and therefore may not be fully absorbed by the fixed effects structure.

Table 3.4: *Impact of LSIs on Occupational Inflexibility among High Skilled Native Women*

Variable	OLS (1)	OLS (2)	IV (1)	IV (2)
Low Skilled Immigrants, (PUMA, Year)	0.009 (0.009)	-0.014 (0.012)	-0.011 (0.026)	-0.052 (0.074)
Graduate Degree	0.019*** (0.000)	0.019*** (0.000)	0.019*** (0.000)	0.019*** (0.000)
Currently Married	0.008*** (0.000)	0.008*** (0.000)	0.008*** (0.000)	0.008*** (0.000)
Were Married	0.005*** (0.000)	0.005*** (0.000)	0.005*** (0.000)	0.005*** (0.000)
Black	-0.003*** (0.000)	-0.003*** (0.000)	-0.003*** (0.000)	-0.003*** (0.000)
Others	-0.003*** (0.000)	-0.003*** (0.000)	-0.003*** (0.000)	-0.003*** (0.000)
Log Mean Income of Males with College, (PUMA, Year)	0.004*** (0.001)	0.004** (0.001)	0.004** (0.001)	0.004** (0.001)
Proportion Married (Age Group, PUMA, Year)	-0.006** (0.002)	0.004*** (0.001)	0.004*** (0.001)	0.005*** (0.001)
Proportion Black (Age Group, PUMA, Year)	0.000 (0.001)	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)
Proportion Other Race (Age Group, PUMA, Year)	0.000 (0.005)	-0.012 (0.011)	-0.012 (0.011)	-0.012 (0.011)
Age FE	Yes	Yes	Yes	Yes
PUMA FE		Yes	Yes	Yes
Survey Year FE	Yes	Yes	Yes	Yes
Region $\times$ Year FE	Yes	Yes	Yes	Yes
First stage Coefficient			0.639*** (0.066)	0.549*** (0.068)
First stage F Statistic			93.352	66.134
Mean Dependent Variable	0.799	0.799	0.799	0.799
Observations	1720363	1720363	1720363	1720363

Notes: Standard errors clustered at the PUMA level are reported in parentheses. The nine age groups used to construct the Age Group by PUMA by Year variables are 20–24, 25–29, ..., 60–64. We also estimate all specifications using a linear and squared term in age instead of age fixed effects and obtain similar results.

To address these concerns, we use an instrumental variables approach. We employ the shift share instrument defined in equation (2) to exploit variation in immigrant inflows that is driven by historical settlement patterns. This variation is less likely to be related to

current economic conditions in the PUMA or to current changes in fertility preferences of high skilled native women.

The first two columns of Table 3.4 report the results for two OLS specifications. The third column reports the results from the 2SLS specification, where we instrument for $LSI_{p,T}$ using the shift share instrument defined in equation (2). In the final column, we again estimate the model using 2SLS, but redefine $LSI_{p,T}$ as the fraction of the PUMA's predicted female working age population made up of recent female low skilled immigrants. The instrumental variable is redefined accordingly for this specification.

The OLS estimate on $LSI_{p,T}$ in the first column, where we do not include PUMA fixed effects, is positive and statistically insignificant, with a coefficient of 0.009. This suggests that a one percentage point increase in the share of recent low skilled immigrants in a PUMA is associated with a 0.009 percentage points increase in the occupational inflexibility index. Using the observed variation in immigration across PUMAs, a one standard deviation increase in $LSI_{p,T}$ (0.014) is associated with a 0.000126 increase in the occupational inflexibility index. Given the sample standard deviation of occupational inflexibility (0.051), this corresponds to approximately 0.25% of one standard deviation.

Once we include PUMA fixed effects in the second column, however, the coefficient becomes negative. This change suggests that the positive relationship in the first specification was likely driven by persistent differences across PUMAs rather than changes within PUMAs over time. In particular, PUMAs with historically larger immigrant populations may also tend to have labor markets where high skilled women are employed in more time inflexible occupations for reasons unrelated to recent immigration. By including PUMA fixed effects, the estimation instead relies on within PUMA changes over time, comparing how changes in low skilled immigration within the same PUMA are associated with changes in occupational inflexibility. The reversal in the sign of the coefficient therefore indicates that cross sectional differences across PUMAs were masking the underlying within PUMA relationship.

The third column presents the 2SLS estimate. The first-stage results indicate that the instrument is strongly correlated with the endogenous variable. The estimated first-stage coefficient is 0.639 and is statistically significant at conventional levels, implying that a

one-percentage-point increase in the instrument is associated with a 0.639-percentage-point increase in the share of recent low-skilled immigrants. Consistent with this strong relationship, the first-stage F-statistic is 93.352, well above conventional thresholds used to assess weak instruments.

The estimated coefficient on $LSI_{p,T}$ is again negative, small in magnitude, and statistically insignificant. The negative sign suggests that, after isolating plausibly exogenous variation in low skilled immigration, increases in immigrant inflows are not associated with greater occupational inflexibility among high skilled native women. More importantly, the lack of statistical significance indicates that we cannot reject the null hypothesis of no effect.

In the final specification, we focus specifically on female low-skilled immigration to examine whether the relationship is stronger when the immigration measure more closely captures workers who are likely to substitute for unpaid household labor performed by women (Cortés, 2023). The first-stage results continue to indicate a strong relationship between the instrument and the endogenous variable. The estimated first-stage coefficient is 0.549 and is statistically significant at conventional levels, implying that a one-percentage-point increase in the instrument is associated with a 0.549-percentage-point increase in the share of recent low-skilled female immigrants. Consistent with this strong relationship, the first-stage F-statistic is 66.134, well above conventional thresholds used to assess weak instruments. The estimated second-stage coefficient remains negative and statistically insignificant, suggesting that even female-specific low-skilled immigration does not lead to a statistically meaningful increase in occupational inflexibility among high-skilled native women.

3.5.1 Heterogeneous Effects by Household Childcare Needs

The sample used in our main analysis includes high skilled women with very different family structures and preferences over work and childcare. It includes younger women who are not yet married or do not have children and therefore may have little need for domestic services to increase labor supply. It also includes mothers with children of different ages. Some women have children who are already of school going age (5 to 17 years old) and may require

relatively less childcare, particularly those with older children. In contrast, mothers with children younger than five years old are likely to have the greatest demand for childcare and other household services. Since the need for domestic services differs substantially across these groups of women, the effect of low skilled immigration on occupational choice is also likely to vary across them.

Table 3.5: *Impact of LSIs on Occupational Inflexibility by Family Structure*

Variable	No Children		School-Aged Children		Children Under 5	
	Baseline	Female	Baseline	Female	Baseline	Female
Low Skilled Immigrants, (PUMA, Year)	-0.037 (0.037)	-0.102 (0.104)	0.016 (0.042)	0.041 (0.116)	0.044 (0.070)	0.071 (0.206)
Graduate Degree	0.019*** (0.000)	0.019*** (0.000)	0.018*** (0.000)	0.018*** (0.000)	0.019*** (0.000)	0.019*** (0.000)
Currently Married	0.007*** (0.000)	0.007*** (0.000)	0.006*** (0.001)	0.006*** (0.001)	0.009*** (0.001)	0.009*** (0.001)
Were Married	0.004*** (0.000)	0.004*** (0.000)	0.005*** (0.001)	0.005*** (0.001)	0.006*** (0.001)	0.006*** (0.001)
Black	-0.003*** (0.000)	-0.003*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)	-0.002* (0.001)	-0.002* (0.001)
Others	-0.003*** (0.000)	-0.003*** (0.000)	-0.003*** (0.001)	-0.003*** (0.001)	-0.001 (0.001)	-0.001 (0.001)
Log Mean Income of Males with College, (PUMA, Year)	0.007*** (0.002)	0.007*** (0.002)	-0.000 (0.002)	-0.000 (0.002)	0.004 (0.004)	0.004 (0.004)
Proportion Married (Age Group, PUMA, Year)	0.006*** (0.002)	0.006*** (0.002)	0.008** (0.003)	0.008** (0.003)	0.001 (0.004)	0.001 (0.004)
Proportion Black (Age Group, PUMA, Year)	0.003 (0.003)	0.003 (0.003)	0.012** (0.004)	0.012** (0.004)	-0.002 (0.008)	-0.002 (0.008)
Proportion Other Race (Age Group, PUMA, Year)	-0.020 (0.017)	-0.020 (0.017)	-0.028 (0.019)	-0.028 (0.019)	-0.011 (0.034)	-0.012 (0.034)
First stage Coefficient	0.625*** (0.069)	0.520*** (0.073)	0.678*** (0.066)	0.621*** (0.063)	0.671*** (0.072)	0.577*** (0.072)
First stage F Statistic	81.753	50.498	104.227	97.856	86.449	64.748
Mean Inflexibility Index	0.796	0.796	0.802	0.802	0.804	0.804
Observations	929987	929987	516984	516984	140507	140507

Notes: Standard errors clustered at the PUMA level are reported in parentheses. All specifications include age fixed effects (by individual year), PUMA fixed effects, survey year fixed effects, and Census region by census year fixed effects. The nine age groups used to construct the Age Group by PUMA by Year variables are 20–24, 25–29, ..., 60–64. We also estimate all specifications using linear and squared terms in age instead of age fixed effects and obtain similar results. The first-stage F-statistics across specifications range from 50.498 to 104.227, indicating a strong relationship between the instrument and the endogenous variable in all specifications.

We consider three sub samples of high skilled native women. The first includes women without children. The second includes women with school age children, defined as children between the ages of 5 and 17. The third includes women with at least one child younger than five years old living at home. Thus, women with both a child younger than five and a child between the ages of 5 and 17 are included only in the second group, making the second and third groups mutually exclusive by construction.

For each of these sub samples, we estimate two 2SLS specifications. The first uses the baseline measure of $LSI_{p,T}$, together with the corresponding instrumental variable defined in equation (2). This specification is labeled “Baseline” in Table 3.5. The second specification instead focuses on female specific immigration exposure. The instrumental variable is modified accordingly for this specification, which is labeled “Female” in Table 3.5.

The first two columns report the results for the subsample of women without children. The estimated coefficients on $LSI_{p,T}$ under both the baseline and female specifications are negative and statistically insignificant. Compared with the IV estimates reported in the last two columns of Table 3.4, the coefficients become somewhat more negative, although they remain statistically insignificant. The absence of a meaningful effect for women without children is consistent with our proposed mechanism. Since these women are less likely to rely on childcare services, increases in low skilled immigration and the resulting decline in childcare costs are unlikely to substantially affect their occupational choices.

The next two columns report the results for the subsample of high skilled native women with school age children. In contrast to the estimates for women without children, the coefficients in both specifications become positive, although they remain statistically insignificant. The estimated effect is 2.6 times larger under the female specification (0.041) than under the baseline specification (0.016).

Using the baseline specification, a one standard deviation increase in low skilled immigration exposure (0.014) is associated with a 0.000224 increase in the occupational inflexibility index. Given that the standard deviation of occupational inflexibility in this sub sample is 0.049, this corresponds to approximately 0.46% of one standard deviation. Under the female specification, a one standard deviation increase in low skilled immigration exposure (0.006) is associated with a slightly larger increase of 0.000246 in the occupational inflexibility index, equivalent to approximately 0.50% of one standard deviation. Although the estimated effects are small relative to the 0 to 1 scale of the index, the similarity in their magnitudes is noteworthy because the one standard deviation increase under the female specification represents only a 0.6 percentage point increase in the immigrant share, compared with 1.4 percentage points under the baseline specification. In other words, an increase in recent female low

skilled immigration of less than half the size generates an effect on occupational inflexibility that is at least as large as that of an equivalent standard deviation increase in overall low skilled immigration. This pattern is consistent with the idea that female low skilled immigrants have a disproportionate influence on household labor allocation by substituting for unpaid household labor, thereby relaxing time constraints and enabling high skilled native women to sort into more occupationally inflexible jobs. Taken together, these estimates are consistent with the idea that women with school age children are more responsive to changes in the availability and cost of household services such as childcare.

The last two columns report results for the subsample of high skilled native women with at least one child under the age of five. This group represents mothers with the highest childcare needs, as their youngest child is not yet of school age and therefore requires intensive time inputs at home. Because of this, we expect this group to be the most sensitive to changes in the availability and cost of childcare services. Consistent with this expectation, the estimated coefficient under the baseline specification (0.044) is approximately 2.8 times larger than the corresponding estimate for women with school-age children. Similarly, the estimated coefficient under the female specification (0.071) is positive and about 1.7 times larger than that for women with school-age children. Although neither estimate is statistically significant, the larger magnitudes are consistent with the hypothesis that reductions in childcare constraints have the strongest effect on mothers with the youngest children.

In terms of economic magnitude, the baseline specification implies that a one standard deviation increase in low skilled immigration exposure (0.014) increases the occupational inflexibility index by 0.000616. Relative to the standard deviation of occupational inflexibility in this sub sample (0.049), this corresponds to approximately 1.26% of one standard deviation. Under the female specification, a one standard deviation increase in low skilled immigration exposure (0.006) increases the occupational inflexibility index by 0.000426, equivalent to approximately 0.87% of one standard deviation. Although these effects are modest relative to the 0 to 1 scale of the index, the female specification produces a response of comparable magnitude despite representing less than half the increase in immigrant exposure. This pattern provides further evidence that female low skilled immigrants play a disproportionately

important role in easing household time constraints by substituting for unpaid household labor, thereby enabling high skilled native women to sort into more inflexible occupations.

Taken together, this pattern is consistent with the idea that women with very young children face the strongest childcare related constraints in their labor market decisions. In such settings, changes in the availability and cost of household services are more likely to matter for occupational choices. However, since the estimates remain statistically insignificant, these patterns should be interpreted as suggestive rather than conclusive evidence.

3.6 Conclusion

Over the past five decades, women’s labor force participation and educational attainment have increased substantially, contributing to higher living standards and greater economic independence. Despite these gains, women, particularly those in high-paying occupations, continue to face trade-offs between employment and fertility, with job time inflexibility emerging as a key constraint. The availability of affordable childcare, which is often influenced by the supply of low skilled immigrants, has the potential to ease these trade-offs by reducing the cost of childrearing and increasing the feasibility of sustained labor market attachment.

This essay contributes to the literature by examining whether changes in the supply of low-skilled immigrants affect the occupational choices of high-skilled native women, with a particular focus on preferences for time-inflexible occupations. We conduct the analysis in two steps.

First, using individual level data from the 2009 and 2019 five year American Community Survey (ACS), we construct PUMA level measures of childcare labor supply and median hourly wages. We then exploit variation in the share of recent low skilled immigrants across PUMAs over time to estimate the impact of immigration on local childcare markets. Using a shift share instrumental variables strategy based on immigrant settlement patterns from the 2000 Census, we find suggestive evidence that increases in low skilled immigration expand childcare labor supply and reduce childcare costs at the local level. We also find that these relationships are stronger when immigration is measured using recent female low

skilled immigrants rather than all recent low skilled immigrants, highlighting the particularly important role of female immigrants in shaping local childcare market outcomes.

Second, we use five occupation-level characteristics from O*NET to construct a measure of occupational inflexibility at the occupation level and then merge this measure with ACS data. Using the same instrumental variable and a 2SLS approach, we find a negative but statistically insignificant effect of low-skilled immigration on occupational inflexibility in the overall sample of high-skilled native women.

We then divide the sample of high-skilled native women based on family structure, since reliance on childcare services differs substantially across these groups. Women without children are unlikely to demand childcare services, mothers with school-aged children rely less on childcare, while mothers with children under the age of five depend most heavily on these services. Accordingly, we split the sample into three subsamples and re-estimate our model separately for each group.

The results across the subsamples reveal a pattern consistent with the proposed mechanism, although the estimates remain statistically insignificant throughout. We find a negative and statistically insignificant effect for women without children, which is expected given their limited reliance on childcare services. Among mothers, however, the estimated relationship between low skilled immigration and occupational inflexibility becomes positive and increases in magnitude, with the largest effects observed for women with children under the age of five. The estimated effect for mothers with children under five is approximately 2.8 times larger than that for mothers with school-aged children, consistent with the greater childcare demands faced by families with younger children. We also estimate the model using a female specific measure of low skilled immigration. Despite representing a smaller increase in immigrant exposure, this specification produces effects of comparable or greater magnitude than the baseline measure, suggesting that female low skilled immigrants play a disproportionately important role. This pattern is consistent with the view that female immigrants substitute for unpaid household labor, thereby easing childcare constraints and enabling high skilled native women to sort into more occupationally inflexible occupations. Taken together, these findings support the hypothesis that childcare needs shape the extent

to which changes in local childcare costs influence occupational choices.

At the same time, the absence of statistical significance across specifications suggests that these findings should be interpreted as suggestive rather than conclusive. Overall, while the estimates do not provide precise causal evidence, the heterogeneity across groups is broadly consistent with the mechanism linking low-skilled immigration, childcare costs, and occupational choice among high-skilled native women.

Overall, the point estimates in all analyses are consistent with the proposed mechanisms, they are imprecisely estimated and do not allow for strong conclusions. The existing literature suggests that such noisy estimates may reflect limitations in the measurement of recent immigration, particularly in the American Community Survey, where continuous data collection and sampling design can introduce substantial measurement error in immigration variables (Furtado, 2016; Grieco, Larsen, and Hogan, 2018). In future work, it would be valuable to re examine these relationships using data sources with more precise and comprehensive coverage of immigrant populations, which may help improve statistical power and reduce measurement error in the key immigration variables.

References

- F. Alvarez, F. Lippi, and J. Passadore. Are state- and time-dependent models really different? *NBER Macroeconomics Annual*, 31(1):379–457, 2017.
- S.B. Aruoba, E. Oue, F. Saffie, and J. Willis. Reviving Micro Real Rigidities: The Importance of Demand Shocks. Technical Report w32518, National Bureau of Economic Research, 2026.
- Aviation at Lazarski University. Seasonality in air traffic. Lazarski Aviation Academy, 27 January, 2025. URL <https://aviation.lazarski.pl/seasonality-in-air-traffic/>. Accessed: January 05, 2026.
- L. Ball and N.G. Mankiw. Asymmetric price adjustment and economic fluctuations. *The Economic Journal*, 104(423):247–261, 1994.
- L. Ball and N.G. Mankiw. Relative-price changes as aggregate supply shocks. *The Quarterly Journal of Economics*, 110(1):161–193, 1995.
- G. Barone and S. Mocetti. With a little help from abroad: The effect of low-skilled immigration on the female labour supply. *Labour Economics*, 18(5):664–675, 2011.
- A.P. Bartel. Where do the new U.S. immigrants live? *Journal of Labor Economics*, 7(4):371–391, 1989.
- S. Bauernschuster, T. Hener, and H. Rainer. Children of a (policy) revolution: The introduction of universal child care and its effect on fertility. *Journal of the European Economic Association*, 14(4):975–1005, 2016.
- J.P. Berkowitz, K.S. Dasari, and R.J. DeLisle. The Effect of Fuel Hedging in the Airline

- Industry on Returns, Volatility, and on the Return-to-risk Relationship Analysis. *The Energy Journal*, 46(3):265–284, 2025.
- D.M. Blau and P.K. Robins. Fertility, employment, and child-care costs.
- A. Blinder, E.R. Canetti, D.E. Lebow, and J.B. Rudd. *Asking about prices: a new approach to understanding price stickiness*. Russell Sage Foundation, 1998.
- L.P. Boustan. Was postwar suburbanization “white flight”? Evidence from the black migration. *The Quarterly Journal of Economics*, 125(1):417–443, 2010.
- P. Bunn, N. Bloom, C. Menzies, P. Mizen, G. Thwaites, and I. Yotzov. State and Time-Dependent Pricing. Technical Report w34666, National Bureau of Economic Research, 2026.
- G.A. Calvo. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398, 1983.
- J.N. Cannon and O.V. Watanabe. Do firms pass commodity cost savings to consumers? Evidence and impacts of asymmetric pricing behavior in the US airline industry. SSRN, 2020. Evidence and Impacts of Asymmetric Pricing Behavior in the US Airline Industry (December 21, 2020).
- D.A. Carter, D.A. Rogers, and B.J. Simkins. Hedging and value in the US airline industry. *Journal of Applied Corporate Finance*, 18(4):21–33, 2006.
- S.H. Chuang. Cost pass-through in the airline industry: Price responses and asymmetries. *Economics Bulletin*, 40(1):639–652, 2020.
- P. Cortés. The effect of low-skilled immigration on U.S. prices: Evidence from CPI data. *Journal of Political Economy*, 116(3):381–422, 2008.
- P. Cortés. Immigration, household production, and native women’s labor market outcomes: A survey of a global phenomenon. *European Economic Review*, 158:104545, 2023.

- P. Cortés and J. Pan. Outsourcing household production: Foreign domestic workers and native labor supply in Hong Kong. *Journal of Labor Economics*, 31(2):327–371, 2013.
- P. Cortés and J. Tessada. Low-skilled immigration and the labor supply of highly skilled women. *American Economic Journal: Applied Economics*, 3(3):88–123, 2011.
- J. Dainauskas. Time-varying exchange rate pass-through into terms of trade. *Journal of International Money and Finance*, 137:102905, 2023.
- F. Dobruszkes, D. Ramos-Pérez, and J.M. Decroly. Reasons for flying. Technical Report 2013/284083, ULB–Universite Libre de Bruxelles, 2019.
- C. East and A. Velasquez. The effect of increasing immigration enforcement on the labor supply of high-skilled citizen women. Technical Report 12029, IZA Discussion Paper, 2019.
- D. Escobari. Asymmetric price adjustments in airlines. *Managerial and Decision Economics*, 34(2):74–85, 2013.
- L. Farré, L. González, and F. Ortega. Immigration, family responsibilities and the labor supply of skilled native women. *The B.E. Journal of Economic Analysis & Policy*, 11(1), 2011.
- E. Forlani, E. Lodigiani, and C. Mendolicchio. Natives and migrants in home production: The case of Germany. *Review of Economics of the Household*, 19(4):1275–1307, 2020.
- D. Furtado. Fertility responses of high-skilled native women to immigrant inflows. *Demography*, 53(1):27–53, 2016.
- P.G. Gayle and Y. Lin. Cost pass-through in commercial aviation: Theory and evidence. *Economic Inquiry*, 59(2):803–828, 2021.
- C. Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119, 2014.

- E.M. Grieco, L.J. Larsen, and H. Hogan. How Period Data Influence the Estimates of Recently Arrived Immigrants in the American Community Survey. *International Migration Review*, 52(1):299–313, 2018.
- B. Hornung. Does Hedging Success Matter? An Empirical Study of Jet Fuel Hedging in the US Airline Industry. USF Tampa Graduate Theses and Dissertations, University of South Florida, 2020.
- Y. Jiang, T. Que, and M. Yu. Price asymmetries in the US airline industry. *Financial Review*, 57(4):793–814, 2022.
- Ò. Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 2005.
- R.K. Kaufmann. Airfares and oil prices: ‘Feathers and rockets’ adjustments. *Energy Economics*, 68:515–521, 2017.
- L. Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069, 2009.
- C. Koopmans and R. Lieshout. Pass-through of airlines cost changes. What does economic research say? Technical report, 2014.
- C. Koopmans and R. Lieshout. Airline cost changes: To what extent are they passed through to the passenger? *Journal of Air Transport Management*, 53:1–11, 2016.
- S.H. Lim and Y. Hong. Fuel hedging and airline operating costs. *Journal of Air Transport Management*, 36:33–40, 2014.
- A. Mas and A. Pallais. Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–3759, 2017.
- R. Merkert and H. Swidan. Flying with (out) a safety net: Financial hedging in the airline industry. *Transportation Research Part E: Logistics and Transportation Review*, 127:206–219, 2019.

- R. Molloy, C.L. Smith, and A. Wozniak. Internal migration in the United States. *Journal of Economic Perspectives*, 25:173–196, 2011.
- P. Morrell and W. Swan. Airline jet fuel hedging: Theory and practice. *Transport Reviews*, 26(6):713–730, 2006.
- D. Pal and S.K. Mitra. Do airfares respond asymmetrically to fuel price changes? A multiple threshold nonlinear ARDL model. *Energy Economics*, 111:106113, 2022.
- V.A. Ramey. Macroeconomic shocks and their propagation. In *Handbook of Macroeconomics*, volume 2, pages 71–162. 2016.
- V.K. Rao. Fuel price risk management using futures. *Journal of Air Transport Management*, 5(1):39–44, 1999.
- J. Reynolds and D.R. Johnson. Don’t blame the babies: Work hour mismatches and the role of children. *Social Forces*, 91(1):131–155, 2012.
- M.W. Reynolds. Statement before the Aviation Subcommittee of the Transportation and Infrastructure Committee, U.S. House of Representatives. U.S. Department of Transportation, Washington, DC, 2008. 14 May.
- E. Samunderu, J.K. Perret, and G. Geller. The economic value rationale of fuel hedging: An empirical perspective from the global airline industry. *Journal of Air Transport Management*, 106:102324, 2023.
- D. Scotti and N. Volta. Price asymmetries in European airfares. *Economics of Transportation*, 14:42–52, 2018.
- Sri. Analysis and Visualization of Airline Flight Delays Data: Unveiling Patterns and Insights. Medium, 2023. URL <https://medium.com/>. 20 May.
- H. Swidan and R. Merkert. The relative effect of operational hedging on airline operating costs. *Transport Policy*, 80:70–77, 2019.

- H. Swidan, R. Merkert, and O.K. Kwon. Designing optimal jet fuel hedging strategies for airlines—Why hedging will not always reduce risk exposure. *Transportation Research Part A: Policy and Practice*, 130:20–36, 2019.
- M. Tabellini. Gifts of the immigrants, woes of the natives: Lessons from the age of mass migration. *The Review of Economic Studies*, 87(1):454–486, 2020.
- M. Tappata. *Consumer search, price dispersion, and asymmetric pricing*. PhD thesis, University of California, Los Angeles, 2006.
- J.B. Taylor. Aggregate dynamics and staggered contracts. *Journal of Political Economy*, 88(1):1–23, 1979.
- K.A. Weeden, Y. Cha, and M. Bucca. Long work hours, part-time work, and trends in the gender gap in pay, the motherhood wage penalty, and the fatherhood wage premium. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2(4):71, 2016.
- A.H. Wolter, T. Ehlers, K. Luetjens, and V. Gollnick. Commodity price pass-through in the US airline industry and the hidden perks of consolidation. *Journal of Air Transport Management*, 95:102100, 2021.