CERGE - EI
Center for Economic Research and Graduate Education –
Economics Institute

Essays in Applied Economics

Aizhamal Rakhmetova

Dissertation

Prague 2026

Dissertation Committee

Andreas Menzel, Ph.D. (University of Padua, Chair)

Sebastian Ottinger, Ph.D. (CERGE-EI, Local Chair)

Christian Ochsner, Ph.D. (CERGE-EI)

Referees

Marc Witte, Ph.D. (Vrije Universiteit Amsterdam)

Jonathan H. Conning, Ph.D. (Hunter College of The City University of New York)

To my mother, Galina Rakhmetova

Table of Contents

Abstract **vi**

Acknowledgments **viii**

Introduction **1**

1 Abortion Restrictions and Gender Gaps in Mortgage Credit **6**

1.1 Introduction 6

1.2 Institutional Background and Policy Variation 11

1.3 Data 13

1.3.1 HMDA Application-Level Data 13

1.3.2 Sample Construction 14

1.3.3 Outcome Variables 15

1.3.4 Borrower and Loan Controls 15

1.3.5 Summary Statistics 17

1.3.6 Lender Identifiers and Transmittal Sheet Data 18

1.3.7 Construction of State-Year Gender Gaps 19

1.3.8 Abortion Policy Classification and Treatment Timing 20

1.3.9 State-Level Controls 24

1.4 Conceptual Framework and Mechanisms 26

1.4.1 Household Balance Sheets and Mortgage Demand 26
1.4.2 Lender Screening and the Approval Margin 27
1.4.3 Heterogeneous Lender Responses 29
1.4.4 Selection into Approval and Conditional Interest Rates 30
1.4.5 Predicted Empirical Patterns 31
1.5 Empirical Strategy 34
1.5.1 Borrower-Level Triple-Difference Design 34
1.5.2 Interpretation of the Main Coefficient 36
1.5.3 Identification Assumption 37
1.5.4 Event-Study Specification 37
1.5.5 Inference 38
1.5.6 Synthetic Difference-in-Differences as Robustness 39
1.5.7 Summary 40
1.6 Results 40
1.6.1 Main Results: Borrower-Level Evidence 40
1.6.2 Event-Study Evidence and Pre-Trends 42
1.6.3 Parallel Trends in Gender Gaps 45
1.6.4 Heterogeneity: Evidence on the Credit Constraint Channel 47
1.6.5 Heterogeneity by Applicant Age 48
1.6.6 Heterogeneity by Political Environment of Lender Headquarters . 50
1.6.7 Observable Applicant Composition and Selection 54
1.6.8 Summary of Mechanism Evidence 57
1.7 Robustness and Placebo 57
1.7.1 Robustness Check: Synthetic Difference-in-Differences 57
1.7.2 Robustness to Excluding Texas 59
1.7.3 Robustness to Gender-Specific State and Year Fixed Effects . . . 61
1.7.4 Randomization Inference 62
1.8 Conclusion 65
Appendix A. State Classification 67
Appendix B. Synthetic Difference-in-Differences Graphical Diagnostics 68
Appendix C. Additional Mechanism and Robustness Tables 71

2 Access to Financial Resources and Environmental Migration of the Poor **75**

2.1 Introduction . . . 75
2.2 Literature Review . . . 78
2.3 Conceptual Framework . . . 82
2.4 Background and Data . . . 84
2.4.1 Measurement of Migration, Shocks, and Financial Resources . . . 88
2.5 Identification Strategy . . . 90
2.5.1 Definition of Rainfall Anomalies . . . 94
2.5.2 Checks for Rainfall Shocks . . . 96
2.5.3 Definition of Temperature Anomalies . . . 98
2.5.4 Checks for Temperature Shocks . . . 100
2.6 Results . . . 102
2.6.1 Rainfall Shocks . . . 102
2.6.2 Temperature Shocks . . . 107
2.6.3 Natural Disasters . . . 109
2.6.4 Change of Occupation . . . 111
2.6.5 Summary of Mechanism Evidence . . . 114
2.7 Robustness Checks . . . 116
2.8 Discussion and Policy Implications . . . 120
2.9 Conclusion . . . 121

3 Weather Shocks and Bride Kidnapping in Kyrgyzstan **123**

3.1 Introduction . . . 123
3.2 Background . . . 127
3.2.1 Climate in Kyrgyzstan . . . 127
3.2.2 Bride Kidnapping in Kyrgyzstan . . . 128
3.3 Data . . . 130
3.4 Identification Strategy . . . 134
3.5 Results . . . 135
3.5.1 Occurrences of Bride Kidnapping . . . 135
3.5.2 Attitudes toward Bride Kidnapping . . . 137
3.5.3 Heterogeneity Analysis . . . 139
3.6 Conclusion . . . 148
3.A Appendix . . . 150

Bibliography **152**

Abstract

This dissertation contains three essays in applied economics on how external shocks affect vulnerable individuals and households when economic, institutional, and social constraints limit their ability to adjust. The chapters study three different settings: abortion restrictions and mortgage credit in the United States, environmental shocks and migration in Indonesia, and weather shocks and bride kidnapping in Kyrgyzstan. A common theme across the dissertation is that shocks have distributional consequences that depend on pre-existing constraints, including credit constraints, liquidity constraints, gendered institutions, and social norms.

The first chapter examines whether post-*Dobbs* abortion restrictions affected gender gaps in mortgage credit. Using HMDA data for 2018–2024 and a borrower-level triple-difference design, I find that abortion bans increased the female–male mortgage denial gap by approximately 0.9 percentage points. The results are consistent with a tightening at the approval margin: female applicants become more likely to be denied relative to male applicants, originations decline, and interest rates among originated loans fall modestly, suggesting selection among approved borrowers.

The second chapter studies whether financial resources change migration responses to environmental stress among poor households in Indonesia. Combining the Indonesian Family Life Survey with climate data, natural-disaster information, and variation in cash-transfer programs, the chapter shows that financial resources can both reduce and increase migration. Cash transfers reduce migration after short-term rainfall shocks and natural disasters, consistent with adaptation in place, but increase migration after persistent rainfall shocks and temperature anomalies, consistent with the relaxation of liquidity constraints.

The third chapter studies whether climate shocks affect harmful marriage practices and related social norms. Using the Life in Kyrgyzstan Survey and monthly weather data, the chapter shows that dry winters increase the share of observed marriages formed through bride kidnapping and shift attitudes in favor of bride-kidnapping marriages relative to love marriages. The attitudinal effects are stronger among lower-income house-

holds, suggesting that economic vulnerability can amplify the effect of climate shocks on social norms.

Overall, the dissertation shows that policy and environmental shocks do not affect all individuals equally. Their effects depend on the constraints people face and on the institutions through which economic and social decisions are made. By studying credit access, migration, and marriage practices, the dissertation contributes to understanding how shocks translate into inequality, vulnerability, and changes in economic behavior.

Acknowledgments

First, I would like to thank my supervisor, Andreas Menzel, for his guidance, unwavering support, and, most importantly, his kindness throughout my Ph.D. studies. I feel incredibly fortunate to have had such an exceptional person and professional as my mentor. Thank you from the bottom of my heart.

I am grateful to my referees, Marc Witte and Jonathan Conning, for their time, effort, and valuable feedback.

I would also like to express my deepest gratitude to Sebastian Ottinger, Christian Ochsner, Vasily Korovkin, Mariola Pytlikova, and Roman Hoffmann for their comments, discussions, and words of encouragement throughout the years.

To Randall K. Filer and Štěpán Jurajda, thank you so much for all the help you so generously provided at the moments when I needed it most.

I am deeply grateful to Jan Švejnar for giving me the opportunity of a lifetime to visit Columbia University in New York for my mobility. It was an absolutely incredible experience, both academically and personally, and one I will always remember.

I am also sincerely grateful to the CERGE-EI Academic Skills Center, particularly Deborah Nováková, Gray Krueger, and Andrea Downing, for their careful editing of my work and for their valuable feedback over the years. I am equally grateful to the Student Affairs Office for all their help and support throughout my Ph.D. studies.

I gratefully acknowledge financial support from the Czech Academy of Sciences under GAČR project No. 20-31615S and from Charles University under GAUK project No. 584120.

I thank my friends and family for their love and belief in me. I love you all dearly. Thank you for being my support system.

To my son, Timur, thank you for changing my life and for giving this journey an entirely new meaning and purpose.

And now to the two most important people.

First, my partner and co-author, Ivan, without whom none of this would have been possible. Thank you for all your love and support, for your help, and for all the sacrifices you made to give me the space and time I needed to work on this dissertation. I share this achievement—quite literally—with you. Thank you.

And last, but certainly not least, my mother, Galina, who, throughout the years of my Ph.D., has been battling stage IV cancer. Mama, you are my hero and my greatest inspiration. There were times when I wanted to give up because everything felt too difficult, but I never did, because you never gave up. Your resilience, strength, love of life, and boundless kindness have inspired me and kept me going. I love you more than words can describe.

This is for you. Thank you.

Prague, Czech Republic Aizhamal Rakhmetova

Introduction

This dissertation consists of three essays in applied economics. The chapters study how policy and environmental shocks shape economic and social outcomes under constraint. The first chapter examines whether abortion restrictions affected gender gaps in mortgage credit in the United States. The second chapter studies how access to financial resources changes migration responses to environmental shocks among poor households in Indonesia. The third chapter examines how climate shocks affect bride kidnapping and attitudes toward bride kidnapping in Kyrgyzstan. Together, the chapters show that shocks do not operate in isolation: their consequences depend on credit constraints, liquidity constraints, social norms, gendered institutions, and the margins through which individuals and households are able to adjust.

The dissertation is motivated by a broad question in applied economics: how do people and institutions respond when external shocks change the economic environment? The answer often depends on constraints. When people have access to financial resources, they may be able to migrate, adapt, or smooth consumption. When they face liquidity constraints, they may remain exposed to environmental stress even when moving would be beneficial. When women face gendered institutions or discriminatory norms, shocks may worsen their access to credit, bargaining power, or protection from harmful practices. In this sense, the chapters are linked by a common focus on vulnerability. They examine settings in which shocks interact with pre-existing constraints and produce unequal

outcomes.

The first chapter, "Abortion Restrictions and Gender Gaps in Mortgage Credit," studies whether abortion restrictions affected gender disparities in mortgage credit after the U.S. Supreme Court's decision in *Dobbs v. Jackson Women's Health Organization*. The *Dobbs* decision generated abrupt and uneven changes in access to reproductive care across U.S. states. Some states implemented near-total abortion bans, while others maintained or strengthened legal access. This policy divergence provides a setting to study whether restrictions on reproductive autonomy affect women's access to mortgage credit relative to men.

Using Home Mortgage Disclosure Act data for 2018–2024, I estimate a borrower-level triple-difference model comparing female relative to male applicants in abortion-ban and abortion-protected states before and after policy adoption. The chapter examines three mortgage outcomes: denial, origination, and interest rates. Denials capture the lender approval margin most directly; originations measure realized access to credit; and interest rates capture pricing conditional on approval. The main empirical design allows for staggered treatment timing, with Texas treated earlier than other abortion-ban states.

The chapter finds that abortion bans widened the female–male gap in mortgage denials. Female applicants in abortion-ban states experienced an increase in denial probabilities of roughly 0.9 percentage points relative to comparable male applicants after abortion restrictions became binding. Originations declined for female applicants relative to male applicants, while borrower-level interest rates among originated loans fell modestly. This pattern is consistent with selection at the approval margin: marginal female applicants are more likely to be denied, leaving a more positively selected pool of approved female borrowers. The denial effect is concentrated among low-income applicants, is not explained by observable changes in applicant composition, and remains robust when excluding Texas, the early-treated state. Staggered synthetic difference-in-differences estimates based on state-year gender gaps provide additional robustness evidence for the denial and origination-rate results.

This chapter contributes to the literature on reproductive policy and women's economic outcomes by showing that abortion restrictions affect not only fertility, health, and labor-market outcomes, but also access to credit. It also contributes to the literature on mortgage credit and discrimination by studying how a major reproductive-policy shock changed gender gaps in approval and pricing. The results suggest that restric-

tions on reproductive autonomy can translate into reduced access to long-term financial opportunities.

The second chapter, "Access to Financial Resources and Environmental Migration of the Poor," co-authored with Mariola Pytlikova (Aarhus University, CERGE-EI) and Roman Hoffmann (IIASA), studies how poor households use migration as an adaptation strategy in response to environmental stress. Climate shocks can affect migration in two opposing ways. On the one hand, adverse environmental conditions may push people to leave affected areas. On the other hand, migration is costly, and poor households may be unable to move when shocks reduce income or destroy assets. Access to financial resources can therefore either facilitate migration, by relaxing liquidity constraints, or reduce it by helping households adapt in place.

The chapter studies this question using micro-data from Indonesia. It combines the Indonesian Family Life Survey with climate data, information on natural disasters, and variation in government cash-transfer programs targeted at poor households. The Indonesian setting is well suited for the analysis because the country is highly exposed to rainfall shocks, temperature variability, and natural disasters, while also having implemented several cash-transfer programs that provide variation in access to financial resources.

The results show that the role of financial resources depends on the type and persistence of the environmental shock. Households receiving cash transfers migrate less in response to current-year rainfall shocks and natural disasters, suggesting that additional resources can help households absorb short-term shocks and adapt in place. At the same time, households with cash transfers migrate more in response to accumulated rainfall shocks and temperature anomalies, suggesting that financial resources can also enable migration when environmental stress is persistent. Thus, the same form of financial support can operate through both a climate-inhibitor mechanism and a climate-driver mechanism.

This chapter contributes to the literature on climate migration by showing that migration responses cannot be understood without accounting for financial constraints. It reconciles two mechanisms often treated separately in the literature: migration as an adaptation strategy and resources as a tool for in-situ adaptation. The findings imply that social protection programs can alter the set of feasible adaptation choices available to poor households, but their effect depends on the nature of the environmental shock.

The third chapter, "Weather Shocks and Bride Kidnapping in Kyrgyzstan," studies how climate shocks affect a harmful marriage practice and attitudes toward it. Bride kidnapping, or marriage by capture, is a form of forced marriage in which a man abducts the woman he wishes to marry. The practice persists in Kyrgyzstan despite legal prohibition and has serious consequences for women's autonomy, safety, and well-being. The chapter asks whether environmental stress changes the composition of observed marriages toward bride kidnapping and shifts social attitudes toward the practice.

The chapter combines monthly weather data with the Life in Kyrgyzstan Survey, a nationally representative panel survey that contains information on marriage formation and attitudes toward bride kidnapping. The empirical strategy exploits variation in dry winters, defined as winters with insufficient precipitation relative to historical norms. In Kyrgyzstan, winter precipitation is important for agriculture because snow accumulation affects water availability, crop conditions, and rural livelihoods. Dry winters therefore provide a plausibly exogenous source of economic stress.

The results show that insufficient winter precipitation increases the share of observed marriages formed through bride kidnapping rather than through love or arranged marriage and affects attitudes toward the practice. Districts exposed to dry winters experience a higher probability of marriages formed through bride kidnapping. The chapter also finds that dry winters make respondents more accepting of bride kidnapping by reducing the perceived negative consequences of marriages that begin this way. The attitudinal effects are stronger among individuals from lower-income households, suggesting that economic vulnerability amplifies the effect of climate shocks on social norms. The presence of daughters in the household is associated with less favorable attitudes toward bride kidnapping, while the analysis does not find systematic heterogeneity by gender or education.

This chapter contributes to the literature on climate shocks, gender inequality, and harmful social practices. While much of the existing literature focuses on child marriage, bride kidnapping is a distinct institution because it does not operate through a standard economic transfer such as bride price or dowry. Instead, it involves coercion and a violation of women's autonomy. By studying both actual marriage outcomes and attitudes toward bride kidnapping, the chapter shows that climate shocks can affect not only material outcomes but also the social norms that sustain harmful practices.

Taken together, the three chapters show that shocks can deepen vulnerability when

they interact with economic and social constraints. In the U.S., abortion restrictions affect women's relative access to mortgage credit, especially at the approval margin. In Indonesia, environmental shocks affect migration differently depending on whether poor households have access to financial resources. Dry winters in Kyrgyzstan affect both the composition of observed marriages and attitudes toward bride kidnapping. Across these settings, the dissertation shows that applied microeconomic analysis can reveal how institutions, resources, and social norms mediate the consequences of shocks. The common message is that policy and environmental shocks have distributional consequences that are shaped by pre-existing constraints, and that these constraints are central for understanding economic behavior and inequality.

Chapter 1

Abortion Restrictions and Gender Gaps in Mortgage Credit

1.1 Introduction

Restrictions on women's reproductive autonomy are not only health or legal issues; they are also shocks to gender equality and economic opportunity. A central claim of the women's rights movement has been that control over fertility is necessary for women's full participation in education, work, family formation, and financial decision-making. When access to abortion is restricted, women may face greater uncertainty over fertility timing, household responsibilities, mobility, and future economic stability (Goldin & Katz, 2002; Myers, 2017; Miller et al., 2023; Londoño-Vélez & Saravia, 2025). These changes can affect not only whether and when women have children, but also their ability to participate on equal terms in markets that shape long-run economic security.

This paper studies one such market: mortgage credit. Mortgage access is a key channel through which individuals build wealth, choose where to live, and make long-term family and labor-market decisions (Munnell et al., 1996; Ladd, 1998). If reductions in reproductive rights affect women's perceived or actual financial risk, they may also reduce women's access to credit and widen gender gaps in economic opportunity. Yet little is known about whether restrictions on women's reproductive rights spill over into

formal credit markets. This paper asks whether abortion restrictions in the United States widened gender gaps in mortgage approval, origination[1], and pricing.

I study this question in the context of the U.S. Supreme Court's 2022 decision in *Dobbs v. Jackson Women's Health Organization* (hereafter *Dobbs*). The decision overturned *Roe v. Wade* and returned abortion regulation to the states, generating sharp and uneven changes in access to reproductive care across the United States (U.S. Supreme Court, 2022). In some states, abortion became effectively unavailable within a short period; in others, access remained broadly protected (Guttmacher Institute, 2026). This policy divergence created a setting in which women in different states faced substantially different reproductive-policy environments after 2022.

The economic relevance of this shock is that abortion restrictions can change both current and expected household resources. An unwanted or mistimed birth can affect income, employment, medical and childcare expenses, geographic mobility, relationship stability, and savings accumulation (Miller et al., 2023; Londoño-Vélez & Saravia, 2025; Kleven et al., 2019). These channels are directly relevant for mortgage markets, where approval and pricing depend on expected resources, debt burden, collateral, and repayment risk (Munnell et al., 1996; Ladd, 1998). Mortgage applications therefore provide a setting in which changes in reproductive constraints and uncertainty about future childbearing can translate into economically consequential credit-market outcomes: restrictions may alter households' financial circumstances and their demand for mortgage credit, while lenders may respond to changes in applicants' observed or perceived repayment risk. Therefore, a reduction in reproductive autonomy may affect mortgage outcomes through both borrower-side and lender-side margins.

A broader channel may also operate through geographic mobility and local housing markets. If abortion restrictions induce some households to relocate from ban states to protected states, the resulting migration may affect housing demand and mortgage-market conditions in both origin and destination states. Such general-equilibrium responses are not separately identified in the empirical design and may therefore form part of the reduced-form effect of the post-Dobbs policy environment.

Using Home Mortgage Disclosure Act (HMDA) data for 2018–2024, I examine three

[1]Originations refer to applications that result in an issued mortgage. This outcome captures realized mortgage access, but it combines lender approval and borrower take-up, so it is less directly tied to lender screening than denial decisions.

outcomes that capture different margins of the mortgage market: denials, originations, and interest rates. Denials measure whether a lender rejects a mortgage application and provide the closest observed proxy for credit approval decisions. Originations indicate whether an application actually results in an issued loan, combining both lender approval and borrower take-up. Interest rates are analyzed only among originated loans and capture the pricing of credit conditional on approval. Distinguishing between these three outcomes allows me to assess whether abortion bans affected (1) access to credit at the approval margin, (2) realized mortgage borrowing, or (3) the terms offered to borrowers who remain in the approved pool.

The main empirical design is a borrower-level triple-difference framework that exploits the staggered post-*Dobbs* policy landscape, including Texas as an earlier treated state. The primary estimates use application-level HMDA data and compare changes in outcomes for female relative to male applicants in abortion-ban states before and after policy adoption. As a robustness check, I also estimate synthetic difference-in-differences (SDID) models based on state-year female–male gaps in mortgage outcomes. This exercise uses the same staggered timing logic as the main design, reweights protected states to better match the pre-treatment gender-gap trajectories of treated states, and provides state-level robustness evidence for the main borrower-level results.

Across specifications, I find that the clearest and most robust result is that abortion bans widened the gender gap in mortgage denials. In the borrower-level triple-difference analysis, the results show that female applicants in abortion-ban states experienced a significant increase in denial rates relative to comparable male applicants after the policy change. The estimated effect is approximately 0.9 percentage points.

The borrower-level estimates also show that originations decline for female applicants relative to male applicants, and the staggered synthetic difference-in-differences estimates based on origination rates align with this result. Interest-rate effects are less robust: the borrower-level estimate is negative and statistically significant but economically small, while the state-year SDID estimate remains negative but is not statistically significant. Because interest rates are analyzed only among originated loans, this conditional pricing outcome is subject to selection into approval. I therefore treat the interest-rate estimate as a secondary result and do not use it to identify the underlying mechanism. Taken together, the strongest evidence indicates that abortion bans reduced women's access to mortgage credit at the approval margin, as reflected in higher denial rates and lower

originations.

Evidence concerning the mechanism provides several additional patterns. The denial effect is substantially larger among low-income applicants, who are more likely to be close to lenders' approval thresholds. The main denial result also remains robust to the observable compositional changes identified in the data and to excluding Texas, the early-treated state. Finer age-group estimates do not show a monotonic decline in the effect with applicant age, weighing against a narrow mechanism based solely on an individual applicant's fertility risk. Finally, the lender analysis shows that a substantial component of the effect remains within lenders. Point estimates are larger in Republican-HQ lender environments, although the formal difference across lender environments is not statistically significant.

Taken together, these patterns suggest that the response extends beyond a narrow channel based solely on an individual applicant's fertility risk. The denial effect is not concentrated only among applicants in prime reproductive ages, and the within-lender response varies across lender institutional environments. These patterns are consistent with several possible mechanisms, including changes in lender screening or perceived repayment risk. However, because HMDA does not observe credit scores or all other dimensions of borrower risk, the analysis cannot distinguish changes in gender-related screening from risk-based underwriting responses to unobserved changes in creditworthiness. I therefore treat the mechanism evidence as suggestive rather than as uniquely identifying gender salience or discriminatory screening.

My paper contributes first to a large literature showing that reproductive control affects women's economic trajectories. Classic work emphasizes how access to contraception and abortion changed fertility timing, marriage, and human-capital investment. Goldin & Katz (2002) showed that early legal access to oral contraception helped women delay marriage and invest in professional careers. Angrist & Evans (1996) used early U.S. abortion reforms to estimate the consequences of teen and out-of-wedlock childbearing for later schooling and labor-market outcomes. Myers (2017) reexamined the role of abortion access in allowing young women to delay marriage and motherhood. Ananat et al. (2007) showed that abortion legalization in the 1970s led to a permanent reduction in fertility. More recent work studies modern restrictions. Jones & Pineda-Torres (2024) find that Targeted Regulation of Abortion Providers (TRAP) laws affect fertility and educational attainment among teens, while Dench et al. (2024) provide early evidence that post-*Dobbs*

abortion bans increased births in affected states. Taken together, this strand of research shows that reproductive control affects women's fertility timing, human-capital investment, labor-market outcomes, and broader economic trajectories. My paper extends this body of work by studying whether abortion restrictions also affect women's access to mortgage credit, a downstream economic outcome that has received much less attention.

A second strand of literature documents that abortion access has consequences far beyond fertility itself. Miller et al. (2023) link the Turnaway Study to administrative credit records and show that being denied a wanted abortion substantially increases financial distress and adverse credit events. Using administrative microdata from Colombia and exogenous judicial variation, Londoño-Vélez & Saravia (2025) show that being denied a wanted abortion worsens women's mortality, health, education, labor supply, poverty, and the well-being of their existing children. Other work traces related channels through health and household conditions. Clarke & Mühlrad (2021) show that legal access to abortion in Mexico decreased maternal morbidity. Farin et al. (2024) show that legal abortion reduced maternal mortality in the United States, especially for nonwhite women. Gardner (2024) finds that restricted abortion access worsens maternal and infant health outcomes in the United States. Durrance et al. (2024) show that mandatory waiting periods increase intimate partner violence. Emerging post-*Dobbs* evidence also points to mental-health consequences of the abortion ban, although some estimates remain imprecise; see, for example, Baser et al. (2025). Relative to this literature, my paper shifts attention from fertility, health, and labor-market outcomes to household finance and the allocation of mortgage credit.

My study also relates to the literature on credit allocation and mortgage discrimination. Research using HMDA and related data has long documented disparities in mortgage approval and pricing and has debated how much of those gaps can be explained by omitted credit-risk variables. Munnell et al. (1996) used detailed Boston data to argue that racial disparities in mortgage denials remain even after controlling for richer underwriting information, while Ladd (1998) surveyed the broader mortgage-discrimination literature and found evidence of discrimination against minorities. Although this literature focuses primarily on race, a smaller body of work studies gender differences in mortgage outcomes. Cheng et al. (2011) show that women pay higher mortgage rates than men on average. Fang & Munneke (2020) find that female borrowers pay higher contract rates in the subprime mortgage market even after accounting for default and pre-

payment risk. More recent evidence suggests that gender bias may also be embedded in credit-scoring systems themselves; see Liu & Liang (2025). I contribute to the mortgage-discrimination literature by asking whether a major reproductive-policy shock widened gender gaps in mortgage approval and pricing. I find that the strongest evidence points to a widening of the female–male denial gap, suggesting that the consequences of abortion bans extend into credit allocation. More broadly, the paper connects to the literature on discrimination and credit allocation by showing how a major policy and institutional shock can change group disparities in credit-market outcomes, while the available data do not allow me to uniquely identify the mechanism generating those changes.

Finally, my analysis is consistent with broader family-economics literature showing that childbearing generates persistent labor-market penalties for women. Adda et al. (2017) estimate large long-run career costs of children; Lundborg et al. (2017) show that fertility affects women's careers using IVF-based quasi-experimental variation; and Kleven et al. (2019) document large child penalties in earnings and hours worked. Thus, if abortion bans increase unwanted or mistimed births, mortgage-market effects are a natural downstream manifestation of the same underlying shock. Furthermore, the age and lender-heterogeneity results of my analysis suggest that the effect of abortion bans is not limited to individual fertility risk. In that sense, my paper brings the economics of abortion policy into direct conversation with household finance, housing markets, and gender inequality, and shows that *Dobbs* may have affected not only whether and when households have children, but also women's ability to translate economic plans into access to mortgage credit.

1.2 Institutional Background and Policy Variation

The legal landscape of abortion access in the United States changed abruptly following the Supreme Court's decision in *Dobbs v. Jackson Women's Health Organization* in June 2022 (U.S. Supreme Court, 2022). The ruling overturned *Roe v. Wade* and eliminated the federal constitutional protection for abortion access, returning authority over abortion regulation to individual states. As a result, states could implement pre-existing laws designed to take effect if *Roe* were overturned, enforce previously blocked restrictions, or enact new legislation that substantially limited or prohibited abortion access.

The discussions around the *Dobbs* case started in 2020. The petition asking the

Supreme Court to review the case was filed by Mississippi on June 15, 2020, the Supreme Court agreed to hear the case on May 17, 2021, and oral arguments took place on December 1, 2021 (Supreme Court of the United States, 2022; Oyez, 2022). The issue became even more visible after a draft opinion was leaked on May 2, 2022, showing that the Court was likely to overturn *Roe v. Wade*. The final decision was issued on June 24, 2022 (Howe, 2022; U.S. Supreme Court, 2022).

The policy changes that followed the final decision were both rapid and highly heterogeneous across states. A subset of states implemented near-total bans shortly after the *Dobbs* decision, often within weeks or months. Other states maintained legal access to abortion and, in some cases, adopted additional statutory or constitutional protections. This divergence created a sharp state-level policy contrast. In some states, abortion became effectively unavailable or severely restricted; in others, access remained broadly protected (Guttmacher Institute, 2026; Dench et al., 2024). This post-*Dobbs* divergence generates the policy variation used in my empirical analysis to classify states into abortion-ban, abortion-protected, and excluded categories. The detailed state classification, treatment timing, and construction of the state-year treatment indicator are described in Section 1.3.8. In brief, the main analysis compares clearly restrictive states to clearly protective states and excludes states with partial, delayed, temporary, or legally uncertain policy regimes.

The economic importance of this legal change extends beyond fertility itself. Abortion access affects households' ability to control the timing of births, and fertility timing is closely related to education, labor supply, earnings, savings, household formation, and mobility. These channels are directly relevant for mortgage credit. A mortgage is a long-term financial contract, and both the decision to apply and the lender's decision to approve depend on expected income, future expenses, household stability, and repayment risk. If abortion restrictions increase the risk of unwanted or mistimed births, raise expected childcare and medical costs, or reduce labor-market flexibility, those restrictions may affect households' ability to qualify for mortgage credit.

The institutional setting therefore motivates a borrower-level triple-difference design, or equivalently a gender-differential difference-in-differences design. The policy shock varies by state and year, while the hypothesized effects are gendered. I compare female applicants to male applicants in abortion-ban and abortion-protected states before and after abortion restrictions become binding. The key empirical question is whether female

applicants in abortion-ban states experience a differential change in mortgage outcomes relative to male applicants after the policy change.

The mechanism is developed more fully in Section 1.4, where I distinguish borrower-side responses, lender screening, and selection into approval.

1.3 Data

1.3.1 HMDA Application-Level Data

The main data source is the Home Mortgage Disclosure Act (HMDA) public loan-level dataset, covering mortgage applications for 2018–2024.[2] HMDA requires covered mortgage lenders in the United States to report detailed information on mortgage applications, loan characteristics, borrower characteristics, and application outcomes. The data include both originated and denied applications, which makes HMDA particularly well suited for studying credit access and lender approval decisions.

HMDA has several advantages for this analysis. First, it provides broad coverage of the formal mortgage market, allowing me to study a large national sample rather than a selected set of lenders or local housing markets. Second, it records both successful and unsuccessful applications, which is essential because the main outcome of interest is denial. Third, it reports applicant sex, race, ethnicity, age, income, loan amount, loan-to-value ratio, debt-to-income ratio, and other underwriting-related variables. These variables allow me to compare female and male applicants and to control for observable differences in borrower and loan characteristics.

The gender indicator is based on the sex of the primary applicant reported in HMDA. Accordingly, an application is classified as female when the primary applicant is female, irrespective of whether a co-applicant is present or of the co-applicant's sex.

An important limitation of HMDA is that it does not report applicants' credit scores or complete credit histories. These variables are central inputs into mortgage underwriting and may themselves respond to changes in households' financial circumstances.

[2]HMDA public loan-level data are retrieved from the Federal Financial Institutions Examination Council / Consumer Financial Protection Bureau HMDA data portal: HMDA data browser.

Consequently, the analysis can condition on a rich set of observable borrower and loan characteristics, but cannot hold constant all dimensions of borrower credit risk.

The unit of observation in the borrower-level analysis is a mortgage application. I use the property state to assign each application to a state-year abortion-policy environment. The main sample includes applications in states classified as either abortion-ban states or abortion-protected states. States with partial, delayed, temporary, or legally uncertain abortion-policy regimes are excluded from the baseline analysis to maintain a clean treatment-control contrast.

1.3.2 Sample Construction

I impose several restrictions to create the estimation sample. First, I restrict attention to home-purchase mortgage applications for owner-occupied properties and first-lien loans. This restriction focuses the analysis on the primary residential home-purchase mortgage market and avoids mixing purchase applications with refinancing, investor, or second-lien borrowing. Second, for the mortgage denial and origination outcomes, I retain applications that receive a definitive lender decision. This "decisioned" sample includes applications that are originated, denied, or approved by the lender but not accepted by the applicant. It excludes applications that are withdrawn, closed for incompleteness, purchased by another institution, or otherwise not directly comparable as lender approval decisions. Because approved-but-not-accepted applications are included in the decisioned sample but are classified as neither originated nor denied, the denial and origination rates do not sum to one. Third, I exclude observations with missing or implausible values for key borrower and loan controls, including income, loan amount, loan-to-value categories, and debt-to-income categories.

For the interest-rate outcome, I impose an additional restriction. I analyze interest rates only among originated loans, because the pricing outcome is intended to measure loan terms conditional on completed origination. I therefore estimate interest-rate specifications on the subset of originated loans with non-missing and valid reported interest rates. This distinction is important for interpretation because denial and origination outcomes are estimated on the decisioned application sample, while interest-rate outcomes are conditional on origination.

The resulting baseline sample contains approximately 19 million mortgage applica-

tions from 2018–2024. The large sample size is useful for estimating gender-specific changes in mortgage outcomes, but it also implies that very small differences can be statistically precise. For this reason, throughout the analysis I interpret the results in terms of both statistical significance and economic magnitude.

1.3.3 Outcome Variables

I study three mortgage outcomes: denial, origination, and interest rates. These outcomes capture distinct margins of the mortgage market.

The first outcome is an indicator for whether the mortgage application is denied. Denial is the primary outcome because it is the closest observed measure of lender rejection in HMDA. If abortion bans affect lenders' willingness to approve female applicants relative to male applicants, this should appear most directly in denial rates.

The second outcome is an indicator for whether the loan is originated. Origination measures realized access to mortgage credit, but it combines several margins. A loan can fail to originate because the lender denies the application, because the borrower does not proceed, or because the application does not otherwise result in a completed loan. For this reason, I interpret originations as a broader measure of realized mortgage access, while denial status is the cleaner measure of lender approval.

The third outcome is the interest rate on originated loans. Interest rates capture the pricing of credit conditional on approval. This outcome is informative, but it must be interpreted carefully because it is observed only among approved and originated loans.

These three outcomes are interpreted jointly in the empirical analysis. Denials provide the closest measure of the approval margin, originations capture realized mortgage access, and interest rates capture pricing conditional on origination.

1.3.4 Borrower and Loan Controls

The borrower-level specifications control for a set of applicant and loan characteristics observed in HMDA. The baseline controls include log applicant income, log loan amount, an indicator for conforming loan status, loan-to-value categories, and debt-to-income categories. These variables capture important dimensions of mortgage underwriting: borrower resources, loan size, collateralization, and repayment burden.

In richer specifications, I also include applicant race, ethnicity, and age categories. These controls are important because female and male applicants differ systematically in observable characteristics, and because demographic composition may vary across states and over time. Including these controls reduces the concern that the estimated gender-specific treatment effect reflects changing applicant composition rather than a change in credit access.

Several variables enter as categorical bins rather than continuous measures. HMDA reports some underwriting variables, such as loan-to-value and debt-to-income ratios, in categories or with special reporting values. I therefore use bin indicators rather than imposing a linear relationship between these variables and mortgage outcomes. In descriptive tables, I also construct approximate numeric versions of LTV and DTI bins to report average values in an interpretable way. Table 1.1 summarizes the main outcome variables, treatment-related applicant characteristics, and control variables used in the empirical analysis.

Table 1.1: Main Variables Used in the Analysis

Variable	Definition	Sample
Denied	Indicator equal to one if the application is denied by the lender	Decisioned applications
Originated	Indicator equal to one if the application results in an originated loan	Decisioned applications
Interest rate	Reported interest rate on originated loan	Originated loans only
Female	Indicator equal to one for female applicant	Application level
Log income	Log of applicant income	Application level
Log loan amount	Log of requested loan amount	Application level
Conforming loan	Indicator for conforming loan status	Application level
LTV bins	Loan-to-value categories	Application level
DTI bins	Debt-to-income categories	Application level
Race, ethnicity, age	Applicant demographic categories	Application level

Notes: LTV and DTI enter the regressions as categorical controls rather than continuous variables. Interest-rate specifications are estimated only among originated loans, so interest-rate results should be interpreted as conditional on approval and origination.

1.3.5 Summary Statistics

Table 1.2 reports summary statistics for the main variables used in the analysis. The sample consists of approximately 19 million mortgage applications from HMDA over the 2018–2024 period after applying the baseline restrictions described above.

Table 1.2: Summary Statistics

	Full sample		Male		Female	
	mean	sd	mean	sd	mean	sd
Denied	0.125	0.331	0.115	0.318	0.144	0.351
Originated	0.841	0.365	0.852	0.355	0.822	0.383
Interest rate	4.598	1.569	4.568	1.553	4.655	1.597
Log income	4.534	0.673	4.599	0.669	4.419	0.666
Log loan amount	12.475	0.720	12.539	0.713	12.364	0.720
Conforming loan	0.945	0.228	0.933	0.251	0.966	0.180
Loan-to-value ratio	89.746	9.048	89.608	8.992	89.994	9.142
Debt-to-income ratio	38.263	10.521	37.878	10.570	38.939	10.400
Female	0.363	0.481				
Black	0.083	0.276	0.064	0.244	0.118	0.322
Asian	0.079	0.269	0.085	0.279	0.068	0.251
Hispanic	0.132	0.339	0.135	0.341	0.128	0.334
Age 18–34	0.398	0.490	0.397	0.489	0.400	0.490
Observations	18954977		12068301		6886676	

Notes: This table reports summary statistics for the main variables used in the analysis. Loan-to-value and debt-to-income ratios are reported as averages. Race, ethnicity, and age variables are reported as shares of applicants in each category. Denial and origination rates do not sum to one because the decisioned sample also includes applications that were approved by the lender but not accepted by the applicant. The observation count reported at the bottom of the table refers to the full decisioned application sample; variable-specific sample sizes may be smaller because of missing values or outcome-specific restrictions. Interest-rate statistics are calculated only among originated loans with a valid reported interest rate.

There are several patterns in the data. First, there are clear gender differences in mortgage outcomes. Female applicants exhibit higher denial rates than male applicants and correspondingly lower origination rates. These differences indicate substantial baseline disparities in credit access. Interest rates among originated loans are similar across groups, with female applicants facing slightly higher average rates. Formal tests of the three main outcome differences are reported in Table 1.3.

Second, female applicants differ systematically in observable financial characteristics. On average, they report lower income and smaller loan amounts than male applicants. In addition, they exhibit slightly higher leverage, with higher average loan-to-value ratios

Table 1.3: Female–Male Differences in Main Mortgage Outcomes

Outcome	Male mean	Female mean	Female–Male difference	Observations
Denied	0.115	0.144	0.030***	18954976
			(0.005)	
Originated	0.852	0.822	-0.031***	18954976
			(0.005)	
Interest rate	4.568	4.655	0.087***	15919233
			(0.010)	

Notes: Denial and origination outcomes are calculated on the decisioned sample. Interest rates are calculated only for originated loans. Female–male differences are estimated from regressions of each outcome on a female indicator. Standard errors, reported in parentheses, are clustered at the state level. $* < 0.10$, $** < 0.05$, $*** < 0.01$.

and higher debt-to-income ratios. These patterns are consistent with female borrowers being closer to credit constraints and therefore potentially more sensitive to changes in underwriting standards.

Third, the descriptive statistics clarify why the analysis focuses on gender gaps rather than only on female outcomes in levels. Female and male applicants differ at baseline in both outcomes and observable characteristics. The empirical strategy therefore asks whether these gender gaps change differentially in abortion-ban states after abortion restrictions become binding, rather than whether female applicants have different average outcomes in levels.

1.3.6 Lender Identifiers and Transmittal Sheet Data

For the lender heterogeneity analysis, I supplement the HMDA loan-level data with the HMDA Transmittal Sheet data.[3] The Transmittal Sheet contains institution-level information reported by filing institutions, including the Legal Entity Identifier (LEI) and the state associated with the filing institution's office. I merge the Transmittal Sheet data to the loan-level HMDA records using the lender LEI and activity year.

The LEI is the key lender identifier in HMDA. In the application-level data, each mortgage application is associated with a reporting institution through its LEI. I use this identifier to attach lender-year characteristics from the Transmittal Sheet. Before merging, I clean the Transmittal Sheet files separately by year, retain the LEI and relevant

[3]HMDA Transmittal Sheet data are retrieved from the FFIEC/CFPB HMDA data publication portal: `https://ffiec.cfpb.gov/data-publication/snapshot-national-loan-level-dataset/2024`.

institution-location fields, and remove duplicate LEI-year observations. I then append the annual Transmittal Sheet files for 2018–2024 and merge them into the HMDA application-level data by LEI and activity year.

The main lender-level variable used in the heterogeneity analysis is the state associated with the filing institution's office. I refer to this variable as the lender headquarters (HQ) or reporting-office state for simplicity. In the empirical analysis, I use this state as a proxy for the conservative or non-conservative institutional environment of the lender. The idea is simple: lenders associated with more conservative headquarters environments may operate under different organizational cultures, market expectations, or assumptions about borrowers. For example, lenders may differ in how much they rely on local loan officers, standardized underwriting systems, or relationship-based lending. These differences could lead to larger negative effects for women in general after *Dobbs*.

This measure is necessarily imperfect. The HMDA Transmittal Sheet reports the state associated with the filing institution's office, which may not always coincide with the ultimate parent company's headquarters or the location of senior management. It also does not observe the beliefs of executives, loan officers, or internal underwriting models. I therefore interpret the Republican-HQ indicator as a proxy for conservative lender culture or institutional environment, not as a direct measure of ideology or intentional discrimination. This type of measurement error is common in empirical work using institutional proxies; the relevant question is whether the gender-specific response to abortion bans is stronger among lenders associated with more conservative environments.

1.3.7 Construction of State-Year Gender Gaps

In addition to the borrower-level dataset, I construct a state-year panel of gender gaps in mortgage outcomes. This panel is used for the synthetic difference-in-differences robustness checks and for descriptive figures showing trends in female–male gaps.

For each state and year, I calculate average outcomes separately for female and male applicants. I then define the gender gap as the female mean minus the male mean. For example, the denial gap is

$$GapDenial_{st} = E[Denied_{ist} \mid \text{Female}_i = 1, s, t] - E[Denied_{ist} \mid Female_i = 0, s, t].$$

Table 1.4: Construction of Lender Political-Environment Variable

Step	Description
HMDA loan-level data	Retain application-level records with lender LEI and activity year.
HMDA Transmittal Sheet	Import annual Transmittal Sheet files and retain LEI, activity year, and filing-institution state.
Merge	Merge Transmittal Sheet information to HMDA applications by LEI and activity year.
State political environment	Classify filing-institution states by Republican or non-Republican political environment.
Lender heterogeneity variable	Define Republican-HQ/reporting-office indicator used in lender heterogeneity specifications.

Notes: The HMDA Transmittal Sheet reports the state associated with the filing institution's office. I use this as a proxy for the lender's headquarters or reporting-office environment and, in the heterogeneity analysis, for the conservative or non-conservative institutional culture of the lender. This variable does not necessarily identify the ultimate corporate headquarters or the political preferences of lender executives.

Analogously, I construct gender gaps in originations and interest rates. Denial and origination-rate gaps are calculated using the decisioned application sample. Interest-rate gaps are calculated using originated loans only, because interest rates are not observed for denied applications.

The state-year gap panel has two advantages. First, it directly corresponds to the level at which abortion policy varies. Second, it provides a transparent way to summarize the evolution of gender disparities across treated and control states. However, the state-year panel is much smaller than the borrower-level dataset and therefore has less statistical power. For this reason, I treat the borrower-level triple-difference estimates as the primary results and use the state-year gap analysis as supporting descriptive and robustness evidence.

1.3.8 Abortion Policy Classification and Treatment Timing

I construct a state-year abortion-policy dataset that classifies U.S. states according to their post-*Dobbs* abortion-policy regime and the timing of policy implementation. This classification is central to the empirical design because the treatment variable is defined at the state-year level and is then assigned to individual HMDA mortgage applications

using the property state and application year. The goal is to create a transparent and empirically usable measure of exposure to binding abortion restrictions, while avoiding ambiguous cases in which treatment status is difficult to define.

Following Dench et al. (2024), I classify states into three groups. First, *ban states* are states that implemented comprehensive abortion restrictions or near-total bans during the post-*Dobbs* period. These states form the treated group. Second, *protected states* are states that maintained legal access to abortion and, in many cases, enacted or preserved additional protections after *Dobbs*. These states form the main comparison group. Third, I define a set of *excluded states* that experienced partial, delayed, temporary, or legally uncertain policy changes during the sample period. These states are omitted from the main analysis.

The distinction between ban, protected, and excluded states is important for identification. The main empirical strategy compares clearly restrictive states to clearly protective states. States with partial or legally contested restrictions may not provide the same treatment intensity as states with binding near-total bans. At the same time, they may not be comparable to protected states if legal uncertainty affects providers, households, or expectations. Including such states could introduce measurement error into the treatment variable and weaken the contrast between treated and control groups. For this reason, the baseline sample prioritizes a sharper comparison between ban states and protected states.

The classification follows recent empirical work on post-*Dobbs* abortion bans (Dench et al., 2024) and is consistent with state-level abortion-policy trackers that document the evolution of state regimes after *Dobbs* (Guttmacher Institute, 2026). The purpose of the classification is not to capture every legal detail of each state's abortion law, but rather to construct a clean treatment indicator for empirical analysis. In particular, the main specification focuses on states with clearly defined policy regimes and excludes states where enforcement was temporary, uncertain, intermediate, or legally contested during the relevant period.

Table 1.5 reports the classification of states into ban, protected, and excluded categories. The map of ban and protected states is reported in Appendix A.

Treatment timing is staggered across ban states. Texas is treated as an early-exposure state, while most other ban states enter treatment after the *Dobbs* decision. In the

Table 1.5: State Abortion Dispositions After *Dobbs*

(1) Total Ban	(2) Protected	(3) Excluded
Alabama	Alaska	Arizona
Arkansas	California	Florida
Idaho	Colorado	Georgia
Kentucky	Connecticut	Indiana
Louisiana	Delaware	Iowa
Mississippi	DC	Nebraska
Missouri	Hawaii	North Carolina
Oklahoma	Illinois	North Dakota
South Dakota	Kansas	Ohio
Tennessee	Maine	Pennsylvania
Texas	Maryland	South Carolina
Wisconsin	Massachusetts	Utah
West Virginia	Michigan	Wyoming
	Minnesota	
	Montana	
	Nevada	
	New Hampshire	
	New Jersey	
	New Mexico	
	New York	
	Oregon	
	Rhode Island	
	Vermont	
	Virginia	
	Washington	

Notes: Classification follows Dench et al. (2024). Ban states are treated states in the main analysis. Protected states form the comparison group. Excluded states are omitted from the main analysis because they experienced partial, delayed, temporary, or legally uncertain policy changes during the sample period.

empirical analysis, I define the post-treatment period as beginning in 2022 for Texas and in 2023 for other ban states. Protected states are never treated during the sample period.

This staggered timing is important. The year 2022 is not the same treatment period for all ban states. Texas is already exposed to binding restrictions, while other ban states are in a transition period or in their last pre-treatment year. A common post indicator for all ban states would therefore misclassify some state-years. In annual HMDA data, this distinction matters because mortgage applications are observed at the year level rather than at the exact date of the policy shock. The treatment variable must therefore reflect the year in which each state is exposed to binding abortion restrictions.

Using this classification, I construct two key variables. First, I define a state-level treatment indicator, BanState_s, equal to one for ban states and zero for protected states. Second, I define a state-year post indicator, Post_{st}, equal to one in state-years after abortion restrictions become binding and zero otherwise. The main treatment indicator is

$$\text{Treated}_{st} = \text{BanState}_s \times \text{Post}_{st}.$$

For Texas, Post_{st} begins in 2022. For other ban states, Post_{st} begins in 2023. For protected states, Treated_{st} is always zero. This construction allows the analysis to exploit both cross-state variation in abortion-policy regimes and within-state variation over time, while accommodating the staggered implementation of abortion restrictions.

Table 1.6 summarizes the timing logic used in the main empirical design.

The same timing logic is used in the event-study analysis. Rather than indexing dynamic effects only by calendar year, I define event time relative to each state's treatment year. Event time 0 corresponds to the year immediately before treatment, and event time 1 corresponds to the first treated year. Thus, for Texas, 2021 is event time 0 and 2022 is event time 1. For other ban states, 2022 is event time 0 and 2023 is event time 1. This normalization allows the event-study analysis to align treated states according to their own policy timing.

The staggered treatment definition is also used in the synthetic difference-in-differences (SDID) robustness check. Because SDID graphical diagnostics are cohort-specific under staggered implementation, the Appendix B reports separate plots for Texas and for the

Table 1.6: Treatment Timing in the Main Empirical Design

State group	Treatment timing	Role in empirical design
Texas	Post period begins in 2022	Early-treated ban state. Used to capture earlier exposure to binding restrictions.
Other ban states	Post period begins in 2023	Main treated group after the post-*Dobbs* policy divergence.
Protected states	Never treated during sample	Comparison group in the borrower-level triple-difference analysis and donor pool in SDID robustness checks.
Excluded states	Omitted from main analysis	States with partial, delayed, temporary, or legally uncertain policy changes.

Notes: The table summarizes the treatment timing used in the main analysis. The staggered timing is incorporated into the borrower-level triple-difference specification, the event-study design, the state-year randomization-inference exercise, and the staggered SDID robustness check.

remaining ban states. In the main text, the SDID results are summarized in Table 1.15 using state-year female–male gaps in mortgage outcomes.

Overall, the treatment construction reflects three empirical choices. First, I compare states with clearly restrictive abortion regimes to states with clearly protective regimes. Second, I exclude states with ambiguous, partial, delayed, or legally uncertain policy changes. Third, I allow treatment timing to differ between Texas and other ban states. These choices are intended to reduce treatment misclassification and make the estimated effects easier to interpret as exposure to binding abortion restrictions rather than as broader regional or political differences across states.

1.3.9 State-Level Controls

In the synthetic difference-in-differences analysis, the unit of observation is the state-year. Accordingly, I augment the state-year panel of gender gaps in mortgage outcomes with a set of macroeconomic controls that capture cross-state differences in labor-market conditions, housing-market dynamics, income levels, and market size. These controls are included in SDID specifications with controls and are intended to absorb time-varying state-level factors that may be correlated both with abortion-policy adoption

and mortgage-market outcomes.

The baseline state-level control vector includes four variables. First, I include the state unemployment rate, which captures local labor-market conditions and household income risk. The unemployment data are taken from the Bureau of Labor Statistics Local Area Unemployment Statistics program.[4] Second, I include a state-level house price index, measured using the Federal Housing Finance Agency House Price Index. This variable captures changes in local housing demand and collateral values, both of which are directly relevant for mortgage approval and pricing.[5] Third, I include personal income per capita from the Bureau of Economic Analysis, which captures differences in state income levels and purchasing power.[6] Fourth, I include total state population from the U.S. Census Bureau, which controls for differences in state size and broader demographic scale.[7]

Because the SDID panel is annual, all controls are harmonized to the state-year level. Variables originally available at higher frequency, such as the FHFA house price index, are aggregated to annual values before being merged with the outcome panel. In the empirical analysis, I use the unemployment rate in levels, while house prices, income per capita, and population enter in logarithms. These transformations reduce skewness and allow the variables to capture proportional differences in local economic conditions.

These state-level covariates are used only in supplementary SDID specifications. Their inclusion provides a sensitivity check for whether the results are robust to conditioning on contemporaneous differences in local labor-market, housing-market, income, and population conditions. However, some of these variables may themselves respond to the post-*Dobbs* policy environment. For example, abortion restrictions may affect migration, local housing demand, labor-market participation, or population growth. In that case, conditioning on contemporaneous house prices, unemployment, income, or population could absorb part of the policy effect rather than remove confounding. I therefore treat

[4]State unemployment rate data are retrieved from the Bureau of Labor Statistics Local Area Unemployment Statistics tables: BLS LAUS tables, and from the BLS data viewer: BLS LAUS time-series files.

[5]FHFA House Price Index data are retrieved from the Federal Housing Finance Agency: FHFA state-level HPI file.

[6]State personal income per capita data are retrieved from the Bureau of Economic Analysis regional data portal: BEA regional data downloads.

[7]State population data are retrieved from the U.S. Census Bureau population estimates program. For earlier years, I use the 2010–2019 state population file; for later years, I use the 2020–2025 state population file.

the SDID specifications without these macroeconomic controls as the primary robustness estimates and the specifications with controls as a sensitivity exercise rather than as a preferred specification.

1.4 Conceptual Framework and Mechanisms

This section lays out the economic channels through which abortion restrictions may affect mortgage credit outcomes. The goal is not to estimate a structural model of household fertility and mortgage choice, but rather to clarify what patterns should be expected in the data under alternative mechanisms. This is useful because mortgage outcomes combine several margins: households decide whether and when to apply, lenders decide whether to approve an application, and interest rates are observed only for loans that survive the approval margin. The joint behavior of denials, originations, and interest rates is therefore informative about the mechanism.

The key idea is that abortion bans may affect mortgage credit through both borrower-side and lender-side channels. Restrictions on abortion access may change expected fertility, expenses, labor supply, savings, and mobility. These changes may affect the applicant pool and credit demand. They may also affect lender screening if lenders revise beliefs about repayment risk, income stability, or household financial fragility. These channels are not mutually exclusive, so the empirical analysis examines which margin of the mortgage market responds most clearly.

1.4.1 Household Balance Sheets and Mortgage Demand

Consider a household deciding whether to apply for a mortgage. A mortgage application is a forward-looking decision: the household compares the expected benefits of homeownership with the expected costs of borrowing, repayment risk, mobility constraints, and future household expenses. Abortion restrictions can affect this decision even before realized fertility changes are observed. If households anticipate higher risk of an unwanted or mistimed birth, they may revise expectations about future income, expenses, and mobility. This can reduce demand for mortgage credit, delay home purchase, or change the type of loan for which the household applies.

Let the latent net benefit of applying for a mortgage be written as

$$A^*_{ist} = B_{ist} - C_{ist},$$

where B_{ist} denotes the expected benefit of applying for a mortgage for applicant i in state s and year t, and C_{ist} denotes the expected cost. A household applies if

$$A^*_{ist} \geq 0.$$

Abortion restrictions may lower B_{ist} or raise C_{ist} for female applicants through several channels. First, expected childcare and medical expenses may increase. Second, expected labor supply or earnings may decline if fertility becomes more difficult to control. Third, the value of geographic mobility may rise if households wish to move to states with greater access to reproductive care, making homeownership less attractive. Fourth, uncertainty itself may increase the option value of delaying a long-term financial commitment.

A pure credit-demand channel has ambiguous implications for observed mortgage outcomes. If fewer financially fragile female applicants apply, the applicant pool may become positively selected, which could lower denial rates. If instead more constrained households apply because they are attempting to secure housing before further economic deterioration, denial rates could increase. Similarly, originations may fall either because fewer female applicants apply or because lenders reject a larger share of applications. For this reason, application and origination outcomes alone cannot distinguish demand-side responses from lender screening. The analysis therefore places particular emphasis on denial rates, which more directly capture lender approval decisions among submitted applications.

1.4.2 Lender Screening and the Approval Margin

The second channel operates through lender screening. Mortgage lenders evaluate whether an applicant is likely to repay the loan, using observable borrower characteristics such as income, loan amount, debt-to-income ratios, loan-to-value ratios, credit history, and loan type.HMDA does not observe all components of the underwriting decision. In particular, it does not report applicant credit scores or complete credit histories, although it observes several other central underwriting inputs and the final decision to approve or deny an application. This limitation means that an increase in denials conditional on

observed characteristics may reflect either a change in lender screening or a deterioration in dimensions of borrower risk that are not observed in HMDA.

A simple way to represent lender screening is through a latent repayment-risk index:

$$R_{ilst} = X'_{ilst}\Gamma + \theta\,(\text{Female}_i \times \text{Ban}_s \times \text{Post}_{st}) + \mu_l + \alpha_s + \gamma_t + \nu_{ilst},$$

where R_{ilst} is perceived repayment risk for applicant i, lender l, state s, and year t. The vector X_{ilst} includes observable borrower and loan characteristics, μ_l captures lender-specific underwriting standards, α_s captures persistent state-level differences, and γ_t captures common year shocks. The parameter θ summarizes the gender-specific change in perceived risk associated with abortion bans.

The lender approves an application if perceived risk is below an approval threshold:

$$\text{Approve}_{ilst} = 1\{R_{ilst} \leq \bar{R}_{lt}\}.$$

An abortion ban can increase denials for female applicants if it shifts the perceived-risk distribution upward for women in treated states, or if lenders effectively tighten the approval threshold for applicants whose household finances are viewed as more exposed to the policy shock. This need not imply explicit animus or direct observation of pregnancy risk. Lenders may respond to changes in expected income stability, household expenses, or local economic uncertainty. The empirical design cannot directly observe beliefs, but it can test whether the approval margin changes differentially for female applicants after abortion bans.

This framework highlights why low-income applicants are especially informative. Applicants close to the approval threshold are most sensitive to small changes in perceived risk or underwriting standards. A modest increase in expected financial fragility may have little effect on applicants who are far above the approval threshold, but it can substantially increase denial probabilities for marginal borrowers. Therefore, a larger effect among low-income applicants would be consistent with an approval-threshold mechanism.

1.4.3 Heterogeneous Lender Responses

The screening channel may also vary across lenders. Mortgage lenders differ in business models, underwriting technologies, risk tolerance, geographic exposure, and institutional culture. I use the political environment of the lender's headquarters or reporting-office state as a proxy for this culture. In particular, lenders headquartered or reporting in Republican-leaning states may be more conservative in their institutional environment, and therefore may respond more strongly to the post-*Dobbs* policy environment.

This proxy does not directly observe the beliefs of loan officers, the internal culture of the lender, or the ultimate owner of the institution. However, the measure is still useful to test whether the gender-specific denial response is stronger among lenders associated with more conservative environments. This also provides a natural follow-up to the age heterogeneity results. If the denial effect also appears among women aged 45 and above, then the mechanism is unlikely to be only a narrow response to individual pregnancy risk. A broader lender-culture or screening interpretation becomes more plausible.

To capture this idea, allow the gender-specific risk shift after abortion bans to vary across lenders:

$$R_{ilst} = X'_{ilst}\Gamma + \theta_l \left(\text{Female}_i \times \text{Ban}_s \times \text{Post}_{st}\right) + \mu_l + \alpha_s + \gamma_t + \nu_{ilst},$$

where θ_l is a lender-specific response to the abortion-policy shock. Let H_l indicate whether lender l is headquartered or reporting in a Republican-leaning state. I write

$$\theta_l = \theta_0 + \theta_1 H_l.$$

The parameter θ_1 captures whether the post-ban increase in perceived risk for female applicants is larger among Republican-HQ lenders. In the empirical implementation, this is tested using a heterogeneous triple-difference specification that compares Republican-HQ and non-Republican-HQ lenders operating in the same borrower state-year, while controlling for permanent lender differences.

I do not interpret headquarters political environment as direct evidence of lender ideology or intentional discrimination. I interpret it as a proxy for company culture or institutional environment. A positive θ_1 would suggest that the gender-specific denial

response to abortion bans is stronger among lenders associated with more conservative environments. This is precisely the type of pattern one would expect if the post-*Dobbs* effect reflects broader screening of female applicants rather than only a response to the fertility risk of younger women.

1.4.4 Selection into Approval and Conditional Interest Rates

Interest rates provide a different type of information because the interest-rate analysis is restricted to originated loans. This creates an important selection issue. If abortion bans cause lenders to deny more marginal female applicants, then the pool of female applicants who still receive loans becomes more positively selected. As a result, interest rates among originated female loans may fall even if women's overall access to credit worsens.

Let the interest rate offered to approved borrower i be

$$r_{ilst} = \rho R_{ilst} + Z'_{ilst}\Pi + \eta_{ilst},$$

where r_{ilst} is the interest rate, R_{ilst} is borrower risk, and Z_{ilst} includes loan characteristics relevant for pricing. This pricing equation is observed only when $\text{Approve}_{ilst} = 1$. If treatment raises denial probabilities for marginal female applicants, then the approved sample excludes more high-risk female borrowers. As a result, the average observed risk among approved female borrowers can decline:

$$\begin{aligned} &E[R_{ilst} \mid Approve_{ilst} = 1,\ Female_i = 1,\ Ban_s \times Post_{st} = 1] \\ &\quad < E[R_{ilst} \mid Approve_{ilst} = 1,\ Female_i = 1,\ Ban_s \times Post_{st} = 0]\,. \end{aligned}$$

This does not imply that the overall position of female applicants improves. Rather, it means that the composition of approved female borrowers becomes more positively selected. This selection mechanism predicts a pattern in which denial rates rise, originations fall, and conditional interest rates do not rise. Conditional interest rates may even decline if the rejected applicants would otherwise have received higher rates.

This point is important for interpreting the empirical results. A decline in the female–male interest-rate gap among originated loans should not be interpreted as evidence that abortion bans improved women's access to mortgage credit. It may instead indicate

that the women who remain in the approved pool after treatment are selected from a stronger part of the applicant distribution. The combination of higher denials, lower originations, and lower conditional interest rates is therefore consistent with screening at the approval margin. This selection mechanism provides one possible explanation for a decline in conditional interest rates, but the interest-rate regression alone cannot establish that this mechanism is operating. Because rates are observed only for originated loans, distinguishing selection into approval from changes in loan pricing would require additional assumptions or a formal selection correction. I therefore treat the conditional interest-rate result as secondary evidence.

1.4.5 Predicted Empirical Patterns

The discussion above implies several possible mechanisms. Before turning to the empirical results, it is useful to state what each mechanism would predict for the three main outcomes. The labels used in Table 1.7 summarize these mechanisms.

The first mechanism is *approval-margin screening*. Under this mechanism, abortion bans increase perceived repayment risk or financial fragility for female applicants relative to male applicants in treated states. The effect should appear most directly in denial rates: female applicants should become more likely to be denied, and originations should fall. Interest rates among originated loans need not increase. If marginal female applicants are screened out, the remaining approved female borrowers may be positively selected, so conditional interest rates may remain flat or even decline.

A narrow version of this mechanism is an *individual fertility-risk channel*. If lenders respond mainly to the possibility that a specific female applicant may experience an unwanted or mistimed birth, the effect should be concentrated among applicants in prime reproductive ages. This prediction motivates the age heterogeneity analysis below. By contrast, if the effect reflects broader gender salience or discriminatory screening, it need not be limited to younger women.

The second mechanism is a *uniform worsening of female credit terms*. This is a stronger form of differential treatment in which female applicants face worse outcomes on both the approval and pricing margins. This mechanism would predict higher denials, lower originations, and higher interest rates among approved female borrowers. This pattern is different from approval-margin screening because it implies that women who

remain approved should also receive worse pricing.

The third mechanism is a *borrower-demand response*. Abortion restrictions may change whether, when, or where women apply for mortgages. For example, some households may delay home purchase, reduce demand for mortgage credit, or move across states. The implications for application-level denial and origination rates are ambiguous because they depend on selection into applying. If financially constrained applicants stop applying, observed denial rates could fall; if more constrained applicants apply before conditions worsen, denial rates could rise.

The fourth mechanism is *applicant-composition change*. Under this explanation, the estimated treatment effect is not driven by a change in lender screening or perceived risk, but by changes in the observable characteristics of female applicants in ban states. For example, female applicants after treatment may have lower income, higher debt-to-income (DTI) ratios, higher loan-to-value (LTV) ratios, or other characteristics associated with higher denial probabilities. This mechanism predicts changes in outcomes only because the applicant pool changes.

The fifth mechanism is a *common state-level mortgage shock*. This is not a gender-specific mechanism, but it is an important benchmark. A local housing-market or credit-market shock may affect all applicants in a treated state-year. If it affects female and male applicants similarly, it should not by itself generate a change in the female–male gap. This is why the main specification compares female applicants to male applicants.

The sixth mechanism is *gender salience or discriminatory screening*. Under this mechanism, abortion bans may affect credit decisions not only by changing actual borrower risk, but also by making female status more salient in treated states. This channel does not require lenders to assess the actual fertility risk of each applicant. Instead, it can operate through broader stereotypes, implicit bias, or statistical discrimination about women's household responsibilities, labor-market attachment, or financial stability. This mechanism predicts higher denials and lower originations for female applicants relative to male applicants. Interest rates among originated loans may remain flat or decline if the effect operates mainly through selection at the approval margin.

Finally, *lender-environment heterogeneity* refers to the possibility that lenders differ in how they respond to the post-*Dobbs* environment. If lender screening responses vary by headquarters or reporting-office environment, the average effect may be stronger for

some lenders than for others. This mechanism does not generate a single prediction for the aggregate sign, but it predicts heterogeneous effects across lender environments.

Table 1.7: Predicted Signs Under Alternative Mechanisms

Mechanism	Denials	Originations	Interest rates
Approval-margin screening	+	−	0/−
Uniform worsening of female credit terms	+	−	+
Borrower-demand response	Ambiguous	Ambiguous	Ambiguous
Applicant-composition change	Ambiguous	Ambiguous	Ambiguous
Common state-level mortgage shock	None	None	None
Gender salience / discriminatory screening	+	−	0/−
Lender-environment heterogeneity	Heterogeneous	Heterogeneous	Ambiguous

Notes: The table reports qualitative predictions for female applicants relative to male applicants in ban states after abortion restrictions become binding. A + sign indicates an increase in the female–male gap; a − sign indicates a decline. "None" means that the mechanism does not by itself predict a gender-specific treatment effect. Interest rates are analyzed only among originated loans and therefore reflect both pricing and selection into approval.

Table 1.7 should be read as a guide for interpreting the empirical patterns, not as a claim that only one mechanism can operate. Several channels may be present at the same time. The most important distinction is between mechanisms that operate through the approval margin and mechanisms that imply uniformly worse credit terms for approved female borrowers. If denials rise, originations fall, and conditional interest rates rise, the evidence would point toward broadly worse credit terms for women. If denials rise, originations fall, and conditional interest rates do not rise or decline, the evidence is more consistent with selection at the approval margin.

The empirical results below show the clearest deterioration at the approval margin. Denial rates increase for female applicants relative to male applicants and originations fall. Conditional interest rates decline modestly in the borrower-level estimates, which is inconsistent with a simple story in which all originated female borrowers receive uniformly worse pricing. However, because the interest-rate outcome is conditional on origination, its decline does not by itself distinguish selection into approval from other changes in the composition or pricing of originated loans.

The framework also motivates the heterogeneity and composition tests. If the mechanism operates through credit constraints, the effects should be larger among applicants closer to the approval threshold, such as low-income applicants. If the mechanism operates narrowly through individual pregnancy risk, the effects should be concentrated among women in prime reproductive ages, while a broader lender-side response need not

be concentrated among women in prime reproductive ages. If the mechanism reflects lender screening or institutional environment, effects may differ across lenders depending on their headquarters or reporting-office environment. Finally, if the results are driven by applicant composition, one should observe post-treatment changes in income, LTV, DTI, or the female share of applicants.

A more detailed mapping from these mechanism labels to their economic interpretation is reported in Appendix C. I keep the main text focused on the testable implications because the same logic is used repeatedly in the results section. The empirical analysis below uses these predictions in two ways. First, the joint movement of denials, originations, and conditional interest rates is used to assess whether the deterioration is concentrated in approval outcomes or is accompanied by worse pricing among originated loans. Because the interest-rate outcome is selection-sensitive, it is not used by itself to identify the underlying mechanism. Second, the heterogeneity and composition tests evaluate whether the effect is concentrated among credit-constrained applicants, whether it is narrowly concentrated among women in prime reproductive ages, whether it differs by lender environment, and whether observable applicant composition changes after treatment.

1.5 Empirical Strategy

1.5.1 Borrower-Level Triple-Difference Design

The primary empirical strategy is a borrower-level triple-difference design that exploits variation in abortion policy across U.S. states and over time. The design compares female applicants to male applicants, in abortion-ban states relative to protected states, before and after abortion restrictions become binding. Equivalently, it can be viewed as a gender-differential difference-in-differences specification: the policy difference-in-differences comparison is interacted with applicant gender. Throughout the paper, I use the term borrower-level triple-difference for the application-level specifications because the coefficient of interest compares female versus male applicants, ban versus protected states, and post- versus pre-treatment periods. When the analysis is conducted on state-year female–male gaps, the gender difference is already embedded in the dependent variable; those specifications are therefore described as state-year DiD or SDID estimates

on gender gaps.

The baseline estimating equation is

$$Y_{ist} = \beta\,(\text{Female}_i \times \text{Treated}_{st}) + \delta \text{Treated}_{st} + \pi \text{Female}_i + X'_{ist}\Gamma + \alpha_s + \gamma_t + \varepsilon_{ist}, \quad (1.1)$$

where Y_{ist} denotes the mortgage outcome for applicant i in state s and year t. The three main outcomes are denial status, origination status, and the interest rate conditional on origination. Denial and origination regressions are estimated on the decisioned application sample, while interest-rate regressions are estimated only on originated loans.

The treatment variable is defined as

$$\text{Treated}_{st} = \text{BanState}_s \times \text{Post}_{st},$$

where BanState_s equals one for abortion-ban states and zero for protected states. The variable Post_{st} is state-specific because treatment timing is staggered. For Texas, the post-treatment period begins in 2022. For other ban states, the post-treatment period begins in 2023. Protected states are never treated during the sample period. Thus, Treated_{st} equals one only for treated state-years after abortion restrictions become binding.

The coefficient of interest is β. Because $\text{Treated}_{st} = \text{BanState}_s \times \text{Post}_{st}$, the coefficient on $\text{Female}_i \times \text{Treated}_{st}$ is equivalent to the coefficient on $\text{Female}_i \times \text{BanState}_s \times \text{Post}_{st}$. It captures the differential post-treatment change in outcomes for female applicants relative to male applicants in ban states, compared with the corresponding female–male change in protected states. In other words, β measures whether abortion bans widened or narrowed the gender gap in mortgage outcomes. I therefore refer to this coefficient as the triple-difference estimate.

The vector X_{ist} includes borrower- and loan-level controls. The baseline specification controls for log applicant income, log loan amount, conforming-loan status, loan-to-value bins, and debt-to-income bins. Richer specifications additionally control for applicant race, ethnicity, and age categories. These controls absorb observable differences in borrower financial capacity, loan characteristics, leverage, repayment burden, and demographic composition. State fixed effects, α_s, absorb permanent differences across states, while year fixed effects, γ_t, absorb national shocks common to all states in a given year,

such as changes in interest rates, macroeconomic conditions, or national mortgage-market trends. I additionally examine robustness to allowing gender gaps to vary flexibly across states and over time by including state-by-female and year-by-female fixed effects; these results are reported in Section 1.7.3.

The specification is estimated as a linear probability model for denial and origination outcomes. This choice makes the coefficients directly interpretable as percentage-point changes and allows the inclusion of high-dimensional fixed effects and interaction terms. For the interest-rate outcome, the same specification is estimated on the originated-loan sample.

1.5.2 Interpretation of the Main Coefficient

The coefficient β is the triple-difference coefficient; it is not the overall effect of abortion bans on all mortgage applicants. Instead, it is the effect on the female–male gap in mortgage outcomes. This distinction is important for interpreting the paper's findings. Abortion bans may affect local mortgage markets generally, for example through housing demand, migration, or macroeconomic changes. The empirical design differences out such common state-year changes to the extent that they affect female and male applicants similarly.

For denial rates, a positive β means that female applicants in ban states became more likely to be denied relative to comparable male applicants after abortion restrictions became binding. For originations, a negative β means that female applicants became less likely to obtain originated loans relative to male applicants. For interest rates, a negative β means that among originated loans, the female–male interest-rate gap declined after treatment.

The sign of β has different interpretations across outcomes. For denials, a positive coefficient indicates a widening female–male denial gap. For originations, a negative coefficient indicates a decline in realized mortgage access for female applicants relative to male applicants. For interest rates, the coefficient is conditional on origination and therefore reflects both pricing and selection into the approved-loan sample. The mechanism section explains why these outcomes should be interpreted jointly.

1.5.3 Identification Assumption

The identifying assumption is that, absent abortion bans, the female–male gap in mortgage outcomes would have evolved similarly in ban and protected states, conditional on controls and fixed effects. Formally, the counterfactual assumption is that the change in outcomes for female applicants relative to male applicants in ban states would have followed the same trend as the female–male gap in protected states if abortion restrictions had not become binding.

This assumption is weaker than requiring female applicants in ban and protected states to have the same outcome levels. The design allows for persistent differences across states through state fixed effects and allows for common national shocks through year fixed effects. It also allows female and male applicants to have different baseline outcomes. The identifying variation comes from whether the gender gap changes differentially in ban states after treatment.

The assumption would be violated if, at the same time as abortion bans were implemented, ban states experienced other shocks that differentially affected female mortgage applicants relative to male applicants and were not present in protected states. For example, a state-specific change in labor-market conditions affecting women more than men could confound the estimates if it coincided with abortion-policy changes. The empirical strategy addresses this concern in several ways. First, the event-study analysis examines whether gender gaps were already diverging before treatment. Second, the composition tests examine whether observable applicant characteristics changed differentially after treatment. Third, the robustness section uses synthetic difference-in-differences and randomization-inference exercises at the state-year level.

1.5.4 Event-Study Specification

To examine dynamic treatment effects and assess pre-treatment trends, I estimate an event-study version of the triple-difference specification. Because treatment timing is staggered, I define event time relative to each state's treatment year. Event time 0 denotes the year immediately before treatment, and event time 1 denotes the first treated year. Thus, for Texas, 2021 is event time 0 and 2022 is event time 1. For other ban states, 2022 is event time 0 and 2023 is event time 1.

The event-study specification is

$$Y_{ist} = \sum_{k \neq 0} \beta_k \left(\text{Female}_i \times \mathbf{1}\{\text{EventTime}_{st} = k\} \right) + \sum_{k \neq 0} \delta_k \mathbf{1}\{\text{EventTime}_{st} = k\} + \pi \text{Female}_i + X_{ist}' \Gamma + \alpha_s + \gamma_t + \varepsilon_{ist}. \tag{1.2}$$

The omitted category is event time 0, the year immediately before treatment. Therefore, each coefficient β_k measures the female–male gap at event time k relative to the female–male gap in the last pre-treatment year. Pre-treatment coefficients provide evidence on whether gender gaps were already changing before abortion bans became binding. Post-treatment coefficients show how the gender gap evolves after treatment.

The event-study graphs are estimated using state-year-sex cells constructed from the HMDA microdata and weighted by the number of applications in each cell. I use this aggregated specification for graphical diagnostics because abortion-policy exposure varies at the state-year level and because the object of interest is the evolution of female–male outcome gaps over event time. The main treatment effects in Table 1.8 are estimated using the full borrower-level data. The event-study controls include cell-level averages of log income, log loan amount, and conforming-loan status, as well as cell-level shares for LTV bins, DTI bins, applicant race categories, ethnicity categories, and age categories.

1.5.5 Inference

Because treatment varies at the state-year level, standard errors are clustered at the state level. This allows for arbitrary correlation in unobserved shocks within states over time. Clustering at the state level is also conservative relative to application-level heteroskedasticity because the identifying policy variation is not at the individual applicant level.

The number of state clusters is moderate, so I also report robustness checks that do not rely only on asymptotic cluster-robust inference. In particular, the robustness section includes randomization-inference exercises using placebo treatment assignments across states. These placebo exercises evaluate whether the estimated effects are unusually large relative to the distribution of effects obtained under random assignment of treatment

status.

1.5.6 Synthetic Difference-in-Differences as Robustness

The borrower-level triple-difference design is the primary empirical strategy because it uses the full application-level HMDA data and allows for rich borrower and loan controls. As a robustness check, I also estimate synthetic difference-in-differences models using state-year female–male gaps in mortgage outcomes. The SDID analysis aggregates the data to the level at which abortion policy varies and constructs a synthetic comparison group from protected states that better matches the pre-treatment trajectory of treated states. Because the outcome is already the female–male gap, the SDID coefficient is the state-year analogue of the borrower-level triple-difference coefficient.

The logic of SDID combines the intuition of difference-in-differences with the reweighting logic of synthetic control methods (Abadie et al., 2010; Abadie, 2021; Arkhangelsky et al., 2021). A standard DiD compares treated and control states using equal weights, after controlling for fixed effects and covariates. This can be restrictive if protected states differ from ban states in baseline outcomes or in the pre-treatment evolution of gender gaps. SDID addresses this concern by constructing a weighted comparison group from the donor pool of protected states. The weights are chosen so that the weighted protected states match the pre-treatment path of the treated ban states as closely as possible. The identifying idea is therefore not that ban and protected states would have followed parallel trends in an unweighted comparison, but that a weighted combination of protected states can approximate the counterfactual trajectory of ban states in the absence of treatment. In this sense, SDID relaxes the standard parallel-trends requirement by replacing it with a synthetic-control-style pre-treatment fit condition.

I perform the SDID exercise to check whether the main denial result remains visible when the analysis is conducted at the state-year level and when protected states are reweighted to match treated states' pre-treatment gender-gap dynamics. The SDID robustness check is implemented using the same staggered treatment timing as the main analysis: Texas is treated beginning in 2022, while the remaining ban states enter treatment in 2023. This timing choice keeps the robustness exercise aligned with the borrower-level design and produces cohort-specific graphical diagnostics. The main text reports the SDID estimates in the Results section, while Appendix B reports separate SDID plots for

Texas and for the remaining ban states. SDID is useful because it relaxes the standard unweighted parallel-trends comparison by requiring that a weighted combination of donor states reproduce the treated states' pre-treatment outcome path.

I estimate the SDID models both without and with contemporaneous state-level macroeconomic controls. Because some of these covariates may themselves be affected by abortion restrictions, the specification without these controls is my primary SDID robustness specification. The controlled estimates are reported to assess sensitivity to contemporaneous state-level economic conditions, but I do not interpret them as necessarily providing a cleaner estimate of the policy effect.

1.5.7 Summary

The empirical strategy therefore proceeds in four steps. First, I estimate borrower-level triple-difference models for denials, originations, and interest rates. Second, I examine event-study estimates to assess pre-treatment trends and the timing of the effects. Third, I study heterogeneity by income, age, and lender headquarters environment to evaluate the mechanisms suggested by the conceptual framework. I also test whether the results can be explained by changes in the applicant pool rather than by changes in approval outcomes conditional on observable characteristics. Fourth, I conduct robustness checks using synthetic difference-in-differences and randomization inference. Together, these exercises assess whether abortion bans changed women's relative access to mortgage credit and whether the effect is concentrated at the approval margin.

1.6 Results

1.6.1 Main Results: Borrower-Level Evidence

Table 1.8 reports borrower-level triple-difference estimates of the impact of abortion bans on mortgage lending outcomes. The coefficient of interest is the interaction term Female $\times$ Ban state $\times$ Post, which captures the differential post-policy change in outcomes for female applicants relative to male applicants in treated states.

The most robust and economically significant effect is observed for denial rates. Across specifications, the estimated coefficient ranges from 0.0090 to 0.0094 and is statistically

Table 1.8: Triple-difference Estimates of Abortion Bans on Mortgage Outcomes

	Originations		Denials		Interest Rate	
	A	B	A	B	A	B
Female	0.0018**	0.0040***	-0.0013*	-0.0033***	0.0181***	0.0178***
	(0.0007)	(0.0008)	(0.0007)	(0.0008)	(0.0024)	(0.0026)
Ban state × Post	-0.0082***	-0.0066***	0.0063***	0.0048***	-0.0747*	-0.0720*
	(0.0026)	(0.0021)	(0.0021)	(0.0017)	(0.0410)	(0.0397)
Female × Ban × Post	-0.0087***	-0.0084***	0.0094***	0.0090***	-0.0259***	-0.0266***
	(0.0025)	(0.0019)	(0.0023)	(0.0018)	(0.0049)	(0.0050)
Observations	17057362	15125238	17057362	15125238	15080021	13422392

Standard errors in parentheses

Notes: Each column reports a borrower-level triple-difference estimate with state and year fixed effects. The reported coefficient of main interest is **Female** × **Ban state** × **Post**, which captures the differential post-ban change in outcomes for female applicants relative to male applicants in treated states. Denial and origination regressions are estimated on the decisioned sample, while interest rate regressions are estimated on originated loans only. Columns labeled *A* include log income, log loan amount, a conforming loan indicator, and LTV and DTI bins. Columns labeled *B* additionally control for applicant race, ethnicity, and age. Standard errors are clustered at the state level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

significant at the 1 percent level. This implies that abortion bans increase the probability that a female applicant is denied relative to a comparable male applicant by approximately 0.9 percentage points.

I also find a statistically significant decline in mortgage originations for female applicants relative to male applicants. The estimated coefficients, between -0.0084 and -0.0087, suggest that abortion bans reduce the likelihood of loan origination for women relative to men. This pattern is consistent with the increase in denial rates and indicates that the primary margin of adjustment operates through credit approval rather than only through application behavior.

For interest rates, the estimates indicate a statistically significant decline in the female–male gap of approximately 2.6 basis points. The magnitude is economically small, and the estimate is conditional on loan origination. Because the set of originated loans may itself change in response to the policy, I do not interpret this coefficient as evidence of improved credit terms for female borrowers or as direct evidence of positive selection at the approval margin.

The main evidence therefore comes from the approval outcomes: abortion bans are

associated with higher denial rates and lower originations for female applicants relative to male applicants. The small decline in conditional interest rates is compatible with selection among originated loans, but without a formal selection correction I cannot distinguish this explanation from other changes in the composition or pricing of originated loans. I consequently treat the interest-rate estimate as a secondary result rather than as evidence identifying the mechanism.

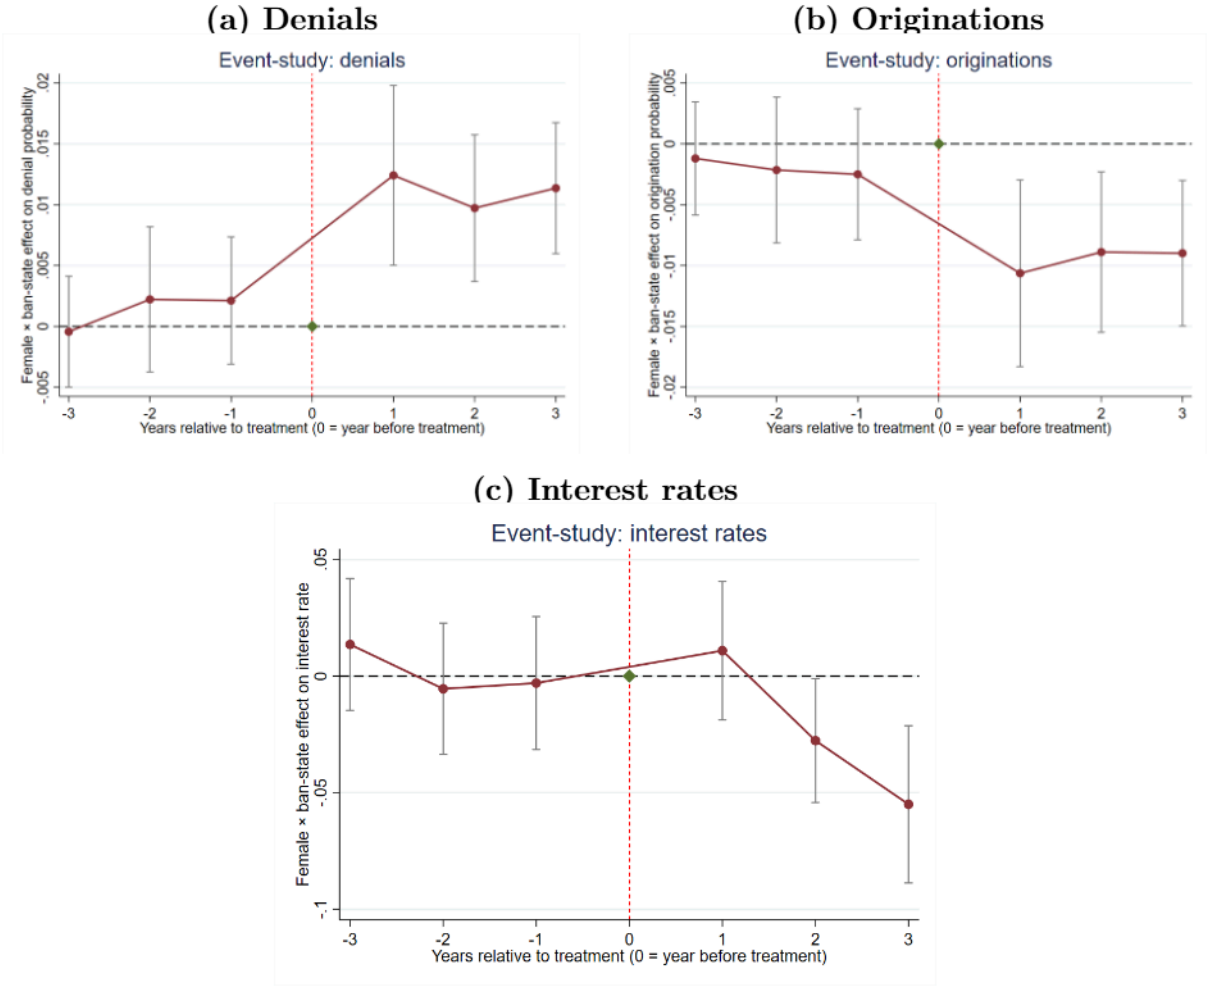
1.6.2 Event-Study Evidence and Pre-Trends

Figure 1.1 presents event-study estimates for the dynamic effects of abortion bans on the three main mortgage outcomes. Because treatment timing is staggered, with Texas treated earlier than other ban states, I index the horizontal axis by event time rather than by calendar year. Event time is defined so that 0 corresponds to the year immediately before treatment and 1 corresponds to the first treated year. Thus, for Texas, 2021 is event time 0 and 2022 is event time 1; for other ban states, 2022 is event time 0 and 2023 is event time 1. The coefficient at event time 0 is omitted from the regression and normalized to zero.

The event-study estimates are based on state-year-sex cells constructed from the HMDA microdata and weighted by the number of applications in each cell. This aggregation produces a transparent state-level diagnostic while preserving the distinction between female and male applicants within each state-year. The specification replaces the single post-treatment interaction from the main triple-difference model with a set of event-time indicators interacted with the female indicator. Event time is defined relative to the implementation of abortion restrictions, with event time 0 denoting the year immediately before treatment and event time 1 denoting the first treated year. Thus, the coefficient at event time 0 is the omitted reference period and is normalized to zero in the figure.

The event-study regressions include state and year fixed effects and a set of cell-level controls. Specifically, the controls include cell-level averages of log income, log loan amount, and conforming-loan status, as well as cell-level shares for LTV bins, DTI bins, applicant race categories, ethnicity categories, and age categories. Denial and origination outcomes are calculated on the decisioned sample, while interest rates are calculated only for originated loans. Standard errors are clustered at the state level, and the vertical bars

Figure 1.1: Event-study estimates of abortion bans on gender gaps in mortgage outcomes



Notes: The figure plots event-study estimates using state-year-sex cells constructed from HMDA microdata and weighted by the number of applications in each cell. Event time is defined so that 0 is the year immediately before treatment and 1 is the first treated year. Texas is assigned treatment year 2022, while other ban states are assigned treatment year 2023. The point at event time 0 is omitted from the regression and normalized to zero. Denial and origination outcomes are calculated on the decisioned sample; interest rates are calculated only for originated loans. Controls include cell-level averages of log income, log loan amount, and conforming-loan status, as well as cell-level shares for LTV bins, DTI bins, applicant race categories, ethnicity categories, and age categories. All regressions include state and year fixed effects. Standard errors are clustered at the state level, and vertical bars report 95 percent confidence intervals. Event time 3 is identified by Texas, the early-treated state.

report 95 percent confidence intervals.

Panel (a) presents the event-study estimates for mortgage denials. The coefficients become larger around the treatment period and remain positive in the post-treatment years. Some movement is visible close to treatment, which is consistent with the fact that the *Dobbs* policy environment became salient before all bans formally took effect. The important point is that the figure does not show a steady long-run increase in the female–male denial gap several years before treatment.

Panel (b) reports the corresponding estimates for originations. The pre-treatment coefficients are close to zero and slightly negative, without evidence of a sharp pre-treatment divergence between female and male applicants in ban states relative to protected states. After treatment, the coefficients become more negative, indicating a decline in female originations relative to male originations after abortion bans. This pattern complements the denial result. If female applicants become more likely to be rejected at the approval margin, then fewer female applicants should ultimately receive originated loans. Because origination combines lender approval and borrower take-up, it is a less direct measure of lender decision-making than denial status. Nevertheless, the post-treatment movement in originations is consistent with reduced realized access to mortgage credit for female applicants.

For both denials and originations, there is some movement (although not statistically distinguishable from zero) before formal implementation, which may reflect anticipation, legal uncertainty, public debate, and salience before bans formally become binding, consistent with the dates discussed in Section 1.2.

Panel (c) shows the event-study estimates for interest rates. These estimates are more volatile, which is expected because interest rates are analyzed only among originated loans. The pre-treatment coefficients do not show a systematic pattern that would threaten the main interpretation. Post-treatment coefficients move downward, consistent with the borrower-level results, but this outcome should be read as selection-sensitive and secondary to the denial margin.

Taken together, the three panels provide evidence consistent with the main results. The denial panel shows the clearest post-treatment increase, which is important because denial status is the outcome most directly tied to lender approval decisions. The origination panel moves in the opposite direction, consistent with fewer female applicants receiv-

ing mortgage credit. The interest-rate panel does not indicate uniformly worse pricing for originated female borrowers, but its post-treatment decline is selection-sensitive and does not separately identify the mechanism. Thus, the dynamic evidence primarily reinforces the central result that abortion bans widened the female–male denial gap and reduced relative originations.

Because the HMDA sample is very large, small differences can sometimes be estimated precisely. I therefore interpret the event-study evidence by considering both statistical significance and economic magnitude. The central diagnostic question is whether the female–male gap was already increasing systematically long before treatment. The event-study evidence does not indicate such a long-running pre-existing increase for the primary outcome, denial rates. Instead, the most pronounced divergence in denials appears after abortion bans become effective. For originations and interest rates, the event-study estimates also provide no evidence of statistically significant pre-treatment divergence. The later post-treatment estimates should be interpreted with some caution because event time 3 is identified by Texas, the early-treated state, but the overall post-treatment pattern is consistent with the borrower-level triple-difference estimates.

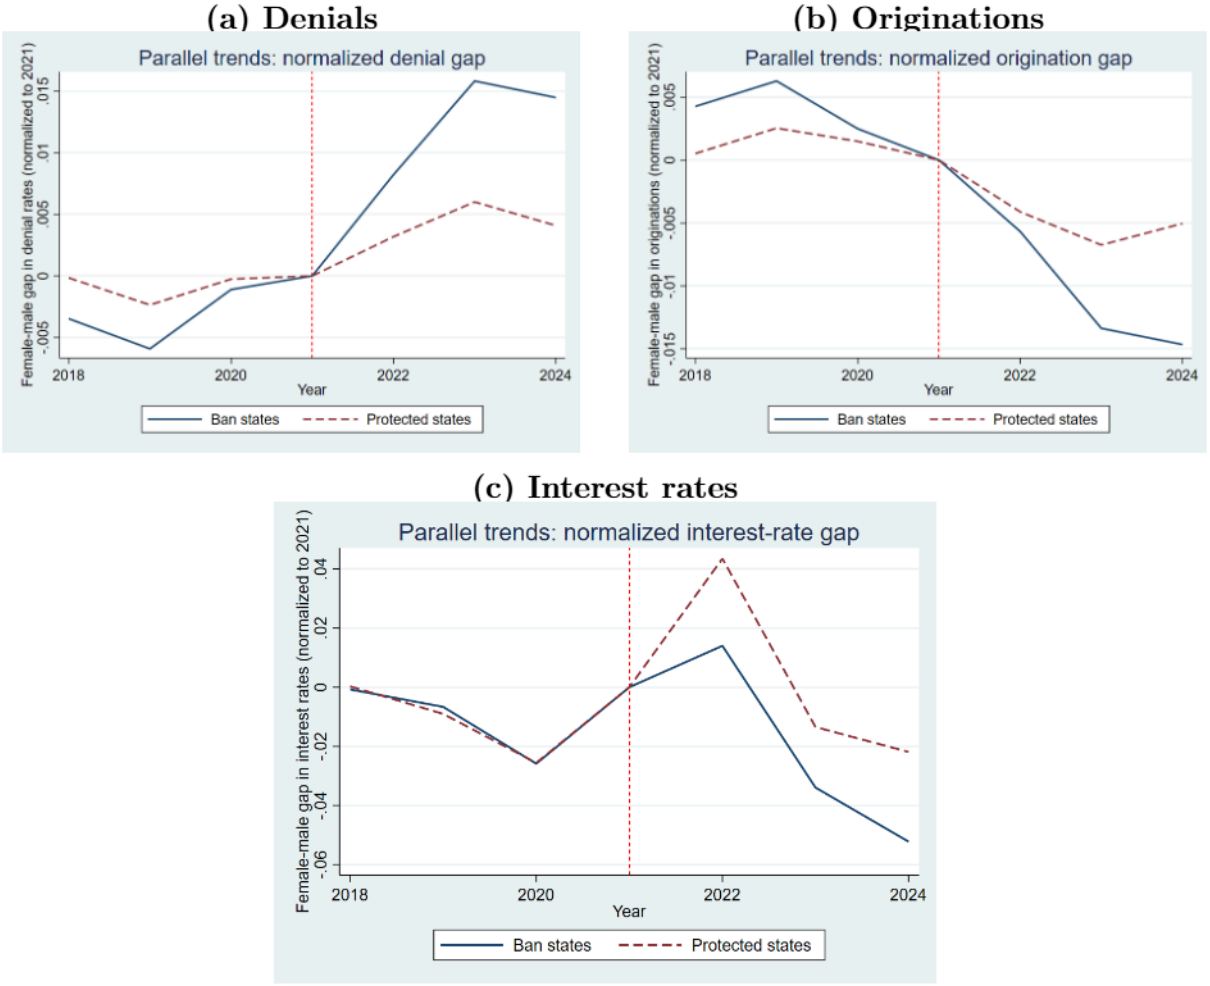
1.6.3 Parallel Trends in Gender Gaps

Figure 1.2 complements the event-study evidence by plotting raw gender gaps in mortgage outcomes for ban and protected states over calendar time. For each outcome, I calculate the female–male gap within each state-year and normalize the series to 2021. This normalization makes it easier to compare the evolution of gender gaps across treated and control states before the post-*Dobbs* policy divergence.

The figure is useful as a descriptive check because it shows the underlying state-year patterns before imposing the borrower-level regression structure. Across the three outcomes, treated and control states follow broadly similar trajectories in the pre-treatment period. Some year-to-year fluctuations are visible, but the figure does not show a systematic widening of the female–male denial gap in ban states before the policy change. This visual pattern is consistent with the event-study evidence in Figure 1.1.

After the policy change, the divergence is most pronounced for denial rates. The female–male denial gap increases in ban states relative to protected states, matching the main triple-difference estimates. The origination-rate gap moves in the opposite

Figure 1.2: Descriptive trends in female–male mortgage outcome gaps for Ban and Protected states



Notes: The figure plots state-year female–male gaps in mortgage outcomes separately for ban and protected states, normalized to 2021. The denial gap is the average denial rate for female applicants minus the average denial rate for male applicants within a state-year. The origination-rate gap and interest-rate gap are defined analogously. Denial and origination-rate gaps are calculated on the decisioned sample; interest-rate gaps are calculated on originated loans only. The figure is descriptive and does not adjust for borrower-level controls. Because treatment timing is staggered, the event-study estimates in Figure 1.1 provide the main formal pre-trends diagnostic.

direction, consistent with a relative decline in female originations. The interest-rate gap declines modestly, which is also consistent with the main borrower-level results. Thus, the descriptive trends reinforce the main pattern in the gender gaps.

Because treatment timing is staggered, the normalized calendar-time trends should not be interpreted as the primary test of parallel trends. Texas is exposed to binding restrictions earlier than other ban states, while other treated states enter treatment later. For this reason, the event-study specification in Figure 1.1, which indexes observations by years relative to treatment, provides the main pre-trends diagnostic. Figure 1.2 instead serves as a transparent descriptive visualization of the raw gender-gap patterns over time.

1.6.4 Heterogeneity: Evidence on the Credit Constraint Channel

To better understand the mechanisms underlying the main results, I examine whether the borrower-level triple-difference effect varies across applicant characteristics. I focus first on heterogeneity by income, which speaks directly to the role of credit constraints. If abortion bans affect mortgage access by increasing perceived financial fragility or by pushing applicants closer to underwriting thresholds, the effect should be stronger among lower-income applicants.

I estimate heterogeneous triple-difference specifications that interact the main treatment effect, $\text{Female}_i \times \text{BanState}_s \times \text{Post}_{st}$, with an indicator for low-income applicants. The resulting four-way interaction, $\text{Female}_i \times \text{BanState}_s \times \text{Post}_{st} \times \text{LowIncome}_i$, measures whether the gender-specific post-ban effect differs between low-income and higher-income applicants. I define low-income applicants as those with income below the sample median.

The results indicate that the increase in denial rates for female applicants is significantly larger among low-income applicants. In the interacted specification, the coefficient on the four-way interaction is positive and statistically significant, with an estimated magnitude of approximately 1.17 percentage points. This means that the post-ban widening of the female–male denial gap is larger among low-income applicants than among higher-income applicants.

The split-sample estimates point to the same pattern. The estimated triple-difference effect for higher-income applicants is smaller and less precisely estimated, while the effect for low-income applicants is positive, larger, and statistically significant. This pattern

Table 1.9: Heterogeneity in Denial Effects by Applicant Income

	Low income	High income	DDDD
Female × Ban state × Post	0.0121***	0.0004	0.0005
	(0.0025)	(0.0006)	(0.0008)
Additional effect: low income			0.0117***
			(0.0026)
Observations	7297879	7827359	15125238

Standard errors in parentheses

Notes: Columns (1) and (2) report split-sample borrower-level triple-difference estimates for low- and high-income applicants, respectively. Column (3) reports the full-sample four-way-interaction specification. The coefficient **Female × Ban state × Post** captures the differential post-ban change in denial probability for female applicants relative to male applicants within the indicated income group. The coefficient **Additional effect: low income** reports the additional post-ban effect for low-income applicants relative to high-income applicants. Low income is defined using the sample median income cutoff. All specifications include state and year fixed effects, log income, log loan amount, a conforming-loan indicator, LTV and DTI bins, and controls for race and ethnicity. Standard errors are clustered at the state level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

suggests that abortion bans disproportionately affect applicants who are closer to lenders' approval thresholds.

Overall, the income heterogeneity analysis supports a credit-constraint interpretation. The effects of abortion bans are not uniform across applicants; instead, the denial effect is concentrated among applicants with lower financial capacity. This is consistent with an approval-margin mechanism in which relatively small changes in perceived repayment risk or household financial fragility have the largest consequences for applicants who are already closer to being denied.

1.6.5 Heterogeneity by Applicant Age

I next examine whether the borrower-level triple-difference effect varies across finer applicant age categories. This provides a sharper test of whether the mortgage-market response is concentrated among women with greater direct fertility exposure. If the effect operates primarily through lenders' assessment of an individual female applicant's probability of an unwanted or mistimed birth, one would expect the effect to be strongest among younger applicants and to decline substantially at older ages. I therefore estimate the baseline denial specification separately for applicants aged below 25, 25–34, 35–44, 45–54, 55–64, and 65 and above.

Table 1.10: Heterogeneity in Denial Effects by Applicant Age

	<25	25–34	35–44	45–54	55–64	65+
Female × Ban state × Post	0.0080***	0.0081***	0.0098***	0.0071***	0.0082***	0.0160***
	(0.0027)	(0.0026)	(0.0023)	(0.0019)	(0.0023)	(0.0033)
Observations	925,763	5,141,054	4,105,856	2,448,817	1,545,369	958,379

Standard errors in parentheses

Notes: Each column reports a borrower-level triple-difference estimate for the indicated applicant age group. The 65+ category combines applicants aged 65–74 and above 74. All specifications include state and year fixed effects, log income, log loan amount, a conforming-loan indicator, LTV and DTI bins, and controls for race and ethnicity. Standard errors are clustered at the state level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.10 shows no evidence that the denial effect declines with applicant age. The estimated Female × Ban state × Post coefficient is 0.0080 for applicants below age 25, 0.0081 for ages 25–34, 0.0098 for ages 35–44, 0.0071 for ages 45–54, and 0.0082 for ages 55–64. For applicants aged 65 and above, the estimated effect is 0.0160. All six estimates are positive and statistically significant.

The absence of a monotonic decline with age weighs against a narrow mechanism operating solely through an individual applicant's own fertility exposure. In particular, the effect remains present among applicants well beyond prime reproductive ages and is not smaller for the oldest group. However, these age patterns do not uniquely identify the underlying mechanism. They are consistent with a response that operates more broadly across female applicants, but may also reflect age-related differences in unobserved borrower risk or other underwriting characteristics not available in HMDA. I therefore interpret the age heterogeneity as evidence against a narrow individual fertility-risk channel rather than as definitive evidence of gender salience or discriminatory screening.

There are also measurement reasons why age heterogeneity may be muted in mortgage data. HMDA records the applicant's age category, but it does not observe pregnancy status, fertility intentions, number of children, childcare responsibilities, marital dynamics, or the age and fertility exposure of other household members. Mortgage applications also reflect broader household financial circumstances, while the age category captures only one applicant's life-cycle position.

Overall, the finer age-group results indicate that the post-ban increase in the female–male denial gap is not concentrated among younger applicants and does not attenuate with age. This weakens a narrow individual fertility-risk interpretation and motivates

the next exercise, which examines whether the within-lender response varies across lender institutional environments.

1.6.6 Heterogeneity by Political Environment of Lender Headquarters

The finer age-group estimates show that the increase in the female–male denial gap does not decline with applicant age and remains present even among applicants aged 65 and above. This pattern weighs against a narrow mechanism operating solely through lenders' assessment of an individual applicant's probability of pregnancy or childbirth. At the same time, it does not by itself identify the underlying mechanism: the persistence of the effect at older ages may reflect broader lender-side responses to female applicants, changes in perceived repayment risk, or other underwriting responses not fully observed in HMDA.

To explore this possibility, I examine whether the effect varies with the political environment of the lender's headquarters or reporting-office state. I interpret this variable as a proxy for lender culture or institutional environment. The idea is that lenders associated with more conservative states may operate in organizational and market environments where gendered assumptions about household roles, labor-market attachment, or repayment risk are more likely to matter after abortion restrictions become binding. This is only a proxy. It does not observe the beliefs of loan officers, the internal culture of the lender, or the ultimate owner of the institution. However, the measure is still informative if it captures whether the lender is associated with a more conservative institutional environment.

The specification extends the main borrower-level triple-difference model by allowing the treatment effect, $\text{Female}_i \times \text{BanState}_s \times \text{Post}_{st}$, to differ with the Republican-HQ/reporting-office indicator. The pooled heterogeneity specification therefore includes the corresponding four-way interaction,

$$\text{Female}_i \times \text{BanState}_s \times \text{Post}_{st} \times \text{RepublicanHQ}_{lt},$$

together with the estimable lower-order interactions required for the heterogeneous treatment model. The borrower state-by-year fixed effects absorb the Ban state $\times$ Post term

itself. The coefficient on the four-way interaction measures whether the post-ban change in the female–male denial gap differs between Republican-HQ and non-Republican-HQ lender environments. The Republican-HQ measure is defined at the lender-year level using the state associated with the institution's HMDA reporting office and may therefore change over time for a given lender.

This heterogeneity exercise requires a stricter specification than the baseline model because Republican-HQ lenders are not randomly distributed across mortgage markets. They may be more active in ban states or in markets with different housing and credit conditions. I therefore include borrower state-by-year fixed effects and lender fixed effects. The borrower state-by-year fixed effects compare lenders operating in the same borrower state and year, while lender fixed effects absorb permanent differences across institutions. The identifying comparison is therefore between Republican-HQ and non-Republican-HQ lenders operating in the same local mortgage market at the same time, net of permanent lender differences.

Before examining heterogeneity by lender political environment, I assess how these additional fixed effects affect the overall denial estimate. This distinction is important because state-by-year fixed effects absorb state-specific time-varying shocks, whereas lender fixed effects additionally condition on the lender receiving the application. Since lender choice may itself respond to the post-Dobbs environment, specifications with lender fixed effects answer a different question from the baseline reduced-form specification.

The specification ladder shows that replacing state and year fixed effects with state-by-year fixed effects has almost no effect on the estimated coefficient: the Female $\times$ Ban state $\times$ Post estimate changes from 0.0090 to 0.0088. Thus, the baseline result is not explained by state-specific contemporaneous mortgage-market shocks. Adding lender fixed effects, however, reduces the estimate to 0.0046 with state and year fixed effects and to 0.0044 with state-by-year fixed effects. The within-lender estimate therefore remains positive and statistically significant, but is approximately one-half of the baseline reduced-form estimate.

I interpret these specifications as answering complementary questions. The baseline estimate captures the total change in the relative probability that a female mortgage application is denied in ban states after the policy change. This total effect may include both changes in screening within lenders and changes in the lenders to which female applicants apply or from which they ultimately receive decisions. By contrast, lender

Table 1.11: Sensitivity of the Denial Effect to State-by-Year and Lender Fixed Effects

	(1)	(2)	(3)	(4)
Female × Ban state × Post	0.0090***	0.0088***	0.0046**	0.0044**
	(0.0018)	(0.0016)	(0.0019)	(0.0017)
State × year FE	No	Yes	No	Yes
Lender FE	No	No	Yes	Yes
Observations	15066875	15066875	15066767	15066767

Standard errors in parentheses

Notes: All columns report borrower-level triple-difference estimates of the effect of abortion restrictions on the female–male mortgage denial gap. Column (1) includes state and year fixed effects. Column (2) replaces these with state-by-year fixed effects. Column (3) adds lender fixed effects to the state and year fixed effects, and Column (4) includes both state-by-year and lender fixed effects. All specifications include log income, log loan amount, a conforming-loan indicator, LTV and DTI bins, and controls for applicant race, ethnicity, and age. Columns (1)–(2) contain 108 additional observations that are singleton lender observations automatically dropped when lender fixed effects are included. Standard errors in Columns (1)–(2) are clustered at the state level; Columns (3)–(4) use two-way clustering by state and lender.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

fixed effects hold constant permanent differences across lenders and therefore isolate the component arising within lenders. Because lender choice is potentially itself affected by the policy environment, I do not treat the lender-fixed-effects estimate as a replacement for the baseline estimate. Instead, the difference between the two specifications indicates that both within-lender responses and across-lender composition may contribute to the overall effect.

Having established how the more restrictive fixed-effects specification changes the overall denial estimate, I next use this specification to examine whether the within-lender response differs across lender political environments. Table 1.12 reports the corresponding heterogeneity estimates.

Columns (1) and (2) estimate the treatment effect separately for non-Republican-HQ and Republican-HQ lender environments. Among non-Republican-HQ lenders, the estimated increase in the female–male denial gap is 0.0025, or approximately 0.25 percentage points (standard error 0.0013). Among Republican-HQ lenders, the corresponding estimate is 0.0049, or approximately 0.49 percentage points (standard error 0.0023). Thus, the Republican-HQ point estimate is larger, although the split-sample estimates alone do not provide a formal test of whether the two effects differ.

Column (3) provides that formal test using the pooled heterogeneous specification.

Table 1.12: Heterogeneity in Denial Effects by Lender Headquarters Political Environment

	Non-Republican HQ	Republican HQ	Pooled test
Female × Ban state × Post	0.0025*	0.0049**	0.0022*
	(0.0013)	(0.0023)	(0.0013)
Additional effect: Republican HQ			0.0031
			(0.0029)
Observations	7,072,475	7,925,352	14,997,839

Standard errors in parentheses

Notes: Columns (1) and (2) report split-sample borrower-level estimates for lenders associated with non-Republican and Republican headquarters or reporting-office states, respectively. Column (3) reports the pooled specification formally testing whether the post-ban female denial effect differs across lender groups. All specifications include borrower state-by-year and lender fixed effects, log income, log loan amount, a conforming-loan indicator, LTV and DTI bins, and controls for applicant race, ethnicity, and age. Standard errors are two-way clustered by borrower state and lender. The pooled specification includes the estimable lower-order interactions required for the heterogeneous triple-difference model.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The estimated Female × Ban state × Post coefficient for the reference lender environment is 0.0022 (standard error 0.0013), while the estimated additional effect for Republican-HQ lenders is 0.0031 (standard error 0.0029). The latter coefficient is not statistically significant. I therefore cannot reject equality of the post-ban female denial effects across the two lender environments. The larger Republican-HQ split-sample estimate should consequently be interpreted as suggestive heterogeneity rather than as evidence of a statistically significant difference across lender groups.

As an additional robustness check, I examine whether this point-estimate difference reflects differences in lender scale. I measure lender size using the log number of HMDA applications received by each lender during the pre-Dobbs period, 2018–2021, and allow the Female × Ban state × Post effect to vary with this measure. Appendix Table C2 reports the results. In the corrected pooled heterogeneity specification, the estimated additional Republican-HQ effect is 0.0031 (standard error 0.0029). After allowing the treatment effect to vary with lender size, the coefficient changes only modestly to 0.0026 (standard error 0.0031). The interaction between the treatment effect and lender size is -0.0003 (standard error 0.0005). Thus, there is little indication that lender scale accounts for the Republican-HQ point-estimate differential. Both the Republican-HQ differential and the lender-size interaction are, however, imprecisely estimated.

More generally, the headquarters measure remains an imperfect proxy for lender institutional environment. Although the lender-size exercise addresses one observable dimension of institutional scale, headquarters or reporting-office location may also be correlated with business model, portfolio composition, geographic specialization, and other institutional characteristics that are not fully observed in HMDA. I therefore interpret the Republican-HQ results as suggestive heterogeneity across lender environments, rather than as evidence that lender political culture itself causes the difference in responses.

Taken together with the age heterogeneity results, the lender evidence weighs against a narrow mechanism operating solely through lenders' assessment of an individual applicant's pregnancy risk. The denial effect does not attenuate with applicant age, and a substantial component of the overall effect remains within lenders. The Republican-HQ point estimates also suggest that responses may differ across lender environments, although the formal heterogeneity estimate is not statistically significant. These patterns do not uniquely identify gender salience or discriminatory screening: they may also reflect differences in perceived repayment risk, unobserved borrower creditworthiness, or other lender-specific underwriting responses. I therefore interpret the evidence more cautiously as showing an important within-lender component of the post-Dobbs increase in the female–male denial gap, while the precise underlying mechanism remains only partially identified.

1.6.7 Observable Applicant Composition and Selection

A key concern in interpreting the main results is that the observed increase in denial rates for female applicants could be driven by changes in the composition of applicants rather than changes in lender behavior. For example, if abortion bans lead more financially constrained women to apply for mortgages, the observed increase in denial rates could reflect a shift in applicant characteristics rather than a tightening of credit supply.

To address this concern, I examine whether observable borrower characteristics change differentially for female applicants in treated states following the policy change. Specifically, I estimate triple-difference models using applicant income, loan-to-value (LTV) ratios, and debt-to-income (DTI) ratios as outcomes. Table 1.13 reports the results.

Across all specifications, I find no economically meaningful or statistically significant changes in these characteristics. Female applicants in treated states do not exhibit lower

Table 1.13: Composition Stress Tests: Borrower Characteristics

	Log income	LTV	DTI
Female × Ban state × Post	-0.0046	-0.0088	0.0431
	(0.0051)	(0.1155)	(0.0490)
Observations	15125238	15126298	15125238

Standard errors in parentheses

Notes: Each column reports borrower-level triple-difference estimates with state and year fixed effects. The reported coefficient is **Female** × **Ban state** × **Post**, which captures the differential post-ban change in borrower characteristics for female applicants relative to male applicants in treated states. The dependent variables are log income in column (1), loan-to-value ratio in column (2), and debt-to-income ratio in column (3). All specifications are estimated on the decisioned sample and include log loan amount, a conforming loan indicator, LTV (for DTI specification) and DTI (for LTV specification) controls as applicable, and demographic controls for race, ethnicity, and age. Standard errors are clustered at the state level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

incomes, higher leverage, or higher debt burdens relative to male applicants after the introduction of abortion bans.

In addition, I check whether the number of women applying for a mortgage relative to men changed after the introduction of abortion bans. This can also be an informative exercise to determine whether the demand side has changed. Table 1.14 provides the estimate for the regression.

Table 1.14: Effect of Abortion Bans on Female Share of Mortgage Applicants

	No controls	Controls
Ban state × Post	0.0006	0.0001
	(0.0031)	(0.0031)
Observations	18954977	15125238

Standard errors in parentheses

Notes: Each column reports borrower-level difference-in-differences estimates where the dependent variable is an indicator equal to one for female applicants. The coefficient **Ban state** × **Post** captures the change in the share of female applicants in treated states after the implementation of abortion bans relative to control states. Column (2) includes borrower-level controls: log income, log loan amount, conforming loan indicator, LTV and DTI bins, and demographic controls for race, ethnicity, and age. All specifications include state and year fixed effects. Standard errors are clustered at the state level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

I find no evidence of a differential increase in the overall share of female mortgage applicants following the introduction of abortion bans.

I also examine whether the results are sensitive to the treatment of applications with

co-applicants. Appendix Table C3 shows that restricting the sample to applications without a co-applicant yields an estimated female × ban-state × post coefficient of 0.0090 for mortgage denial (standard error 0.0020), essentially identical to the baseline estimate. In the full sample, additionally controlling for co-applicant sex yields an estimate of 0.0084 (standard error 0.0018). These results show that the main denial effect remains when joint applications are excluded and when co-applicant sex is explicitly controlled, suggesting that differences in joint-application structure or the designation of the primary applicant do not account for the estimated effect.

I further extend the observable-composition analysis to applicant race, ethnicity, finer age categories, and joint-applicant status. Appendix Table C4 reports separate triple-difference estimates using indicators for these characteristics as outcomes. Most demographic changes are small. The estimated changes in the White, Asian, and Hispanic shares are statistically insignificant, while the Black share increases by about 1.8 percentage points and is marginally significant. Across the finer age categories, the estimates do not display a systematic shift toward either younger or older applicants, although the shares of applicants below age 25 and aged 55–64 decline modestly. The clearest observable compositional change is in joint applications: the probability that an application includes a co-applicant increases by about 4.6 percentage points for female applicants in treated states after the policy change.

Taken together, these exercises do not imply that the observable composition of female applicants is completely unchanged after abortion restrictions become binding. Rather, they show that the main denial result is robust to the observable compositional changes identified in the data. The preferred borrower-level specification already conditions on applicant race, ethnicity, and age, while the denial effect remains similar when co-applicant characteristics are included and when the sample is restricted to sole applicants. Together with the absence of economically meaningful changes in applicant income, LTV, DTI, and the overall female share of applicants, this evidence makes observable applicant composition an unlikely explanation for the main denial result. However, these tests address selection on observables only. HMDA does not allow me to rule out changes in unobserved applicant characteristics, so selection on unobservables remains a limitation of the analysis.

1.6.8 Summary of Mechanism Evidence

The mechanism evidence points to an approval-margin interpretation. Denials increase and originations fall for female applicants relative to male applicants, while conditional interest rates do not rise. The effect is stronger among low-income applicants and remains robust to the observable applicant-composition changes identified in the data. The age results indicate that the response is not confined to women in prime reproductive ages, while the Republican-HQ results show heterogeneity across lender institutional environments. The lender fixed-effects analysis further shows that approximately half of the baseline denial effect remains within lenders, while the attenuation relative to the baseline estimate points to an additional across-lender component. However, HMDA does not observe credit scores or all other dimensions of borrower credit risk. I therefore cannot distinguish whether the within-lender increase in denials reflects changes in lender screening of female applicants or risk-based underwriting responses to changes in unobserved creditworthiness. The mechanism evidence is consequently suggestive of an important lender-side response, but does not uniquely identify gender salience or discriminatory screening.

For a compact table linking each test to the mechanism it addresses, see Appendix C. I keep the main text shorter here because the individual pieces of evidence have already been discussed above.

1.7 Robustness and Placebo

1.7.1 Robustness Check: Synthetic Difference-in-Differences

I next examine whether the main results are robust to using a state-year design based on female–male outcome gaps. The synthetic difference-in-differences estimates use the state-year gap panel described above and compare ban states to a reweighted combination of protected states. I treat this exercise as a robustness check rather than as a separate core empirical design, because the borrower-level specification remains the main analysis and uses substantially more information.

I place greater weight on the specifications without contemporaneous macroeconomic controls because these variables may themselves lie on causal pathways from abortion

Table 1.15: Synthetic Difference-in-Differences Estimates of Abortion Bans on Gender Gaps

	Origination rate gap		Denial rate gap		Interest-rate gap	
	No controls	Controls	No controls	Controls	No controls	Controls
Ban state × Post	-0.0073***	-0.0068**	0.0073***	0.0068***	-0.0108	-0.0105
	(0.0027)	(0.0027)	(0.0023)	(0.0023)	(0.0128)	(0.0138)
Observations	266	266	266	266	266	266

Standard errors in parentheses

Notes: Each column reports a synthetic difference-in-differences estimate using state-year female–male gaps in mortgage outcomes. The coefficient reports the effect of abortion-ban exposure on the corresponding female–male gap. The origination outcome is the female–male gap in origination rates, not log originations. Denial and origination gaps are calculated on the decisioned sample, while interest-rate gaps are calculated on originated loans only. Columns labeled *No controls* include no additional macro controls. Columns labeled *Controls* include the state unemployment rate, log house price index, log personal income per capita, and log population. Inference uses placebo standard errors with 500 replications.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

restrictions to mortgage-market conditions. The controlled specifications are therefore interpreted as sensitivity checks rather than as preferred estimates. The similarity of the estimates across the two specifications is reassuring: for denials, for example, the estimated effect changes only from 0.73 to 0.68 percentage points when the macroeconomic controls are included.

Table 1.15 reports staggered SDID estimates for denial-rate gaps, origination-rate gaps, and interest-rate gaps. The clearest result is for denials. Abortion bans increase the female–male denial-rate gap by about 0.73 percentage points without macro controls and 0.68 percentage points with macro controls. Both estimates are statistically significant and close to the borrower-level triple-difference estimate. This supports the central finding that abortion bans widened gender disparities in mortgage approval.

The origination-rate estimates are also negative and statistically significant. This is consistent with the borrower-level evidence that female applicants experience a relative decline in realized mortgage credit after abortion bans. The interest-rate estimates remain negative, but they are small and not statistically significant. This might be due to the higher volatility of this variable, as evident in the Event-study exercise in Section 1.6.2. In addition, the SDID analysis uses a much smaller state-year panel, while the borrower-level interest-rate regressions use millions of originated loans.

The denial and origination SDID estimates are nearly symmetric in magnitude, but

this symmetry is not imposed mechanically. The denial-rate and origination-rate gaps are constructed separately from the underlying numbers of denied and originated applications. The two outcomes are close to complements because denials and originations account for the large majority of the decisioned sample; the remaining category consists primarily of applications approved by the lender but not accepted by the applicant. Consequently, changes in the female–male denial and origination gaps can be approximately equal and opposite without being algebraically constrained to be so.

Overall, the staggered SDID exercise supports the main borrower-level findings. The denial result remains positive, statistically significant, and similar in magnitude when the analysis is aggregated to state-year gender gaps and protected states are reweighted to match treated states' pre-treatment trajectories. The origination result also aligns with the borrower-level evidence. The interest-rate result has the expected sign but is less precise, so I treat it as supportive rather than as independent confirmation of the pricing result.

1.7.2 Robustness to Excluding Texas

Texas plays a special role in the empirical design because it is the early-treated state. While most ban states enter treatment in 2023, Texas is assigned to the post-treatment period beginning in 2022. This timing is substantively justified, but it raises a natural robustness concern: the main estimates could partly reflect Texas-specific mortgage-market dynamics rather than the broader effect of abortion bans across treated states. This concern is empirically relevant because Texas experienced a pronounced housing-market adjustment during the same period in which it is coded as the early-treated state. Texas Real Estate Research Center reports from 2022 describe the state housing market as cooling sharply as rising mortgage rates reduced affordability. In September 2022, Texas home sales fell by more than 15 percent year-over-year, and by December 2022 nearly 30 percent of total home sales had disappeared relative to December 2021 (Texas Real Estate Research Center, 2022b,a). These dynamics could affect mortgage applications, originations, and denials independently of abortion policy, making it important to verify that the main gender-specific result is not driven by Texas alone.

To address this concern, I re-estimate the main borrower-level triple-difference specifications after excluding Texas from the sample. The treatment group in this exercise

consists only of the remaining ban states, which are assigned to treatment beginning in 2023. The specification is otherwise identical to the main analysis: denial and origination regressions are estimated on the decisioned sample, interest-rate regressions are estimated on originated loans, and the same sets of controls are included in columns A and B.

Table 1.16: Triple-difference Estimates Excluding Texas

	Originations		Denials		Interest Rate	
	A	B	A	B	A	B
Female	0.0015**	0.0038***	-0.0010	-0.0031***	0.0194***	0.0187***
	(0.0007)	(0.0008)	(0.0007)	(0.0008)	(0.0024)	(0.0027)
Ban state × Post	-0.0045	-0.0043	0.0039	0.0034	-0.0100	-0.0118
	(0.0036)	(0.0033)	(0.0032)	(0.0029)	(0.0353)	(0.0349)
Female × Ban × Post	-0.0124***	-0.0107***	0.0124***	0.0109***	-0.0354***	-0.0353***
	(0.0039)	(0.0030)	(0.0038)	(0.0030)	(0.0051)	(0.0054)
Observations	14545465	12932245	14545465	12932245	12912648	11517245

Standard errors in parentheses

Notes: Each column reports a borrower-level triple-difference estimate excluding Texas from the estimation sample. All specifications include state and year fixed effects. The coefficient of main interest is **Female × Ban state × Post**, which captures the differential post-ban change in outcomes for female applicants relative to male applicants in treated states. Denial and origination regressions are estimated on the decisioned sample, while interest-rate regressions are estimated on originated loans only. Columns labeled *A* include log income, log loan amount, a conforming-loan indicator, and LTV and DTI bins. Columns labeled *B* additionally control for applicant race, ethnicity, and age. Standard errors are clustered at the state level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.16 shows that the main gender-differential results are robust to excluding Texas. The coefficient on *Female* × *Ban state* × *Post* remains negative and statistically significant for originations, positive and statistically significant for denials, and negative and statistically significant for interest rates. In the preferred specifications with demographic controls, abortion bans are associated with a 1.07 percentage point decline in female originations relative to male originations and a 1.09 percentage point increase in female denials relative to male denials. The interest-rate coefficient is also negative, at approximately 3.5 basis points, consistent with the approval-margin selection interpretation.

The results are, if anything, slightly larger than in the baseline table. In the main specification, the estimated increase in the female–male denial gap is approximately 0.9 percentage points. Excluding Texas, the corresponding estimate is approximately 1.1 percentage points. This indicates that the main denial result is not driven by Texas

alone. Instead, the widening of the female–male denial gap is also present among the later-treated ban states.

An important difference from the baseline table is that the coefficient on *Ban state* × *Post* becomes statistically insignificant after Texas is excluded. Since this coefficient captures the post-treatment change for the male reference group, the result suggests that the significant overall tightening for male applicants in the baseline specification was partly driven by Texas. This is consistent with the concern that Texas may have experienced distinctive mortgage-market dynamics around 2022. Once Texas is removed, there is no statistically significant evidence that male applicants in ban states experienced a broad post-treatment change relative to male applicants in protected states.

This pattern helps clarify the interpretation of the main results. The baseline estimates showed both a general post-treatment tightening in ban states and an additional deterioration for female applicants. The no-Texas robustness check suggests that the general tightening component was more sensitive to the inclusion of Texas, while the gender-specific component is not. The coefficient of interest, *Female* × *Ban state* × *Post*, remains large, statistically significant, and economically meaningful after excluding Texas. Thus, the evidence that abortion bans widened gender disparities in mortgage approval does not depend on the early-treated Texas case.

Overall, the no-Texas robustness check strengthens the central conclusion of the paper. It shows that the increase in the female–male denial gap is not an artifact of Texas-specific housing-market conditions or early treatment timing. Rather, the result persists among the broader set of ban states treated after *Dobbs*, supporting the interpretation that abortion restrictions worsened women's relative access to mortgage credit.

1.7.3 Robustness to Gender-Specific State and Year Fixed Effects

I examine whether the borrower-level results are robust to allowing the female–male gap in mortgage outcomes to vary more flexibly across states and over time. Specifically, I augment the preferred specification with state-by-female and year-by-female fixed effects. The former absorb persistent state-specific differences in gender gaps, while the latter absorb national changes in gender gaps that are common across states in a given year. This specification, therefore, addresses the possibility that treated and protected states differ systematically in their underlying gender gaps or that national mortgage-market

developments after 2022 affected female and male applicants differently.

Table 1.17: Robustness to Gender-Specific State and Year Fixed Effects

	Origination		Denial		Interest rate	
	Baseline	Gender FE	Baseline	Gender FE	Baseline	Gender FE
Female × Ban state × Post	-0.0084***	-0.0045**	0.0090***	0.0050***	-0.0266***	-0.0156**
	(0.0019)	(0.0018)	(0.0018)	(0.0017)	(0.0050)	(0.0064)
Observations	15125238	15125238	15125238	15125238	13422392	13422392

Standard errors in parentheses

Notes: The table reports borrower-level triple-difference estimates. Baseline columns reproduce the preferred specifications from Table 1.8 and include state and year fixed effects. Gender FE columns additionally allow the female–male outcome gap to vary flexibly across states and years by absorbing state-by-female and year-by-female fixed effects. All specifications include log income, log loan amount, a conforming-loan indicator, LTV and DTI bins, and controls for applicant race, ethnicity, and age. Denial and origination models use the decisioned sample; interest-rate models use originated loans with valid interest rates. Standard errors are clustered at the state level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.17 shows that the estimated effects attenuate under the more flexible specification but remain statistically significant and retain the same signs as in the preferred specification. The denial effect declines from 0.0090 to 0.0050, while the origination effect changes from -0.0084 to -0.0045. The conditional interest-rate estimate declines in magnitude from -0.0266 to -0.0156. Thus, allowing for persistent state-specific gender gaps and national year-specific changes in gender gaps reduces the estimated magnitudes, but does not overturn the qualitative pattern of the main results.

1.7.4 Randomization Inference

As an additional robustness check, I conduct a randomization-inference exercise based on placebo treatment assignments. The purpose of this exercise is to evaluate whether the estimated gendered DiD effects are unusually large relative to the effects that would be obtained if treatment status were randomly assigned across states. This provides a complementary inference procedure that does not rely only on conventional cluster-robust standard errors.

Given the computational scale of the borrower-level HMDA data, which contains almost 19 million mortgage applications after the baseline restrictions, conducting hundreds of placebo re-estimations at the individual-application level would be computationally

burdensome. I therefore implement the randomization-inference exercise using state-year female–male gaps in mortgage outcomes. This aggregation is also conceptually appropriate because abortion policy varies at the state-year level, not at the individual-applicant level. The state-year gap specification corresponds directly to the main gender-gap estimand in the borrower-level triple-difference model: whether the female–male gap changes differentially in ban states after treatment.

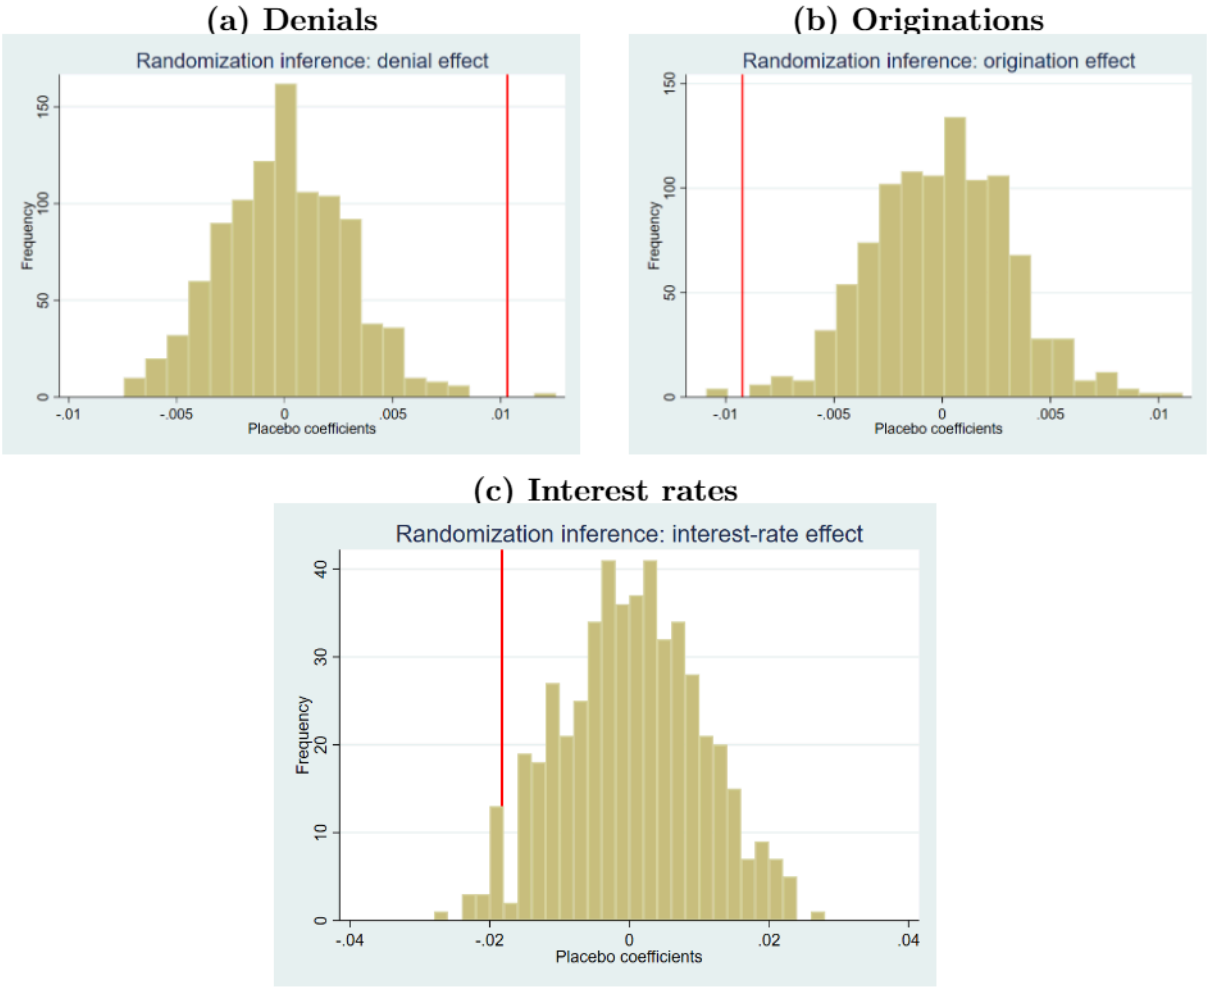
For each state and year, I construct the female–male gap in denials, originations, and interest rates. I then repeatedly assign placebo treatment status across states, holding fixed both the number of treated states and the staggered timing structure of treatment adoption. For each placebo assignment, I re-estimate the gap DiD specification using the placebo treatment indicator and store the placebo coefficient. This produces a distribution of treatment effects that would be observed under random assignment of treatment status.

Figure 1.3 plots the distribution of placebo estimates across 500 random assignments. The vertical line in each panel marks the actual estimated gap DiD effect using the true abortion-ban classification. The logic of the exercise is as follows. If the observed estimate is similar to estimates generated by random treatment assignments, the result could plausibly arise by chance. If, instead, the observed estimate lies in the tail of the placebo distribution, this suggests that the estimated effect is unusual relative to random assignment and provides additional support for the empirical design.

Panel (a) shows the placebo distribution for denial rates. The placebo estimates are centered close to zero, while the actual estimate lies far in the right tail of the distribution. This indicates that the observed increase in the female–male denial gap is unlikely to be generated by a random assignment of treatment status across states. This is important because the denial result is the central finding of the paper and the outcome most directly tied to the lender approval margin.

Panel (b) reports the corresponding exercise for originations. The actual estimate lies in the left tail of the placebo distribution, consistent with the main borrower-level finding that female applicants experience a relative decline in originations after abortion bans. The evidence is somewhat less visually sharp than for denials, which is expected because originations combine lender approval, borrower take-up, and broader mortgage-market activity.

Figure 1.3: Randomization inference: placebo distributions for state-year gap DiD estimates



Notes: The figure reports randomization-inference distributions for DiD estimates using state-year female–male gaps in mortgage outcomes. In each iteration, placebo treatment status is randomly assigned across states while holding fixed the number of treated states and the staggered treatment-timing structure. The DiD model is then re-estimated using the placebo assignment, and the placebo coefficient is stored. The vertical line indicates the estimate obtained using the true abortion-ban classification. Panel (a) corresponds to the female–male denial-rate gap, panel (b) to the female–male origination-rate gap, and panel (c) to the female–male interest-rate gap. Denial and origination-rate gaps are calculated on the decisioned sample, while interest-rate gaps are calculated on originated loans only. The exercise evaluates whether the observed DiD estimates are unusually large in magnitude relative to the distribution generated by random treatment assignment.

Panel (c) presents the placebo distribution for interest rates. The actual estimate is also located away from the center of the placebo distribution, although the distribution is wider than for denials. This is consistent with the interpretation that interest-rate estimates are noisier because they are observed only for originated loans and therefore depend on selection into approval.

Overall, the randomization-inference results reinforce the main conclusion. The denial effect, in particular, is difficult to explain as a chance result generated by arbitrary assignment of treatment across states. The placebo evidence therefore supports the interpretation that abortion bans widened the female–male denial gap, while the originations and interest-rate placebo exercises provide additional, though noisier, evidence consistent with the approval-margin mechanism.

1.8 Conclusion

This paper asks whether post-*Dobbs* abortion restrictions affected gender gaps in mortgage credit. Using HMDA data from 2018–2024, I find that abortion bans widened these gaps, with the strongest evidence appearing at the approval margin. Female applicants in abortion-ban states became more likely to be denied mortgage relative to comparable male applicants after abortion restrictions became binding. Originations move in the same direction, while economically small interest-rate results are consistent with selection among approved borrowers.

Several pieces of evidence support the findings. Treated and control states exhibit similar pre-treatment trends in gender gaps, and event-study estimates show no evidence of systematic differential pre-trends before policy adoption. The no-Texas robustness check shows that the gender-specific mortgage denial effect is not driven by the early-treated Texas case. Randomization-inference exercises also indicate that the denial effect is unusual relative to placebo treatment assignments. Finally, staggered synthetic difference-in-differences estimates based on state-year gender gaps provide support for the mortgage denial and origination-rate results, although the interest-rate estimates are less precise.

The heterogeneity analysis provides additional evidence on possible mechanisms. The mortgage denial effect is substantially larger among low-income applicants, consistent

with a credit-constraint channel. Across finer age categories, the estimates do not display a monotonic decline with age and remain positive even among applicants aged 65 and above. This weighs against a narrow mechanism operating solely through an individual applicant's fertility exposure, although it does not uniquely identify the alternative mechanism. The lender analysis further shows that approximately half of the baseline denial effect remains within lenders. Point estimates are larger among Republican-HQ lender environments, but the formal heterogeneity estimate is not statistically significant. I therefore interpret these patterns as suggestive evidence of heterogeneous lender responses rather than as direct evidence of gender salience, lender ideology, or discriminatory screening.

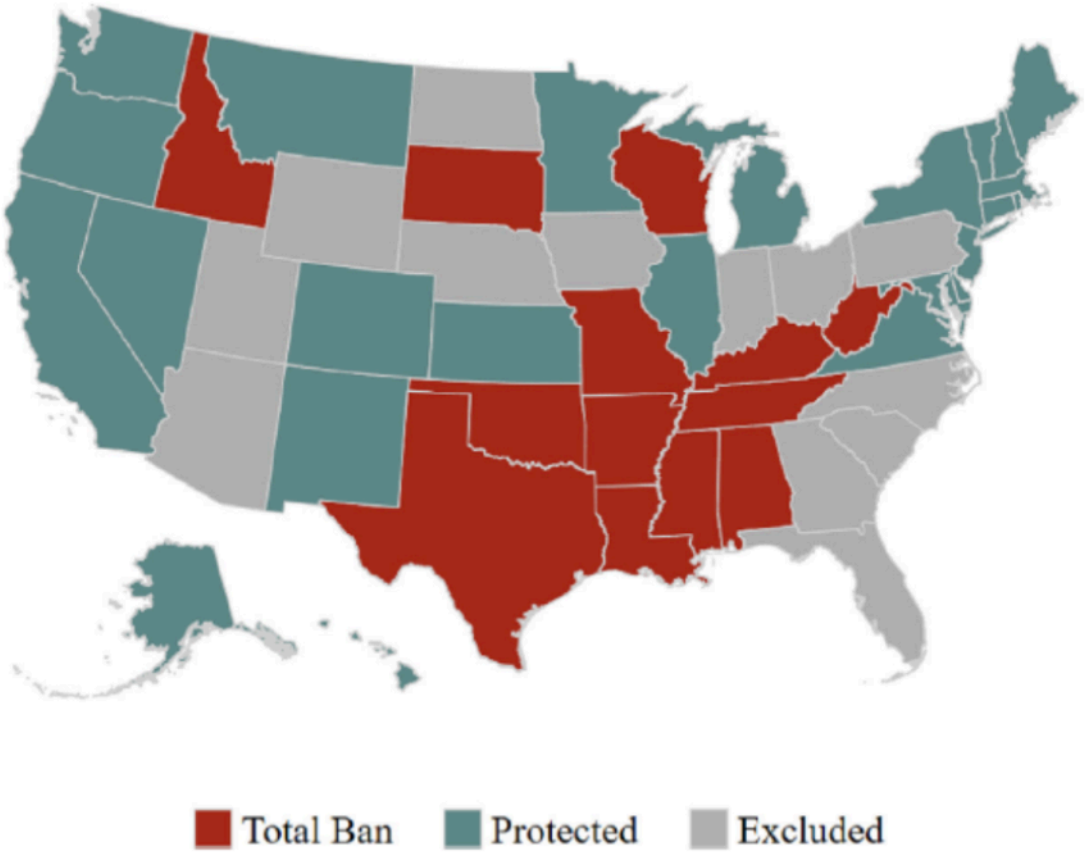
These findings contribute to the economics of abortion policy and to the study of mortgage credit. The results show that abortion restrictions can affect not only fertility, health, and labor-market outcomes, but also women's access to long-term financial opportunities. They also add a new perspective to the mortgage-discrimination literature by showing that gender gaps in mortgage approval outcomes can widen in response to a major reproductive-policy shock.

The paper also has limitations. HMDA is rich for studying mortgage applications, but it does not observe all dimensions of borrower risk, household expectations, or lender underwriting models. The reduced-form design therefore identifies how gender gaps changed after abortion bans, but it cannot fully separate shifts in lender beliefs, household behavior, and unobserved financial risk. Future work should connect these mortgage-market responses to downstream outcomes such as homeownership, refinancing, delinquency, mobility, and longer-run household wealth.

Overall, the evidence indicates that the post-*Dobbs* abortion bans had consequences beyond reproductive healthcare itself. The findings show that abortion restrictions can reinforce gender inequality not only in the labor market and within the household, but also in the allocation of credit.

Appendix A. State Classification

Figure A1: State Classification Used in the Analysis

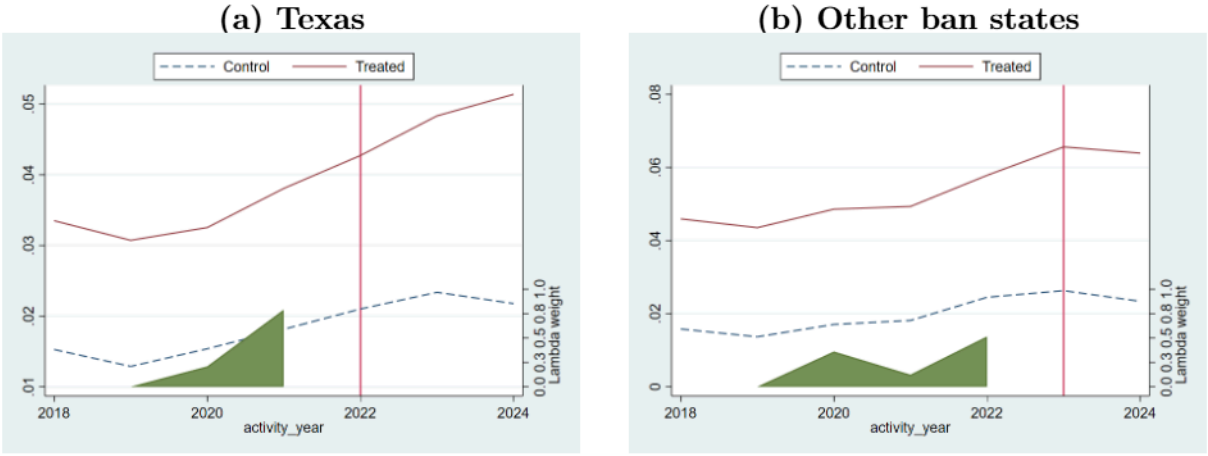


Notes: The classification follows Dench et al. (2024). Total-ban states are those that enacted a total abortion-ban by the end of 2022. Protected states either kept abortion access in place or expanded abortion rights. Excluded states include states that enacted bans too late to be effective for fertility in my analysis, states with enforced gestational-age bans, and states with legally uncertain or intermediate policy environments.

Appendix B. Synthetic Difference-in-Differences Graphical Diagnostics

Because treatment timing is staggered, the SDID graphical diagnostics are reported separately by treatment cohort. Texas is treated beginning in 2022, while the remaining ban states are treated beginning in 2023. For each outcome, I therefore present two SDID plots: one for Texas and one for the remaining ban states. In each panel, the red line traces the observed path of the treated unit or treated cohort, while the blue dashed line traces the corresponding synthetic comparison path constructed by the SDID estimator. The green shaded areas represent the pre-treatment time weights chosen by the SDID procedure. These weights indicate which pre-treatment years receive greater weight in matching the treated and synthetic-control trajectories; the synthetic comparison group itself is constructed using unit weights over protected donor states.

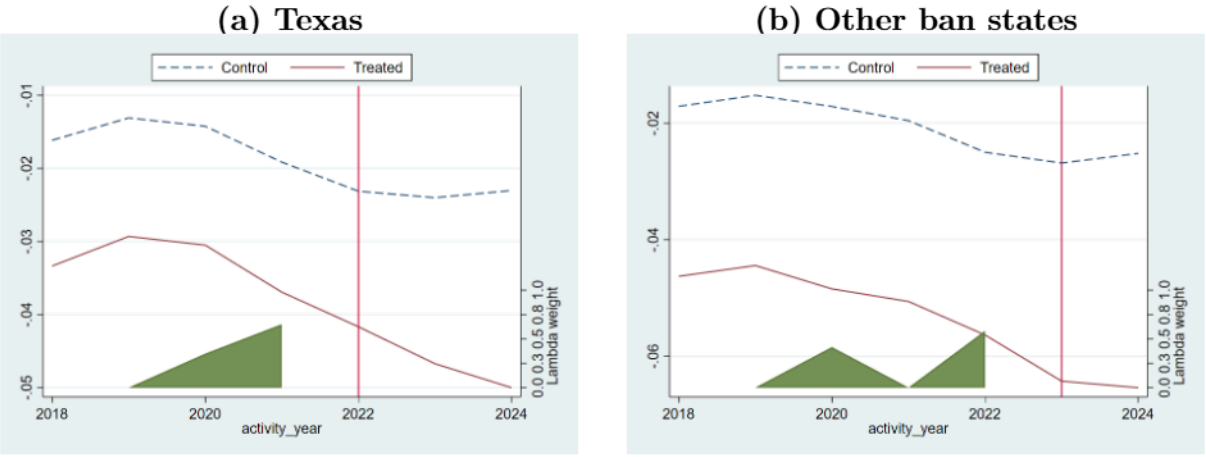
Figure B1: SDID graphical diagnostics: female–male denial-rate gap



Notes: The figure reports SDID graphical diagnostics for the female–male denial-rate gap. Panel (a) shows the Texas cohort, which is treated beginning in 2022. Panel (b) shows the remaining ban states, which are treated beginning in 2023. The outcome is the difference between the average denial rate for female applicants and the average denial rate for male applicants within a state-year. Denial rates are calculated on the decisioned sample. In each panel, the black line is the observed treated path and the red line is the SDID synthetic comparison path. The green shaded areas indicate the pre-treatment time weights selected by the SDID procedure.

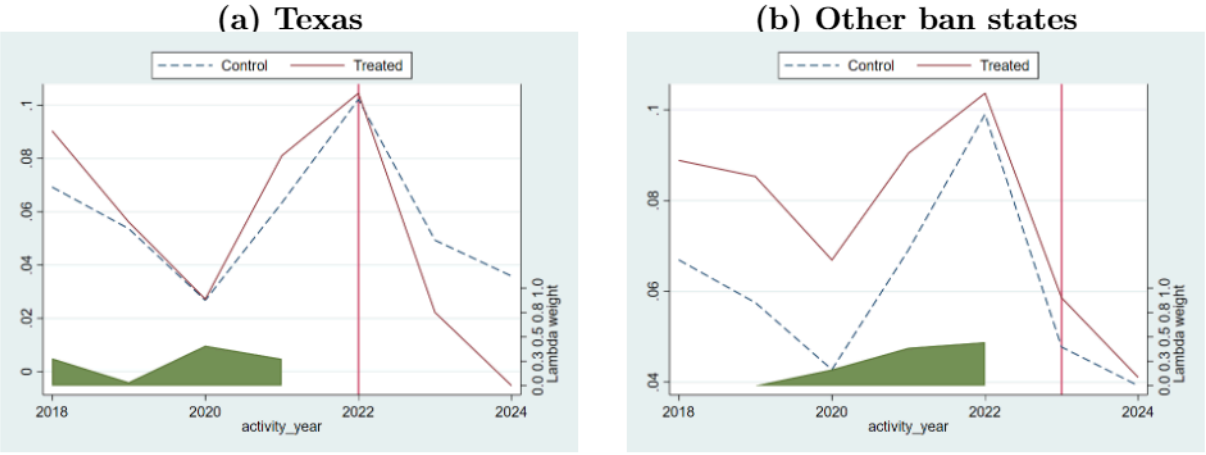
Taken together, these cohort-specific SDID plots provide a visual diagnostic for the robustness analysis reported in Table 1.15. The denial-gap figures are the most informative, as they correspond to the paper's primary outcome and provide the clearest visual support for a post-treatment widening of the female–male denial gap. The origination-rate plots are also useful because the revised SDID outcome is an origination-rate gap, matching the borrower-level estimand more closely than the earlier log-originations specification. The

Figure B2: SDID graphical diagnostics: female–male origination-rate gap



Notes: The figure reports SDID graphical diagnostics for the female–male origination-rate gap. Panel (a) shows the Texas cohort, which is treated beginning in 2022. Panel (b) shows the remaining ban states, which are treated beginning in 2023. The outcome is the difference between the average origination rate for female applicants and the average origination rate for male applicants within a state-year. Origination rates are calculated on the decisioned sample. In each panel, the black line is the observed treated path and the red line is the SDID synthetic comparison path. The green shaded areas indicate the pre-treatment time weights selected by the SDID procedure.

Figure B3: SDID graphical diagnostics: female–male interest-rate gap



Notes: The figure reports SDID graphical diagnostics for the female–male interest-rate gap. Panel (a) shows the Texas cohort, which is treated beginning in 2022. Panel (b) shows the remaining ban states, which are treated beginning in 2023. The outcome is the difference between the average interest rate for female applicants and the average interest rate for male applicants within a state-year. Interest rates are calculated only on originated loans. In each panel, the black line is the observed treated path and the red line is the SDID synthetic comparison path. The green shaded areas indicate the pre-treatment time weights selected by the SDID procedure.

interest-rate plots should be read more cautiously. interest rates are analyzed only among originated loans and are therefore sensitive to selection into approval; in the staggered SDID estimates, the point estimates remain negative but are not statistically significant.

Appendix C. Additional Mechanism and Robustness Tables

This appendix reports additional tables supporting the mechanism and robustness analyses. Table C1 summarizes how the main results, heterogeneity tests, and composition checks map into the mechanism discussion.

Table C1: Summary of Mechanism Tests

Mechanism / test	Empirical prediction	Result	Interpretation
Approval-margin selection	If bans tighten access through lender screening, denials should rise, originations should fall, and conditional interest rates need not increase.	Denials increase for female applicants relative to male applicants; originations fall; conditional interest rates decline.	The joint pattern is consistent with a response operating primarily at the approval margin. Because interest rates are analyzed only among originated loans, the negative conditional-rate estimate is selection-sensitive and does not by itself identify positive selection or a particular mechanism.
Credit-constraint channel	Effects should be larger among applicants closer to approval thresholds, especially low-income borrowers.	The denial effect is substantially larger and statistically significant among low-income applicants.	The results support a credit-constraint interpretation: abortion bans have the strongest effect on borrowers with less financial slack.
Narrow fertility-risk channel	If lenders respond mainly to individual pregnancy risk, effects should be concentrated among women in prime reproductive ages.	Positive and statistically significant denial effects are estimated in all six age groups, from below 25 through 65+, with no monotonic decline in the point estimates with age.	The age pattern weighs against a narrow mechanism based solely on an individual applicant's fertility exposure, but does not uniquely identify the alternative mechanism.
Lender institutional environment	If lenders differ in their response to the post-*Dobbs* environment, effects may vary by lender headquarters or reporting-office environment.	The Republican-HQ split-sample point estimate is larger, but the pooled interaction testing the difference across lender environments is not statistically significant.	The results provide suggestive evidence of heterogeneous lender responses, but do not establish a statistically significant difference across political environments.
Applicant composition	If the main effect is driven by changing applicant composition, treated female applicants should exhibit systematic post-treatment changes in observable characteristics associated with mortgage outcomes.	Income, LTV, DTI, and the overall female applicant share show little change. Additional checks identify some demographic shifts and a larger increase in co-applicant presence, but the main denial estimate remains similar when co-applicant characteristics are controlled for and when the sample is restricted to sole applicants.	Observable composition changes do not appear to account for the main denial result, although selection on unobserved characteristics cannot be ruled out.

Notes: The table summarizes the mechanism and heterogeneity tests used to interpret the borrower-level triple-difference results. The coefficient of primary interest is Female × Ban state × Post, which captures the differential post-ban change in outcomes for female applicants relative to male applicants in treated states.

As an additional robustness check for the lender-environment heterogeneity analysis, I examine whether the estimated difference between Republican-HQ and non-Republican-

HQ lenders is related to lender scale. I measure lender size using the number of HMDA mortgage applications received by each lender during the pre-Dobbs period, 2018–2021, and use the logarithm of this measure in the regression. Because lender size is measured before the policy change, it is not itself affected by the post-Dobbs treatment environment.

Table C2 compares the corrected pooled Republican-HQ heterogeneity specification with a specification that additionally allows the Female × Ban state × Post effect to vary with pre-period lender size. The estimated additional Republican-HQ effect changes only modestly after introducing this interaction, while the lender-size interaction is small and statistically insignificant. The exercise therefore provides little evidence that lender scale accounts for the Republican-HQ point-estimate differential. However, both the Republican-HQ differential and the lender-size interaction are imprecisely estimated, and the exercise does not rule out confounding by other lender characteristics such as business model, portfolio composition, or geographic specialization.

Table C2: Robustness of Republican-HQ Heterogeneity to Lender Size

	Baseline HQ heterogeneity	Allowing for lender size
Female × Ban state × Post	0.0022*	0.0041
	(0.0013)	(0.0031)
Additional effect: Republican HQ	0.0031	0.0026
	(0.0029)	(0.0031)
Treatment effect × lender size		-0.0003
		(0.0005)
Observations	14,997,839	14,997,839

Standard errors in parentheses

Notes: Both columns report pooled borrower-level heterogeneity specifications for mortgage denial. Column (1) estimates whether the post-ban female denial effect differs between Republican-HQ and non-Republican-HQ lenders. Column (2) additionally allows the treatment effect to vary with lender size, measured as the centered log number of HMDA applications received by the lender during 2018–2021. Lender size is centered at the mean log size across lenders. Both specifications include borrower state-by-year and lender fixed effects, log income, log loan amount, a conforming-loan indicator, LTV and DTI bins, and controls for applicant race, ethnicity, and age. Standard errors are two-way clustered by borrower state and lender.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C3 examines whether the main denial estimate is sensitive to the treatment of applications with co-applicants. The estimated effect is essentially unchanged when the sample is restricted to applications without a co-applicant and remains similar when co-applicant sex is additionally controlled for in the full sample.

Table C3: Robustness to Co-Applicant Status

	Baseline	Sole applicants	Co-applicant controls
Female × Ban state × Post	0.0090***	0.0090***	0.0084***
	(0.0018)	(0.0020)	(0.0018)
Observations	15125238	8250990	15125238

Standard errors in parentheses

Notes: The table reports borrower-level triple-difference estimates for mortgage denial. Column (1) reports the baseline specification. Column (2) restricts the sample to applications without a co-applicant. Column (3) uses the full sample and additionally controls for co-applicant sex, with applications without a co-applicant as the reference category. The coefficient **Female × Ban state × Post** captures the differential post-ban change in denial probability for female applicants relative to male applicants in treated states. All specifications include state and year fixed effects, log income, log loan amount, a conforming-loan indicator, LTV and DTI bins, and controls for applicant race, ethnicity, and age. Standard errors are clustered at the state level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C4 extends the applicant-composition analysis to additional demographic and application characteristics available in the analysis data. Each row reports a separate borrower-level triple-difference regression in which the dependent variable is an indicator for the corresponding applicant characteristic.

Table C4: Expanded Observable-Composition Checks

Outcome	Female × Ban state × Post	SE	Observations
White	-0.0074	0.0166	17,218,257
Black	0.0183*	0.0107	17,218,257
Asian	-0.0109	0.0088	17,218,257
Hispanic	0.0003	0.0021	16,784,790
Age < 25	-0.0047*	0.0027	18,951,056
Age 25–34	-0.0021	0.0037	18,951,056
Age 35–44	0.0082	0.0055	18,951,056
Age 45–54	0.0018	0.0017	18,951,056
Age 55–64	-0.0040***	0.0010	18,951,056
Age 65+	0.0009	0.0012	18,951,056
Co-applicant present	0.0462***	0.0066	18,954,977

Notes: Each row reports a separate borrower-level triple-difference regression. The dependent variable is an indicator for the applicant characteristic listed in the first column. The reported coefficient is Female × Ban state × Post. All specifications include state and year fixed effects, and standard errors are clustered at the state level. The estimation sample contains decisioned applications in abortion-ban and protected states. Differences in observation counts reflect missing values in the corresponding demographic variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Chapter 2

Access to Financial Resources and Environmental Migration of the Poor[1]

2.1 Introduction

Anthropogenic climate change is causing increases in average temperature and changes in precipitation patterns, rising sea levels and more frequent extreme weather events, including droughts, heatwaves, and floods (IPCC, 2013; Jones & O'Neill, 2016). Climate change affects all countries, but the impact will be felt more strongly by less developed ones (IPCC, 2014). Deterioration in the quality of life and hazardous environmental conditions are expected to lead to an increase in migration flows from affected areas. Since the 1980s, scholars and international organizations have predicted that environmental change will lead to a substantial increase in migration flows worldwide, with predictions ranging from 50 to 200 million additional climate migrants per year (Myers, 2002; Renaud et al., 2007; Stern, 2007; Biermann & Boas, 2010).

Despite an increasing number of studies, no scientific consensus currently exists as to what extent and under which conditions environmental factors influence migration. In particular, little is known about the role resources play in enabling or hindering mobility

[1]**Co-authored with Mariola Pytlikova (Aarhus University, CERGE-EI) and Roman Hoffmann (IIASA)**

under environmental stress. While some households are found to migrate in response to environmental hazards, others remain in place, potentially being trapped due to lack of resources and liquidity constraints (Black et al., 2013; Zickgraf & Perrin, 2016). While this pattern has been observed in previous work, little is known empirically about how access to financial resources influences the decision of a household to remain or migrate. Theoretically, additional resources can both increase or reduce environmental migration as they support *in-situ* adaptation. Likewise, the loss of resources due to an environmental shock can both trigger migration (climate-driver mechanism) and reduce out-migration (climate-inhibitor mechanism).

In this study, we investigate household migration behavior as a response to climatic shocks conditional on different possibilities to access financial resources. We focus specifically on the situation of resource-constrained poor households. Our study uses data from the Indonesian Family Life Survey (IFLS), which is an on-going longitudinal survey. The sample is representative of about 83% of the Indonesian population and contains over 30,000 individuals living in 13 of the 26 provinces in the country. The survey contains detailed information about the migration history of individual household members. Additionally, due to the longitudinal panel nature of the data, we can trace back the migration of entire households, thereby allowing us to estimate the impact of environmental conditions on mobility outcomes. As a country particularly affected by environmental changes and rapid onset disasters, Indonesia offers an interesting opportunity to study the moderating effect of resources on environmental migration.

To identify the impact of slow- and fast-onset environmental events on migration processes, we combine the IFLS data with additional secondary data sources. First, we use publicly available weather data from the Climate Research Unit (CRU) of the University of East Anglia to estimate the impact of changing environmental conditions on migration. In addition, we use the information on natural disasters from the IFLS to look into fast-onset events. To test for the moderating role of resources in influencing migration outcomes due to environmental stress, we analyze several cash-transfer programs that were implemented in Indonesia from 2005. In particular, we examine three government programs targeted to poor households: two unconditional and one conditional cash transfer programs. We exploit temporal variation in the receipt of the unconditional cash transfer, and temporal and spatial variation in the conditional one. We make use of this variation and the longitudinal data of the IFLS to estimate the moderating impact of

resource access on environmental migration.

Our results demonstrate heterogeneous responses to different kinds of climate shocks and the differential role of cash transfers in adaptation. We find that households that received cash transfers migrate less in response to current-year rainfall shocks. Under accumulated rainfall exposure, the positive cash-transfer interaction is concentrated in moves involving part or the whole household rather than in the overall migration indicator. For the temperature indicator, the cash-transfer interaction is positive but not statistically significant. Another finding suggests that households affected by natural disasters, such as floods, landslides, and droughts, tend to migrate less if they receive cash transfers, which is consistent with the evidence we find for short-term rainfall shocks. Finally, our results suggest that members of a household that received cash transfers are less likely to change jobs and enter the risky agricultural sector, and are more likely to leave an occupation in the agricultural sector as a response to rainfall and temperature shocks.

Thus, our research contributes to better understanding of adaptation strategies to climate shocks, particularly in the form of migration, depending on having access to financial resources. From the policy perspective, our study sheds light on how the choice of adaptation mechanisms of poor households changes conditional on having additional financial resources.

A useful feature of our setting is that we observe several margins of adjustment in the same data. We are not limited to one aggregate migration rate, but can distinguish between all moves, moves undertaken by one person, and moves undertaken by part or whole households. This distinction is important because these decisions may reflect different adaptation strategies. Sending one migrant can be a way of diversifying income while keeping the household in the original location. Moving with a larger part of the household is a stronger and more costly response. We therefore interpret the results not only as evidence on whether migration itself increases or decreases, but also as evidence about which type of mobility is more likely to be used when households face environmental stress and receive additional resources.

More broadly, the analysis emphasizes that cash transfers should not be interpreted only as poverty-alleviation instruments. In environmentally-exposed settings, they may also shape the set of adaptation strategies available to poor households. This distinction is important because a lower probability of migration after a shock does not necessarily

imply lower vulnerability; it may reflect an improved ability to cope locally. Conversely, a higher probability of migration after repeated exposure may indicate that resources help households choose a more costly but potentially more effective adjustment.

The remainder of this paper is organized as follows. Section 2.2 reviews the related literature. Section 2.3 presents the conceptual framework. Section 2.4 describes the institutional background and data. Section 2.5 presents the empirical and identification strategy. Section 2.6 reports the results. Section 2.7 presents robustness checks. Section 2.8 discusses policy implications, and Section 2.9 concludes.

2.2 Literature Review

This study builds on two strands of literature. The first investigates the determinants of migration and the climate-migration link. Although this research area is developing rapidly, there are still gaps that have not been addressed by existing studies. We contribute to one of them: how household's decision on whether to migrate as a response to a climate shock is related to their access to resources.

The second strand of the literature addresses how access to resources influences a household's decision to migrate. We provide an overview of several studies in this area, and identify that, although the findings differ, there is evidence that accessibility to financial resources can be a significant factor in the migration decision. Below, we firstly introduce the literature on the climate-migration nexus, followed by the literature on the influence of access to financial resources on migration.

There is a growing body of literature that investigates the influence of climate shocks on migration, part of which examines the effect on migration flows of fast-onset events, including floods, landslides, earthquakes, and hurricanes. The majority of existing studies agree that migration in response to fast-onset events tends to be over short distances and temporary (Cattaneo et al., 2019). Halliday (2006) exploits earthquakes in El Salvador as an exogenous shock for households and finds that migration to the United States increased at all wealth levels. However, the author finds that wealthier households tend to migrate more often, since they are less liquidity-constrained. Using municipality-level data in Mexico, Saldana-Zorrilla & Sandberg (2009) demonstrate that, from 1990-2000, municipalities that experienced disasters (droughts, floods, hurricanes) more often exhibit

higher migration rates than less exposed municipalities. Carvajal & Pereira (2010) use panel data of households in Nicaragua to study how people adapt to natural hazards. The study shows that hurricane Mitch, which hit Nicaragua in 1998, induced affected households to migrate more than those who were unaffected.

Another part of the literature examines the effect of slow-onset events on migration, including temperature rises, changes in rainfall patterns, and droughts. Here, the most relevant study for us is the paper by Thiede & Gray (2017), which shows how migration in Indonesia is influenced by climate shocks, in particular by rising temperatures and delayed monsoons. The authors find that increasing temperatures decrease out-migration from affected areas, while monsoon delay does the opposite. Their results vary significantly for different groups of people depending on gender, age and other characteristics, which suggests the heterogeneous use of migration as an adaptation strategy. Dillon et al. (2011) use panel data of Nigerian households to investigate how households respond to variability in temperature and the number of days suitable for growing crops. They show that males decide to migrate as a response to ex-ante and ex-post risk. Mueller et al. (2014) study the effect of heat stress on rural migration in Pakistan. They find that heat-stress consistently increases male migration due to its adverse effects on farm and non-farm income.

Cattaneo & Peri (2016) conduct a cross-country study to demonstrate that an increase in temperature decreases the migration flow from poor countries, consistent with the hypothesis of the presence of liquidity constraints. Nawrotzki & Bakhtsiyarava (2017) use census and climate data for Senegal and Burkina Faso to understand whether the population adapts to climate shocks *in situ* or decides to migrate. They show that excessive precipitation increases international out-migration from Senegal, while increasing temperature decreases migration out-flow from Burkina Faso. The authors conclude that adverse climate shocks can trap poor populations, thereby undermining their ability to migrate. Gröschl & Steinwachs (2017) show in a cross-country study that droughts increase migration for middle-income countries consistent with the presence of liquidity constraints for migration in poor countries. They also show that the rich countries do not exhibit increased out-migration because they have developed insurance schemes.

There is also a vast literature that considers the link between access to financial resources and migration. Stecklov et al. (2005) use data from the Progresa/Oportunidades conditional cash-transfer program in Mexico to study its effect on internal and external

out-migration. The authors exploit the random program assignment to the treatment and control communities during a pilot phase. Their results demonstrate that conditional cash transfers reduce out-migration to the United States but do not change internal migration patterns. Angelucci (2015) also uses data on the same conditional cash-transfer program to study labor migration to the United States. In contrast to the paper by Stecklov et al. (2005), the author finds that better access to financial resources for poor households in Mexico increases labor migration to the United States. Specifically, conditional cash transfers allow less-skilled migrants to move to the United States. The results of the two papers seem to be contradictory; however, the studies focus on two different types of migration: overall migration and labor migration.

Bryan et al. (2014) develop a model with risky migration and test it using data from a randomized trial. The authors randomly assign an $8.50 incentive to households in rural Bangladesh to temporarily out-migrate to the urban areas during the "lean season"[2]. They show that financial incentives induce more people to migrate and increase their consumption. In this case, access to financial resources can work as a complement to migration because, without it, households were liquidity constrained or too risk-averse.

Poggi (2019) uses a quasi-random experiment in Thailand to study credit availability and internal migration. The author studies the effect of the Thailand Village and Urban Community Fund Program (VFP). In 2001, the Thai government introduced this micro-finance initiative, which distributed one million Baht (US $24,000 in 2001 prices) to each village. A group of village members provided short-term credit to fellow villagers. The exploited variation comes from the equal amount of money distributed to all of the villages regardless of their size. Employing an instrumental variables approach, the author finds that borrowing reduces internal migration in the medium-term, but it does not affect migration when the policy is first introduced.

The reviewed literature on links between financial resources and migration provides evidence in support of two different hypotheses. The studies by Stecklov et al. (2005) and Poggi (2019) support the hypothesis that better access to financial resources improves living conditions at the place of residence, making it less profitable to out-migrate. As a result, they observe decreasing migration flows. The studies by Angelucci (2015) and Bryan et al. (2014) support another view that better access to financial resources induces

[2]the lean season is the time of year between planting and harvest when food runs out

out-migration due to relaxing the liquidity constraints of the population. In our study, we try to reconcile the research question in a slightly different framework that also incorporates climate shocks. We attempt to understand whether better access to financial resources induces people to migrate after a climate shock (climate-driver mechanism), or whether it results in the household deciding to adapt to the climate shock at the place of residence due to better adaptation capacity (climate-inhibitor mechanism).

In this sense, the existing evidence does not necessarily contradict itself. It shows that the effect of resources depends on the reason for migration and on the outside options available to the household. If the household can improve its welfare by remaining in place, an additional transfer may reduce migration. If the household wants to move but cannot pay the fixed costs of doing so, the same transfer may increase migration. Climate shocks make this distinction more important because they can change both sides of the decision at the same time: they can reduce the value of remaining in place and also reduce the resources needed to leave.

This distinction is central for the interpretation of cash transfers as an adaptation policy. The same transfer can reduce migration after a temporary shock if it helps households smooth consumption and repair losses, but it can also increase migration when households face repeated shocks and need liquidity to finance relocation. Therefore, the expected sign of the interaction between climate shocks and cash transfers is ambiguous ex ante and depends on the type, duration, and perceived persistence of the shock.

The studies by Halliday (2006), Nawrotzki & Bakhtsiyarava (2017), Cattaneo & Peri (2016),and Gröschl & Steinwachs (2017) elaborate on the climate-inhibitor mechanism and liquidity constraints that poor countries and poor households face after adverse climate shocks, thus making people unable to migrate. However, their findings are based on the comparison between rich and poor and do not reveal how poor households respond to climate shocks when they gain better access to financial resources. Our study contributes to the literature by addressing this gap. Specifically, we use cash-transfer programs in Indonesia to identify household adaptation strategies in reaction to environmental distress when they have better access to financial resources.

2.3 Conceptual Framework

This section clarifies why access to financial resources can either increase or decrease migration after environmental shocks. The main idea is that migration is both an adaptation strategy and a costly investment. A household may want to leave an environmentally affected area, but migration requires money, information, networks, and the ability to bear risk. At the same time, remaining is also not costless. Households that remain in place may need to repair damaged assets, smooth consumption, change occupations, or invest in local adaptation. Therefore, better access to financial resources can support two different responses.

The first response is the *climate-driver mechanism*. Under this mechanism, environmental shocks make the current place of residence less attractive or less economically viable. If a household receives additional resources, they can relax liquidity constraints and make migration feasible. This mechanism is most likely when shocks are persistent, when they change expectations about future living conditions, or when the household believes that local adaptation will not be sufficient. In this case, cash transfers should increase migration after environmental shocks.

The second response is the *climate-inhibitor mechanism*. Under this mechanism, environmental shocks create short-run losses, but the household prefers to remain in place if it has enough resources to cope locally. Cash transfers can help households repair damages, smooth consumption, avoid distress sales of assets, and maintain local livelihoods. In this case, better access to financial resources reduces migration after environmental shocks because households are able to adapt in place.

These two mechanisms imply that the sign of the interaction between climate shocks and cash transfers is ambiguous ex ante. A positive interaction means that cash transfers increase migration in response to a shock, which is consistent with the climate-driver mechanism. A negative interaction means that cash transfers reduce migration in response to a shock, which is consistent with the climate-inhibitor mechanism. The interpretation depends on the type of shock.

Distinguishing between short-term and persistent shocks is important for interpreting how households respond to environmental stress. A current-year rainfall anomaly or a natural disaster may be perceived as a temporary shock. In this case, households may use additional resources to repair losses or smooth consumption, so migration may fall

among cash-transfer recipients. By contrast, accumulated rainfall shocks or repeated temperature anomalies may signal a more persistent deterioration in local conditions. In this case, migration becomes a more attractive adaptation strategy, and cash transfers may provide the liquidity needed to move.

This framework also helps interpret why poor households are the relevant population for the analysis. Wealthier households may already have savings, credit access, insurance, or migration networks. For them, receiving a transfer may not substantially change the feasible set of choices. Poor households, by contrast, may be close to both liquidity and subsistence constraints. A transfer can therefore meaningfully change whether they can cope locally or finance migration. For this reason, the analysis focuses on resource-constrained poor households.

The empirical analysis tests which mechanism dominates for different environmental shocks. If cash transfers reduce migration after current rainfall shocks and natural disasters, this suggests that households use transfers to remain in place and absorb short-run losses. If cash transfers increase migration after accumulated rainfall shocks or temperature anomalies, this suggests that transfers relax liquidity constraints and allow households to leave when environmental stress is more persistent.

Table 2.1: Predicted Migration Responses Under Alternative Mechanisms

Shock type	Expected role of cash transfers	Predicted interaction	Interpretation
Current rainfall shock	Smooth short-run losses and maintain local consumption	Negative	Cash transfers support in-place adaptation and reduce immediate migration.
Accumulated rainfall shock	Finance relocation after repeated environmental stress	Positive	Cash transfers relax liquidity constraints and make migration feasible.
Temperature anomaly	Respond to broader deterioration in local productivity or living conditions	Positive	Cash transfers help households overcome mobility constraints.
Natural disaster	Repair losses and cope locally after rare, unexpected events	Negative	Cash transfers reduce the need for longer-term relocation after temporary shocks.

Notes: The table summarizes the expected signs of the interaction between cash transfers and environmental shocks. A positive interaction is consistent with the climate-driver mechanism, in which financial resources enable migration. A negative interaction is consistent with the climate-inhibitor mechanism, in which financial resources support adaptation in place.

2.4 Background and Data

Climate change has affected Indonesia with more frequent droughts, heatwaves, and floods occurring in the country. These shocks pose an increasing threat to the country's development (Ministry of Foreign Affairs of the Netherlands, 2018). The World Bank has ranked Indonesia 12th among 35 countries that face high mortality risks due to multiple hazards such as increasing temperature, changing rainfall patterns, and, more importantly, exposure to various natural disasters (World Bank, 2010). Considering the high exposure of Indonesia to environmental hazards, the country offers a well-suited setting for our analysis.

Our choice of Indonesia is also based on the presence of the country's governmental cash-transfer programs: Bantuan Langsung Tunai (BLT), Program Keluarga Harapan

(PKH), and Bantuan Langsung Sementara Masyarakat (BLSM) conducted by the government of Indonesia (World Bank, 2012a,b, 2017; SMERU Research Institute, 2013). BLT and BLSM are unconditional cash transfer programs, while PKH is a conditional one. In our research, the reception of cash transfer serves as a case of better access to financial resources for poor households. The goal of the unconditional cash-transfer programs was to increase the consumption of the poor population by providing them with some external finances. Although there is no spatial variation in receiving these transfers, there are temporal differences: we observe in the data that different households received the transfers in different years in a small window. Common national timing is absorbed by year fixed effects, so the remaining variation for the unconditional programs comes from differences in household receipt over time.

The goal of the PKH is the alleviation of poverty by providing financial resources to families with children if they fulfill basic obligations by utilizing health and education services (World Bank, 2012c). The useful (for our research design) features of this particular conditional cash-transfer program is the gradual introduction of the program to other provinces. The PKH first started with a pilot-program and only after that gradually started to expand over the country, which, for our study, creates time and spatial variation in the access to this program. The seven provinces that participated in the pilot program are West Java, East Java, West Sumatra, North Sulawesi, Gorontalo, East Nusa Tenggara, and DKI Jakarta. These provinces were chosen to represent various types of areas that are present in Indonesia. PKH additionally provides geographic variation through its phased provincial rollout. The estimating specification therefore combines PKH rollout variation with household-level receipt variation rather than relying on national program timing alone.

We concentrate on the population that remains consistently poor in our studied period in order to have a best possible treatment and control groups. Thus, we focus on households with similar characteristics that are eligible for cash-transfer programs. This restriction is important because the object of interest is not the average Indonesian household, but the group for which financial constraints are most likely to shape adaptation choices. Poor households tend to have fewer savings, weaker access to formal credit, and fewer ways to insure themselves against shocks. As a result, a government transfer can represent one of the few sources of liquidity available when an environmental shock occurs. This restriction defines the estimand as applying to households that remain poor

over the study period. Because later poverty status can itself be affected by shocks or migration, the restriction can exclude households that subsequently leave poverty.

This study uses variation in cash transfer receipt to examine how better access to financial resources affects household adaptation to climate shocks, with a particular focus on migration as an adaptation strategy. For households with sufficient resources, both local adaptation and migration may be financed through savings, credit, or family networks. For poorer households, however, these options are more limited. This makes cash transfers especially relevant for understanding whether households remain in place or migrate in response to environmental stress.

The same focus also helps to narrow the interpretation of the cash transfer variable. In this sample, transfers are not treated as a general income shock for all Indonesian households, but as an additional resource available to households that are close to the eligibility margin of government support. This makes the empirical question more specific: conditional on being poor and exposed to environmental stress, do households use additional public resources to remain in place, or do they use them to finance migration? The answer to this question is not obvious ex ante, because both responses are costly and both can be interpreted as adaptation.

To conduct the analysis, we use data from several sources, the first of which is the Indonesian Family Life Survey (IFLS)[3], an on-going longitudinal survey in Indonesia. The IFLS is representative of about 83% of the Indonesian population (Strauss et al., 2016) and has 5 waves: IFLS1 in 1993-1994, IFLS2 in 1997-1998, IFLS3 in 2000, IFLS4 in 2007-2008, and IFLS5 in 2014-2015. Since the PKH was introduced in 2007, the last three waves of the IFLS are appropriate sources of information for our analysis. The IFLS data is especially suitable for studying migration due to the high effort to track respondents over time, which leads to low rates of attrition. This data set records all the moves longer than 6 months by people older than 12 years of age. In addition to detailed information about migration, the IFLS also has data on socio-economic and demographic characteristics of households and individuals. Migration timing is reconstructed from retrospective histories, while cash-transfer receipt is observed through the IFLS survey information; the data therefore do not establish the exact within-year ordering of transfer receipt and migration.

[3]Retrieved from `https://www.rand.org/well-being/social-and-behavioral-policy/data/FLS/IFLS/download.html`.

The migration information in the IFLS is useful for our purpose because it allows us to separate different forms of movement. In the main analysis, we use an indicator for whether a household has any recorded move. In additional specifications, we distinguish moves by one person from moves by part or whole households. This distinction is not only technical. A move by one person may be closer to a labor-migration strategy, where the household remains in the origin location and one member leaves to reduce risk or earn income elsewhere. A move by a larger part of the household is closer to relocation and may therefore require more financial resources. We use this difference when interpreting the accumulated-shock results. The one-person and family-move indicators can overlap within the same household-year, so they are not an exhaustive additive decomposition of the indicator for any migration.

The six-month threshold in the survey also shapes the interpretation of the outcomes. We do not observe very short moves, emergency evacuation, or seasonal movements that last less than six months. Therefore, the migration outcomes in this chapter capture more persistent responses to environmental stress. This is useful because it allows us to study decisions that are costly enough to be related to financial constraints, but it also means that the results should not be interpreted as a complete description of all temporary displacement after disasters.

Sources that provide data about weather characteristics are abundant. They include the dataset produced by the Climatic Research Unit at the University of East Anglia, weather data obtained from the Center for Climatic Research at the University of Delaware, and NASA MERRA-2, among others. All of these datasets provide information about precipitation and temperature on grids of 0.5 by 0.5 degrees, which corresponds to about 50 by 50 kilometers in Indonesia. We use weather data obtained from the Climatic Research Unit at the University of East Anglia[4], particularly data about temperature and precipitation, which we use to investigate the effect of slow-onset events on migration.

For natural disasters we use the data from the IFLS, since this survey also has information about natural disasters occurring in different parts of Indonesia. This data turned out to be more precise for our purposes than our initial choice, the EM-DAT. Data from the International Disaster Database (EM-DAT) record natural disasters and

[4]Retrieved from `https://www.uea.ac.uk/groups-and-centres/climatic-research-unit`.

can be used to study the effect of the fast-onset events. The EM-DAT contains data on the occurrence and effects of mass disasters in the world from 1900 until today.

The difference between the weather data and the natural-disaster data is useful for the empirical analysis. Rainfall and temperature anomalies capture deviations from long-run local climate conditions and can be interpreted as slow-onset environmental pressure. Natural disasters, by contrast, are more sudden events that may damage assets or local infrastructure directly. We therefore use these two sources of environmental variation to examine whether cash transfers have similar or different roles across different types of shocks.

2.4.1 Measurement of Migration, Shocks, and Financial Resources

Table 2.2 summarizes the main variables used in the analysis. The table is included to make the empirical object transparent. The outcome variables capture migration decisions at the household-year level. The main explanatory variables combine exposure to environmental shocks with receipt of cash transfers, which proxy for improved access to financial resources.

The migration outcomes should be interpreted as relatively durable moves because the IFLS records moves lasting longer than six months. This is useful for studying costly and meaningful relocation decisions, but it also means that short seasonal moves and very short displacement episodes are not fully observed. The climate variables capture both slow-onset shocks, including rainfall and temperature anomalies, and fast-onset events, including floods, landslides, and droughts. The cash transfer variable captures whether a poor household receives additional public resources in a given year.

Table 2.2: Main Variables Used in the Indonesia Analysis

Variable	Definition	Interpretation
Migration	Indicator equal to one if the household sends a migrant in a given year	Main mobility outcome
One-person migration	Indicator equal to one if one household member migrates	Partial household adjustment
Family migration	Indicator equal to one if the whole household or a larger part of the household migrates	More complete relocation response
Cash transfer	Indicator for receiving a government cash transfer in a given year	Proxy for improved access to financial resources
Rainfall shock	Deviation in rainfall conditions relative to historical local patterns	Slow-onset environmental stress
Accumulated rainfall shock	Repeated rainfall deviation over several years	Persistent environmental stress
Temperature shock	Temperature anomaly relative to historical local conditions	Slow-onset environmental stress
Natural disaster	Indicator for reported floods, landslides, droughts, or related disasters	Fast-onset environmental shock

Notes: The table summarizes the main outcome, transfer, and shock variables. Cash transfer receipt proxies for access to financial resources. Rainfall and temperature shocks capture slow-onset environmental stress, while natural disasters capture fast-onset events.

After combining the data about precipitation and temperature with different waves from the IFLS we now have the household-year panel dataset, which contains information about poor households in Indonesia with several socio-economic characteristics as well as the history of weather and natural disasters that they experienced each year. Table 2.3 shows the summary statistics of our panel dataset. The table shows that although the sample contains households that are poor, there are still differences in demographic structure, household size, assets, and mobility-related characteristics. Our analysis is based on a panel in which the same household can be observed under different environmental and transfer conditions. This panel structure is important for the fixed-effects specification used below.

Table 2.3: Summary Statistics

	Mean	SD	Min	Max	N
Migration	0.08	0.26	0.00	1.00	17,546
1 Person Migration	0.05	0.23	0.00	1.00	17,546
Family Migration	0.03	0.17	0.00	1.00	17,546
Cash Transfer	0.16	0.36	0.00	1.00	17,546
Temperature Shock	0.15	0.36	0.00	1.00	17,546
Rainfall Shock	0.08	0.27	0.00	1.00	17,542
Flood	0.03	0.18	0.00	1.00	17,546
Landslide	0.01	0.10	0.00	1.00	17,546
Tsunami	0.00	0.01	0.00	1.00	17,546
Drought	0.02	0.13	0.00	1.00	17,546
Assets	2.00	0.93	1.00	4.00	17,544
Urban	0.33	0.47	0.00	1.00	17,546
Household Size	4.88	2.47	1.00	16.00	17,536
Female household head	0.13	0.34	0.00	1.00	17,354
Share of children	0.35	0.23	0.00	1.00	17,534
Share of females	0.53	0.20	0.13	1.00	17,411
Share of mobile	0.37	0.16	0.08	1.00	15,814
Muslim household	0.85	0.35	0.00	1.00	17,534

2.5 Identification Strategy

The empirical strategy estimates whether access to financial resources changes the migration response to environmental shocks. The main object of interest is not only the direct effect of a climate shock on migration, or the direct effect of receiving a cash transfer. Instead, the key coefficient is the interaction between climate shocks and cash-transfer receipt. This interaction asks whether households with better access to financial resources respond differently to the same environmental shock than households without such access.

The baseline specification is

$$\begin{aligned} M_{ist} =& \alpha + \beta_1 \text{ClimateShock}_{ist} + \beta_2 \text{CashTransfer}_{ist} + \beta_3 \left(\text{ClimateShock}_{ist} \times \text{CashTransfer}_{ist}\right) \\ &+ X'_{ist}\Gamma + \mu_i + \mu_s + \mu_t + \varepsilon_{ist}, \end{aligned} \tag{2.1}$$

where M_{ist} is a migration outcome for household i in sub-district s and year t. The variable ClimateShock$_{ist}$ captures exposure to rainfall shocks, temperature shocks, or natural disasters. The variable CashTransfer$_{ist}$ is equal to one if the household receives a cash transfer in that year. The vector X_{ist} includes household controls. Household fixed effects, μ_i, absorb time-invariant household characteristics; sub-district fixed effects, μ_s, absorb persistent local differences; and year fixed effects, μ_t, absorb common shocks across households in a given year. Reported standard errors are clustered at the sub-district level; spatial correlation at broader geographic scales is not separately modeled.

The coefficient of interest is β_3. A positive β_3 means that cash-transfer households migrate more in response to environmental shocks than comparable non-recipient households. This is consistent with the climate-driver mechanism: financial resources relax liquidity constraints and make migration feasible. A negative β_3 means that cash-transfer households migrate less in response to environmental shocks. This is consistent with the climate-inhibitor mechanism: financial resources help households cope locally and reduce the need to migrate.

The identifying variation comes from two sources. The first is variation in environmental shocks across locations and time. Rainfall and temperature anomalies vary geographically and temporally, and natural disasters affect different households in different years. The second is variation in access to cash transfers. Indonesia's cash-transfer programs generate variation in whether poor households receive additional financial resources in a given year. The empirical specification combines these sources of variation to test whether the migration response to environmental shocks differs by access to financial resources. For cash transfers, common national timing is absorbed by year fixed effects; the remaining variation combines PKHś geographic rollout with household-level receipt variation within locations and over time.

Several features of the specification are important. Household fixed effects compare households to themselves over time and therefore remove permanent household differences, including baseline migration propensity, long-run wealth, risk preferences, or persistent migration networks. Year fixed effects remove common macroeconomic, political, or climatic shocks that affect all households in a given year. With household fixed effects included, time-invariant sub-district characteristics are already absorbed for households that do not move; we therefore do not interpret the sub-district fixed effect as a separate source of identification.

The remaining identification concern is that cash-transfer receipt may be correlated with time-varying household conditions that also affect migration. This is why the interpretation of the interaction term requires caution. The estimates should be read as evidence on whether migration responses to shocks differ by access to financial resources among poor households, not as a fully randomized cash-transfer experiment in every specification. The analysis is nevertheless informative because it focuses on poor households, includes household fixed effects, and compares responses to environmental shocks across households with different access to resources. In addition, because transfer receipt and migration are measured at annual frequency, the specification does not establish their within-year ordering; this is particularly relevant for PKH because participation is linked to local health and education services.

This interpretation is also why the analysis distinguishes different types of shocks. If cash transfers had the same effect for all environmental stress, one would expect the interaction coefficient to have the same sign across rainfall shocks, temperature anomalies, and natural disasters. Instead, the conceptual framework predicts that the sign may differ. Short-run shocks may induce local coping, while persistent shocks may induce migration if resources make relocation feasible. The empirical strategy is therefore designed to test not only whether cash transfers matter, but how their role changes with the nature of the environmental shock.

Changing rainfall patterns demonstrate more variability in Indonesia than temperature variation, which is visible in Figures 2.1 and 2.2, with larger variation in rainfall than temperature. As a result, floods and landslides are two of the most frequent disastrous events that occur in Indonesia, both of which are precipitated by heavy rains. For this reason, we concentrate more on rainfall variation, but we also explore temperature variation and natural disasters including floods, landslides and droughts. Thus, we explore two types of climate shocks: slow-onset events (temperature and precipitation anomalies) and fast-onset events (natural disasters).

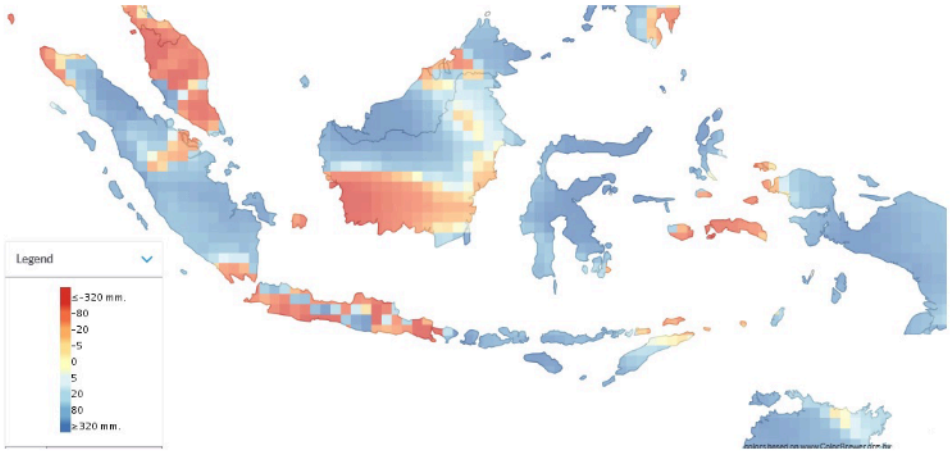


Figure 2.1: Trends in Rainfall, 1901–2012



Figure 2.2: Trends in Temperature, 1901–2012

Notes: Figure 2.1 presents areas in Indonesia with less rain than the historical mean (red) and more rain than the historical mean (blue). Figure 2.2 does the same for temperature trends—higher temperatures (red) and lower temperatures (blue).
Source: Climatic Research Unit at the University of East Anglia.

2.5.1 Definition of Rainfall Anomalies

To measure rainfall shocks, we construct standardized rainfall anomalies at the sub-district-year level. This approach allows each sub-district to be compared to its own historical rainfall distribution rather than to a national average. This is important because normal rainfall levels differ substantially across regions in Indonesia.

For sub-district s in year t, the rainfall anomaly is defined as

$$RA_{st} = \frac{Rainfall_{st} - \overline{Rainfall}_s}{SD_s}, \tag{2.2}$$

where $Rainfall_{st}$ is total rainfall in sub-district s in year t, $\overline{Rainfall}_s$ is the historical mean rainfall in sub-district s, and SD_s is the historical standard deviation of rainfall in the same sub-district. The variable RA_{st} therefore measures how unusual rainfall in a given year is relative to the local historical distribution.

We define a rainfall shock as an extreme deviation from normal local rainfall conditions. A household is classified as exposed to a rainfall shock in year t if it lives in a sub-district where

$$|RA_{st}| > 2.$$

Because the indicator uses the absolute standardized anomaly, it pools unusually wet and unusually dry years. The coefficient should therefore be interpreted as the response to extreme rainfall deviations rather than as a drought-specific or flood-specific effect.

Figure 2.3 shows the distribution of rainfall anomalies during the study period. The red vertical lines indicate the thresholds used to define rainfall shocks. Observations outside these thresholds are classified as rainfall-shock observations.

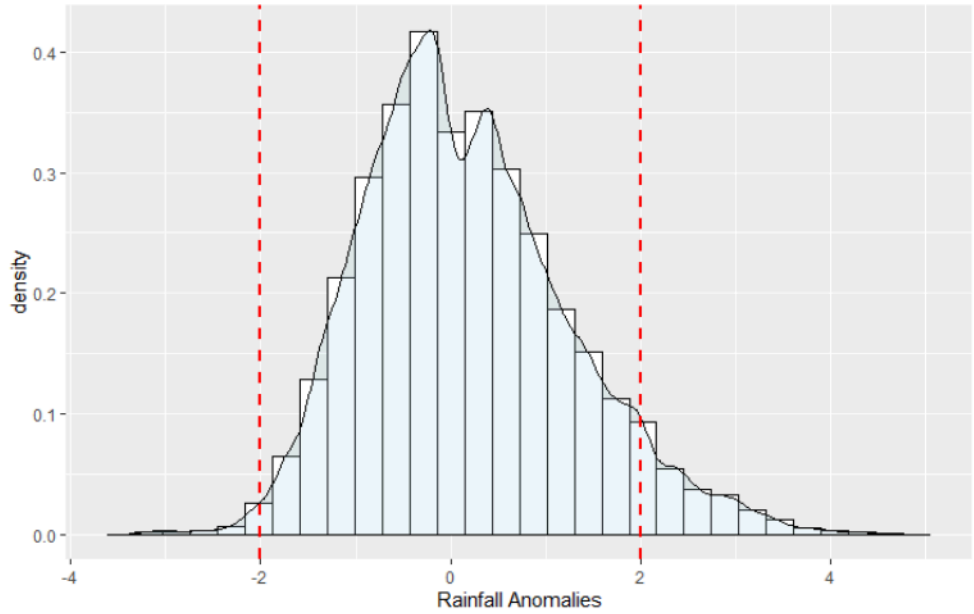


Figure 2.3: Distribution of Rainfall Anomalies

Notes: The figure presents the distribution of standardized rainfall anomalies during the study period. The red vertical lines indicate the thresholds used to construct the rainfall-shock variable. A sub-district-year is classified as experiencing a rainfall shock if the rainfall anomaly is greater than two standard deviations above or below the sub-district's historical mean.

Source: Authors' calculations based on data from the Climatic Research Unit at the University of East Anglia.

Figure 2.4 maps rainfall anomalies across Indonesia in selected years. The spatial distribution of anomalies does not show a persistent concentration in the same areas over time. This supports the empirical strategy, which relies on unexpected deviations from local historical rainfall patterns rather than on permanent differences between wetter and drier regions.

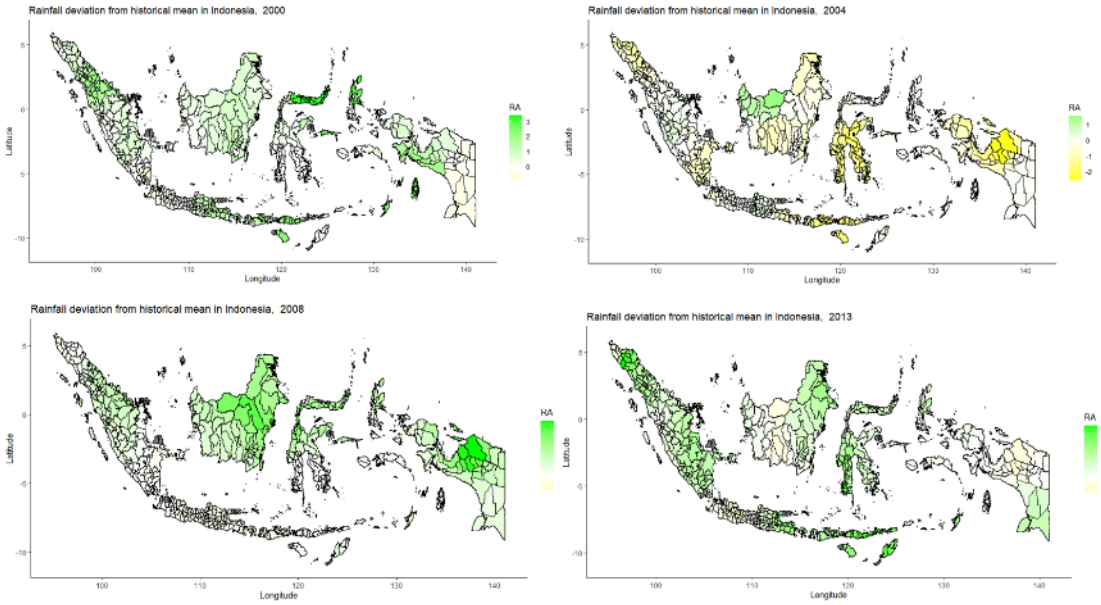


Figure 2.4: Rainfall Anomalies in Indonesia

Notes: The figure presents maps of standardized rainfall anomalies in selected years. Green areas indicate rainfall above the local historical mean, while yellow areas indicate rainfall below the local historical mean. The absence of a stable spatial pattern suggests that rainfall anomalies vary across both locations and years.

Source: Authors' calculations based on data from the Climatic Research Unit at the University of East Anglia.

2.5.2 Checks for Rainfall Shocks

Before moving to the main results, we check whether the rainfall shocks behave as expected for the empirical strategy. The first check is whether future rainfall shocks predict current migration. If households were able to anticipate the shocks and migrate before they occur, the estimated response would be difficult to interpret as a response to the realized shock. Table 2.4 reports this check. The coefficients on the leads do not suggest that migration is systematically increasing before the rainfall shock, which supports the interpretation that the shock is not simply capturing pre-existing migration trends.

Table 2.4: Leads of Rainfall Shocks and Migration

	Migration	
	(1)	(2)
Rainfall Shock t+1	−0.0092	
	(0.0126)	
Rainfall Shock t+2		−0.0127
		(0.0134)
Observations	17099	15937
Controls	*Yes*	*Yes*
Time Fixed Effects	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates for leads of rainfall shocks and their influence on migration decisions. None of the leads predicts migration. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

The second check is whether receiving a cash transfer is systematically related to rainfall shocks. This is important because our empirical question relies on comparing the migration response of households with and without access to transfers under similar environmental shocks. Table 2.5 reports the orthogonality check. The results do not indicate that rainfall shocks predict the receipt of a cash transfer after conditioning on the control variables and fixed effects. This reduces the concern that the interaction term is mechanically driven by the targeting of transfers to currently shocked locations.

Table 2.5: Orthogonality Check: Rainfall Shocks

	Cash Transfer
	(1)
Rainfall Shock	−0.0126
	(0.0181)
Observations	18291
Controls	*Yes*
Time Fixed Effects	*Yes*
Household Fixed Effects	*Yes*
Sub-district Fixed Effects	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates for the relationship between rainfall shocks and receipt of a cash transfer. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

2.5.3 Definition of Temperature Anomalies

To measure temperature shocks, we construct standardized temperature anomalies at the sub-district-year level. As with rainfall, this approach compares each sub-district to its own historical temperature distribution rather than to a national average. This is important because average temperature levels differ across regions in Indonesia.

For sub-district s in year t, the temperature anomaly is defined as

$$TA_{st} = \frac{Temperature_{st} - \overline{Temperature}_s}{SD_s}, \tag{2.3}$$

where $Temperature_{st}$ is the average temperature in sub-district s in year t, $\overline{Temperature}_s$ is the historical mean temperature in sub-district s, and SD_s is the historical standard deviation of temperature in the same sub-district. The variable TA_{st} therefore measures how unusual temperature in a given sub-district-year is relative to local historical conditions.

Unlike rainfall anomalies, the distribution of temperature anomalies during the study period is not centered around zero. As shown in Figure 2.5, temperature in Indonesia during the sample period was generally higher than historical averages. For this reason, we do not use a symmetric two-standard-deviation rule for temperature shocks. Following the literature, we define high-temperature shocks as observations for which

$$TA_{st} > 2.$$

To define the lower-temperature comparison group, we use the threshold

$$TA_{st} < 0.4.$$

The second cutoff is chosen so that the share of observations below 0.4 is approximately the same as the share above 2. Because temperatures during the sample period are shifted upward relative to the historical mean, the lower category should not be interpreted as an extreme-cold shock. The temperature indicator is therefore an asymmetric composite comparing very hot years with a relatively lower-temperature group.

Figure 2.5 presents the distribution of temperature anomalies during the study period. The red vertical lines indicate the thresholds used to construct the temperature-shock variable. Observations beyond these thresholds are classified as temperature-shock observations.

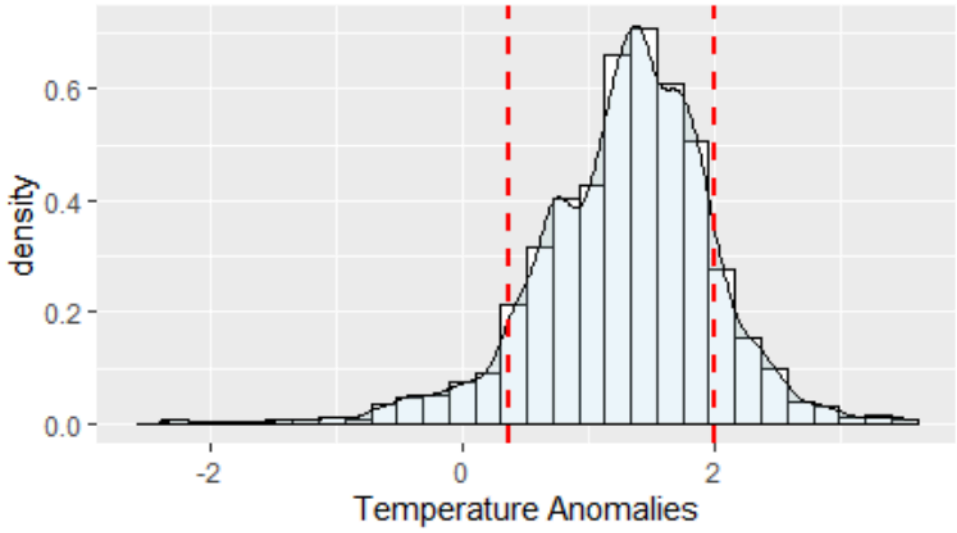


Figure 2.5: Distribution of Temperature Anomalies

Notes: The figure presents the distribution of standardized temperature anomalies during the study period. The red vertical lines indicate the thresholds used to construct the temperature-shock variable. High-temperature shocks are defined as sub-district-years with $TA_{st} > 2$. The lower-temperature comparison group is defined by $TA_{st} < 0.4$; this cutoff is chosen to create a group of approximately the same size and should not be interpreted as an extreme-cold tail. *Source:* Authors' calculations based on data from the Climatic Research Unit at the University of East Anglia.

Figure 2.6 maps temperature anomalies across Indonesia in selected years. The maps show that temperature anomalies vary across both space and time, although the spatial variation is somewhat less pronounced than for rainfall. As with rainfall shocks, this variation is useful for the empirical strategy because it allows the analysis to compare household responses across locations and years with different exposure to abnormal temperature conditions.

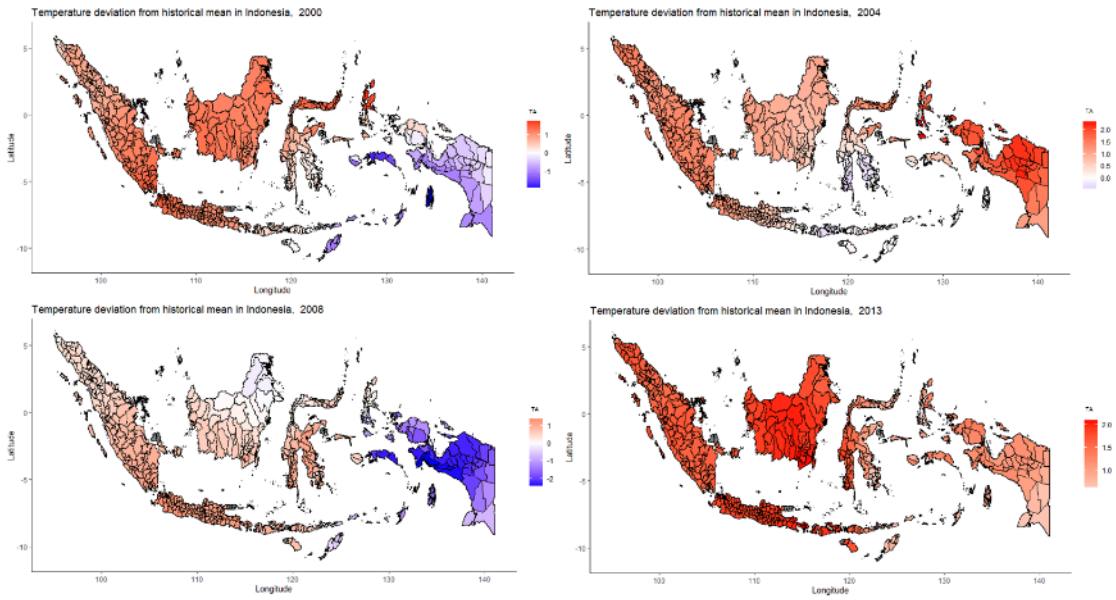


Figure 2.6: Temperature Anomalies in Indonesia

Notes: The figure presents maps of standardized temperature anomalies in selected years. Red areas indicate temperatures above the local historical mean, while blue areas indicate temperatures below the local historical mean. The maps show that temperature anomalies vary across locations and years, although less sharply than rainfall anomalies.
Source: Authors' calculations based on data from the Climatic Research Unit at the University of East Anglia.

As with rainfall, we also check that the temperature-shock results are not driven by anticipation of shocks or by a mechanical relationship between shock exposure and cash-transfer receipt. These checks are important because the empirical interpretation relies on temperature anomalies representing unexpected deviations from local historical conditions rather than predictable differences across locations.

2.5.4 Checks for Temperature Shocks

The same two checks are useful for temperature shocks. Table 2.6 reports the estimates with leads of temperature shocks. The purpose is to see whether households move before the temperature shock is realized. The absence of a systematic lead pattern is important because the temperature variable is based on deviations from long-run local conditions, and these deviations should not predict migrations before they occur.

Table 2.6: Leads of Temperature Shocks and Migration

	Migration	
	(1)	(2)
Temperature Shock t+1	−0.0159	
	(0.0097)	
Temperature Shock t+2		−0.0072
		(0.0103)
Observations	16275	14912
Controls	*Yes*	*Yes*
Time Fixed Effects	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates for leads of temperature shocks and their influence on migration decisions. None of the leads predicts migration. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

Table 2.7 then checks whether temperature shocks predict the receipt of a cash transfer. This check is especially relevant because temperature shocks are more persistent than rainfall shocks in the period studied. If transfer receipt were systematically higher in districts with unusually high temperatures, the interaction term would be harder to interpret. The results do not show such a pattern, which supports using the interaction term as evidence on how financial resources moderate the response to temperature shocks.

Table 2.7: Orthogonality Check: Temperature Shocks

	Cash Transfer
Temperature Shock	0.0147
	(0.0125)
Observations	17406
Controls	*Yes*
Time Fixed Effects	*Yes*
Household Fixed Effects	*Yes*
Sub-district Fixed Effects	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates for the relationship between temperature shocks and receipt of a cash transfer. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

2.6 Results

The results are organized by type of environmental shock. This structure follows the conceptual framework. We first examine current-year rainfall shocks, which are more likely to capture short-term deviations in local conditions. We then examine accumulated rainfall shocks, which better capture persistent environmental stress. Next, we turn to temperature anomalies and natural disasters. Finally, we examine occupational adjustment as an additional margin of adaptation for households that remain in place.

The same cash transfer can support different adaptation responses depending on the nature of the shock. A negative interaction between cash transfers and a current-year rainfall shock suggests that transfers help households cope locally. A positive interaction after accumulated rainfall shocks or temperature anomalies suggests that transfers relax liquidity constraints and make migration possible. The results should therefore be interpreted as evidence on the mechanism, not simply as evidence that cash transfers increase or decrease migration in general.

2.6.1 Rainfall Shocks

In this subsection, we concentrate on the precipitation shocks in the data. Based on the literature (Thiede & Gray (2017); Mueller et al. (2014); Nawrotzki & Bakhtsiyarava (2017)), we construct the climate shock variable as follows: if average yearly precipitation in *kecamatan* (sub-district in Indonesia) deviated from the historical mean (calculated on the time span from 1901-2015) by more than 2 standard deviations, then *kecamatan* is considered to have had a shock.

Table 2.8 shows the relationship between rainfall shocks, cash-transfer status, and migration decisions. The identification checks in Tables 2.4 and 2.5 support the interpretation that the coefficient of interest is not driven by anticipation of the shock or by a mechanical relationship between rainfall shocks and cash-transfer receipt.

In line with Thiede & Gray (2017), precipitation variation in the form of rainfall shocks is not statistically significant predictor of migration Table 2.8 shows that the interaction term of cash transfer and rainfall shocks is decreasing migration in treated areas: under the shock, households that had access to a government cash transfer in that year decided to migrate less frequently than those who did not receive a cash transfer. The result is

Table 2.8: Rainfall Shocks, Cash Transfers, and Migration

	Migration and Rainfall Shocks		Migration, CT, and Rainfall Shocks	
	(1)	(2)	(3)	(4)
Rainfall Shock	−0.0151	−0.0146	−0.0093	−0.0090
	(0.0127)	(0.0126)	(0.0142)	(0.0141)
Cash Transfer			0.0084	0.0075
			(0.0077)	(0.0078)
Rainfall Shock × Cash Transfer			−0.0402**	−0.0391**
			(0.0190)	(0.0191)
Observations	18440	18291	18440	18291
Controls	*No*	*Yes*	*No*	*Yes*
Time Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table presents the baseline regression results. The dependent variable is the indicator for all migration decisions. Columns (1) and (2) report estimates for the rainfall shock specification without and with controls. Columns (3) and (4) report estimates from the main specification without and with controls. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

stable in size and significance in specifications both with and without controls. In the controlled specification, the interaction coefficient is about -0.039, corresponding to a 3.9 percentage-point difference; Table 2.3 reports an average migration rate of about 8 percent in the summary sample. We conclude that the effect of interest - the interaction between rainfall shock and receiving a cash transfer - has a significant negative effect on the migration of poor households, which suggests that the interaction term facilitates the climate-inhibitor mechanism in the short term. We believe that the logic is similar to an insurance mechanism: since migrating is a costly strategy in terms of both financial and mental costs, affected households use the finance provided by the government as a means to adapt to a short-term shock in the place of residence rather than choosing costly migratory behaviour. To strengthen this result, in the Robustness checks section, we estimate several different specifications. First, we change the definition of the shock from a yearly dummy to how many months in a given year the rainfall exceeds the monthly historical mean (Table 2.17). We find a similar pattern, with a cash transfer decreasing the probability of migration as a response to a rainfall shock. Subsequently, we try the same baseline specification, but we change the period studied to a larger one - from

1997 to 2014 (Table 2.18). We use a smaller period in the main specification because cash transfer programs started from 2005. To have data more or less symmetrically distributed around those years, we use three waves of the IFLS beginning from 2000. In this check too, the result survived the inclusion of additional time periods and we still observe a negative effect of the interaction term on the migration decision.

This result should not be interpreted as evidence that rainfall shocks are irrelevant for poor households. Rather, it suggests that the observed migration response depends on whether households have enough resources to cope with the shock locally. A rainfall anomaly can damage agricultural production, affect income, or increase the need for short-term expenditures. If the cash transfer helps the household cover these costs, the immediate need to move may become lower. This is why the negative interaction term is consistent with local adaptation rather than with the absence of climate-related pressure.

Another reason why the short-term response may be negative is that migration itself requires planning. Households need information about destinations, travel costs, housing, and expected employment. A one-year rainfall shock may not be enough to make this costly decision attractive, especially if the household expects conditions to return to normal. In this case, the transfer may be used first to smooth the shock at the place of residence. The interpretation is therefore different from the interpretation of accumulated shocks below, where repeated exposure may change the household's expectations about future risk.

We find some evidence of a behavioral change from a cash transfer as a response to a current shock. However, we are also interested in the influence of being exposed to a rainfall shock several times. Thus, we next explore accumulated rainfall shocks and their connection with migration and the moderating effect of cash transfers.

This distinction between a current shock and accumulated exposure is important. A single rainfall anomaly may be interpreted by households as temporary and may not be sufficient to justify the financial, informational, and psychological costs of migration. Repeated exposure, however, can change expectations about the future quality of the location. In this case, migration may become a more attractive adaptation strategy, but only for households that have enough resources to finance it. The accumulated-shock specification is designed to capture this difference between temporary coping and longer-term adjustment.

We construct a measure of accumulated shock for several periods that occur if a household receives a sum of the rainfall anomalies exceeding a certain threshold. We define Accumulated shock=1 if $\sum_{t-n}^{t} |RA_{st}| > Cutoff$. Because this measure also uses absolute anomalies, accumulated exposure likewise combines unusually wet and unusually dry years. The baseline results of the model with n=3 and *Cutoff*=5 are in Table 2.9. To ensure that our results are not dictated by chosen thresholds, we conduct robustness checks with a different number of years for which the shock is calculated and different cutoffs, see Table 2.19 in the Robustness Checks section. The orthogonality check for accumulated rainfall shocks is reported in Table 2.10.

Table 2.9: Accumulated Rainfall Shocks, Cash Transfers, and Migration

	Migration, Accumulated Shock for 4 periods		
	(1) All	(2) 1 person	(3) Family
Accumulated Shock for 4 Periods	−0.0234**	−0.0197**	0.0066
	(0.0093)	(0.0085)	(0.0071)
Cash Transfer	0.0066	0.0035	−0.0163*
	(0.0077)	(0.0069)	(0.0095)
Accumulated Shock for 4 Periods × Cash Transfer	−0.0124	−0.0096	0.0997**
	(0.0132)	(0.0114)	(0.0389)
Observations	21114	21114	21114
Controls	*Yes*	*Yes*	*Yes*
Time Fixed Effects	*Yes*	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates from the main model. The outcome variables are all moves, moves undertaken by one person, and moves undertaken by part or whole households in Columns (1), (2), and (3), respectively. The one-person and family-move indicators can overlap within a household-year and therefore are not expected to add to the all-move outcome. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

Several things are of interest here. First, the accumulated shocks have a heterogeneous effect on the migratory decisions of the poor population depending on the types of moves. It seems to decrease the probability of sending 1 migrant from the household, but the sign for migration of the whole household is different. The interaction term seems to have a negative effect on all moves and moves undertaken by 1 person, but it has a positive significant effect on moves undertaken by families, which implies that for this group a cash

transfer facilitates the climate-driver mechanism. The family-move interaction is 0.0997, or about 10 percentage points; Table 2.3 reports a family-move mean of about 3 percent in the summary sample. The one-person and family-move indicators overlap within some household-years, so their coefficients are not expected to sum to the coefficient on all migration. The accumulated result therefore concerns the family-relocation margin rather than an increase in migration overall. This might be explained by the households changing their beliefs about how hazardous the place of residence is: a one-period shock might not be enough for households to change their beliefs, but being exposed to a shock several times leads to a reassessment of the risk of being hit by a rainfall shock in the future, which leads to the household choosing the migration option as an adaptation strategy.

The difference between individual moves and moves of part or whole households is also informative. Sending one migrant can be interpreted as a diversification strategy that allows the household to maintain its original residence while gaining access to income from another location. Moving with a larger part of the household is a stronger form of adjustment and may require more resources. The positive interaction term for family moves is therefore consistent with the idea that repeated rainfall shocks create incentives to relocate, while cash transfers help households overcome the fixed costs of doing so.

The accumulated-shock result also clarifies why the effect of cash transfers cannot be summarized by one sign. For current rainfall shocks, the transfer appears to reduce migration, which is consistent with coping in place. For repeated rainfall shocks, the transfer can increase the probability of moving with part or whole household, which is consistent with a relocation response. The same policy instrument can therefore facilitate different forms of adaptation depending on whether the shock is temporary or repeated.

As an additional check, Table 2.10 reports whether accumulated rainfall shocks predict cash-transfer receipt. This is a stricter concern than in the case of one-year shocks because accumulated exposure may be correlated with more persistent local characteristics. The check is therefore useful for showing that the accumulated-shock interaction is not simply capturing the fact that households in repeatedly affected areas are more likely to receive transfers.

Table 2.10: Orthogonality Check: Accumulated Rainfall Shocks

	Cash Transfer
	(1)
Accumulated Shock	−0.0142
	(0.0097)
Observations	21114
Controls	*Yes*
Time Fixed Effects	*Yes*
Household Fixed Effects	*Yes*
Sub-district Fixed Effects	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates for the relationship between accumulated rainfall shocks and receipt of a cash transfer. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

2.6.2 Temperature Shocks

In this subsection, we move to another slow-onset event that threatens Indonesia, which is increasing temperature. We attempt to unpack whether a cash transfer distorts the behaviour of poor households as a response to drastic temperature shocks. Table 2.11 shows our baseline estimations that replicate the same models we estimated for rainfall shocks.

Table 2.11: Temperature Shocks, Cash Transfers, and Migration

	Migration and Temperature Shocks		Migration, CT, and Temperature Shocks	
	(1)	(2)	(3)	(4)
Temperature Shock	−0.0113	−0.0129	−0.0168*	−0.0184*
	(0.0091)	(0.0092)	(0.0101)	(0.0103)
Cash Transfer			0.0068	0.0060
			(0.0083)	(0.0083)
Temperature Shock=1 × Cash Transfer=1			0.0280	0.0284
			(0.0197)	(0.0193)
Observations	17541	17406	17541	17406
Controls	*No*	*Yes*	*No*	*Yes*
Time Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table presents the baseline regression results. The dependent variable is the indicator for all migration decisions. Columns (1) and (2) report estimates for the temperature shock specification without and with controls. Columns (3) and (4) report estimates from the main specification without and with controls. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

The temperature indicator is negatively associated with migration, while the cash-transfer interaction is positive but does not reach conventional levels of statistical significance. Because the indicator combines very hot years with a relatively lower-temperature comparison group, the coefficient should be interpreted as a difference across these temperature regimes rather than as an extreme-heat effect alone.

One possible interpretation is that temperature shocks affect household resources through a more persistent deterioration in working conditions and productivity. If high temperatures reduce income at the same time as they increase the cost of remaining in place, poor households may be pushed into a liquidity trap in which migration would be desirable but unaffordable. In such a setting, the positive sign of the interaction term is consistent with cash transfers partly relaxing this constraint. Although the estimate is not sufficiently precise to be interpreted as conclusive evidence, it is consistent with the broader mechanism that resources can turn environmental stress into actual mobility.

Table 2.12 shows the effect of accumulated temperature shocks and cash transfers. Although non-significantly, the main pattern remains as in the case of accumulated rainfall shocks: migration of the whole or part of the household is facilitated by a cash transfer as an adaptation strategy to a shock. The accumulated temperature results are less precise than the rainfall results, but they are still useful because they show that the di-

rection of the interaction is not inconsistent with the liquidity-constraint interpretation. If temperature shocks reduce the ability of poor households to finance migration, then additional resources may matter more for larger moves than for small individual moves. For this reason, we interpret the temperature results cautiously, but we retain them in the analysis because they represent a different form of slow-onset environmental stress.

Table 2.12: Accumulated Temperature Shocks, Cash Transfers, and Migration

	Migration					
	n=2		n=3		n=4	
	All	Family	All	Family	All	Family
Accumulated Shock	−0.0018	0.0048	−0.0034	0.0032	−0.0111	−0.0029
	(0.0107)	(0.0062)	(0.0105)	(0.0075)	(0.0127)	(0.0071)
Cash Transfer	−0.0005	0.0018	0.0020	0.0028	0.0029	0.0009
	(0.0084)	(0.0070)	(0.0089)	(0.0071)	(0.0100)	(0.0079)
Accumulated Shock × Cash Transfer	0.0314*	0.0005	0.0238	−0.0036	0.0139	0.0032
	(0.0162)	(0.0101)	(0.0177)	(0.0093)	(0.0174)	(0.0099)
Observations	17406	17406	17406	17406	17406	17406
Controls	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Time Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates from the main model. The outcome variables are all moves and moves undertaken by part or whole households. Columns (1) and (2) present accumulated shocks calculated over two periods. Columns (3) and (4) present accumulated shocks calculated over three periods. Columns (5) and (6) present accumulated shocks calculated over four periods. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

2.6.3 Natural Disasters

In this subsection, we examine how natural disasters influence the migration decisions in our data. Natural disasters are assumed to be unexpected events and are often used in different areas of research as a source of exogenous variation (Halliday (2006) Carvajal & Pereira (2010)). We incorporate the interaction term into the same model to understand if recipients of cash-transfer programs react to natural disasters differently than non-recipients. The results are presented in Table 2.13. Unlike the rainfall and temperature specifications, the chapter does not report parallel anticipation and transfer-orthogonality checks for the disaster measures.

Table 2.13: Natural Disasters, Cash Transfers, and Migration

	Migration					
	Flood		Landslide		Drought	
	All	Family	All	Family	All	Family
Event	−0.0066	−0.0086	−0.0029	−0.0059	0.0077	0.0065
	(0.0132)	(0.0058)	(0.0176)	(0.0132)	(0.0158)	(0.0109)
Cash Transfer	0.0047	0.0024	0.0049	0.0034	0.0042	0.0030
	(0.0070)	(0.0068)	(0.0071)	(0.0068)	(0.0070)	(0.0067)
Event × Cash Transfer	−0.0035	−0.0176	−0.0112	−0.0626***	0.0172	−0.0636*
	(0.0205)	(0.0205)	(0.0236)	(0.0207)	(0.0312)	(0.0347)
Observations	21114	21114	21114	21114	21114	21114
Controls	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Time Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates from the main model using three types of natural-disaster shocks. The outcome variables are all moves and moves undertaken by part or whole households. Columns (1) and (2) present results for floods. Columns (3) and (4) present results for landslides. Columns (5) and (6) present results for droughts. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

Table 2.13 shows that the interaction term has a mostly negative effect on migration decisions, meaning that poor households use additional money to adapt to the shock at the place they reside rather than move away. This is supported by evidence from the literature that records short-term movements out of affected areas after a natural disaster but returns to the area in the longer term. The findings of Table 2.13 also mimic the results from rainfall shocks (short-term).

The comparison with slow-onset shocks is useful here. Natural disasters are more visible and more sudden than rainfall or temperature anomalies, but their effect on long-term migration can still depend on whether households can recover locally. A household that loses assets in a flood or landslide may need immediate resources for repair and replacement. If these resources are available, the household may remain in the same area, even though it was affected by a serious shock. This is consistent with the negative sign of the interaction term and with the idea that cash transfers may reduce persistent migration after sudden events.

The similarity between short-term rainfall shocks and natural disasters is plausible because both can create immediate losses that require liquidity to cope with them. Affected households may need to repair housing, replace damaged assets, smooth food consump-

tion, or recover local sources of income. If the transfer helps cover these costs, it can reduce the need to move for a period long enough to be recorded in the IFLS migration module. This interpretation does not rule out temporary displacement, but it suggests that cash transfers may reduce more persistent migration responses after sudden shocks.

2.6.4 Change of Occupation

Next, we investigate another adaptation strategy that household use to cope with climate shocks - occupational change. We also attempt to uncover how the presence of a cash-transfer program influences occupation decisions and how households act when they are exposed to a climate shock and receive a transfer. We assume that agriculture is the sector that suffers from rainfall and temperature shocks the most, and we look into two outcome variables: change out of the agricultural sector and change into the agricultural sector, expecting the opposites signs from our interaction term. We consider stayers only (households that did not have migration recorded) because migration and change of occupation mostly occur simultaneously. Tables 2.14 and 2.15 examine the relationship between the change into and out of agriculture, cash transfers, and rainfall and temperature shocks.

Table 2.14: Rainfall Shocks, Cash Transfers, and Occupational Choices in Agriculture

	Out of Agriculture	Into Agriculture
	(1)	(2)
Rainfall Shock	−0.0054	0.0092
	(0.0098)	(0.0128)
Cash Transfer	0.0073	0.0180**
	(0.0082)	(0.0084)
Rainfall Shock × Cash Transfer	0.0346	−0.0426***
	(0.0282)	(0.0150)
Observations	8594	8594
Controls	*Yes*	*Yes*
Time Fixed Effects	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates from the model in which the outcomes are whether a household member changes occupation out of agriculture (Column 1) or into agriculture (Column 2). Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

Table 2.15: Temperature Shocks, Cash Transfers, and Occupational Choices in Agriculture

	Out of Agriculture	Into Agriculture
	(1)	(2)
Temperature Shock	−0.0161**	−0.0071
	(0.0065)	(0.0074)
Cash Transfer	0.0081	0.0175**
	(0.0079)	(0.0082)
Temperature Shock × Cash Transfer	0.0153	−0.0082
	(0.0145)	(0.0162)
Observations	8437	8437
Controls	*Yes*	*Yes*
Time Fixed Effects	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates from the model in which the outcomes are whether a household member changes occupation out of agriculture (Column 1) or into agriculture (Column 2). Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

Tables 2.14 and 2.15 indicate that receiving a cash transfer decreases the probability of household members changing their job into the risky agricultural sector after a shock. Conversely, although not as evident, the tables show that receiving a cash transfer increases the likelihood of household members leaving the agricultural sector. Both interaction terms are of the same sign for rainfall (Table 2.14) and temperature shocks (Table 2.15).

The occupational-choice results are important because they show that remaining is not necessarily passive. A household can remain in the same location but still reduce exposure to climate risk by changing the sector of employment. In this sense, the occupational results help interpret the negative migration response found for short-term rainfall shocks and natural disasters. If cash transfers allow households to avoid moving while also helping them reallocate away from risky activities, then lower migration does not necessarily mean lower adaptation. It may instead mean that adaptation takes place through a different margin.

These results should be interpreted as complementary evidence rather than as a full mechanism test. We observe whether household members move into or out of agriculture, but we do not observe all possible forms of local adaptation, such as changing crops, repairing housing, relying on informal networks, or adjusting consumption. Still, the agriculture outcomes are informative because it is one of the sectors most directly exposed to rainfall and temperature shocks.

The occupational results, therefore, complement the migration results. Migration is not the only way households can adjust to environmental stress; households that remain in place may still reduce exposure to climate risk by reallocating labor across sectors. The fact that the signs are similar for rainfall and temperature suggests that cash transfers may help households avoid increasing their dependence on a climate-sensitive sector after a shock, and may also help some members move away from agricultural employment.

2.6.5 Summary of Mechanism Evidence

Table 2.16 summarizes the mechanism evidence across shock types. The main lesson is that access to financial resources does not have a uniform effect on migration. Instead, it changes the way poor households respond to different types of environmental stress. Short-run shocks and natural disasters are associated with lower migration among

cash-transfer households, consistent with in-place adaptation. Persistent rainfall shocks and temperature anomalies are associated with higher migration among cash-transfer households, consistent with liquidity constraints being relaxed.

Table 2.16: Summary of Mechanism Evidence

Shock or outcome	Effect of cash transfers on response	Mechanism supported	Interpretation
Current rainfall shock	Reduces migration response	Climate-inhibitor	Transfers help households cope locally with short-run rainfall stress.
Accumulated rainfall shock	Increases family migration response	Climate-driver	Transfers help finance migration after repeated environmental stress.
Temperature anomaly	Increases migration response in the direction predicted by liquidity constraints	Climate-driver	Transfers relax mobility constraints when temperature stress reduces local viability.
Natural disaster	Reduces migration response	Climate-inhibitor	Transfers support recovery and reduce the need for longer-term relocation.
Occupational adjustment	Reduces entry into risky agriculture and increases exit from agriculture	In-place adaptation / risk reduction	Transfers help households adjust livelihoods without necessarily moving.

Notes: The table summarizes the interpretation of the main empirical results. The climate-driver mechanism means that financial resources increase migration by relaxing liquidity constraints. The climate-inhibitor mechanism means that financial resources reduce migration by supporting adaptation in place.

Taken together, the results show that migration is not the only form of adaptation and that lower migration is not necessarily evidence of lower climate pressure. For short-run shocks, cash transfers appear to reduce migration because they allow households to remain in place. For persistent shocks, cash transfers appear to increase migration because they make relocation feasible. This distinction is central for interpreting social protection in environmentally exposed settings.

2.7 Robustness Checks

We conduct several robustness checks to ensure that the patterns we observe in our baseline results are robust to several modifications. First, to strengthen our result, we also consider a second discrete measure of rainfall shock. For a sub-district s, we define its monthly rainfall anomaly as:

$$RA_{sm} = \frac{Rainfall_{sm} - \overline{Rainfall_s}}{SD_s}$$

where s indexes sub-districts, and m indexes months.

We define Rainfall Shock for a given year=Number of months in a given year with $|RA_{sm}| > 2$. Table 2.17 summarizes the results from our baseline and from the new measure of the shock.

Table 2.17: Alternative Rainfall Shock Measures, Cash Transfers, and Migration

	Migration			
	Rainfall Shock (Yearly Measure)		Rainfall Shock (Monthly Measure)	
	(1)	(2)	(3)	(4)
Rainfall Shock	−0.0093	−0.0090	−0.0023	−0.0022
	(0.0142)	(0.0141)	(0.0031)	(0.0031)
Cash Transfer	0.0084	0.0075	0.0126	0.0119
	(0.0077)	(0.0078)	(0.0085)	(0.0086)
Rainfall Shock × Cash Transfer	−0.0402**	−0.0391**	−0.0111	−0.0114*
	(0.0190)	(0.0191)	(0.0069)	(0.0069)
Observations	18440	18291	18440	18291
Controls	*No*	*Yes*	*No*	*Yes*
Time Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates from the main model using two different measures of rainfall shocks. Columns (1) and (2) use the baseline shock definition from the main results. Columns (3) and (4) use an alternative rainfall shock measure based on the number of months in the year with severe rainfall anomalies. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

As indicated in Table 2.17, even if we change the way we define rainfall shock, we still obtain the same pattern in the results. We still find that households that received a cash

transfer migrated less as a response to a current rainfall shock than households without a cash transfer.

In our second robustness check we include the earlier waves of the IFLS. We concentrate on information from the last 3 waves of the IFLS since cash-transfer programs were introduced only after 2000. However, to check the robustness of the main results on rainfall shocks, we present results with all waves in Table 2.18.

Table 2.18: Rainfall Shocks, Cash Transfers, and Migration: Extended Survey Window

	Migration	
	(1)	(2)
	Rainfall Shock, 2000-2014	Rainfall Shock, 1997-2014
Rainfall Shock	−0.0153	−0.0152
	(0.0126)	(0.0111)
Cash Transfer	0.0122	0.0141*
	(0.0078)	(0.0078)
Rainfall Shock × Cash Transfer	−0.0353*	−0.0339*
	(0.0205)	(0.0205)
Observations	17850	20545
Controls	*Yes*	*Yes*
Time Fixed Effects	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*

Standard errors in parentheses
Sub-district Clustered SE
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates from the main model for two survey windows. Column (1) reports the baseline results for 2000–2014. Column (2) extends the period to 1997–2014 by using one additional wave of the IFLS survey. Controls include household size, the share of mobile household members, assets, Muslim households, and urban indicators.

Again, our result survives the robustness check and we find the same pattern and even the same magnitude of the effect of a cash transfer on migration decisions after a rainfall shock. Therefore, we conclude that our result for a short-term rainfall shock is stable and does not change with modifications.

The two robustness checks address different possible concerns. The alternative rainfall measure checks whether the results depend on the annual dummy definition of a shock. The extended survey window checks whether the estimates are sensitive to excluding the earlier IFLS wave. Since the main coefficient of interest remains similar across these

specifications, the interpretation of the short-term rainfall result does not appear to be driven by one narrow measurement choice or by one particular selection of survey years.

Next, we conduct a battery of robustness checks for accumulated shocks. We vary the number of years for which a shock is calculated as well as the threshold. This is important because accumulated exposure is not directly observed as a single event. It has to be constructed from a sequence of rainfall anomalies, and the empirical result should not depend on one arbitrary choice of the number of years or cutoff. The results are presented in Table 2.19 and show that the pattern we observe for our baseline results is also present in our robustness checks for different modifications of the model.

Table 2.19: Accumulated Rainfall Shocks, Cash Transfers, and Migration

	Migration									
	n=2		n=3		n=4, cutoff=5		n=4, cutoff=4.5		n=4, cutoff=5.5	
	All	Family	All	Family	All	Family	All	Family	All	Family
Accumulated Shock	−0.0010	−0.0004	−0.0019	0.0087	−0.0234**	0.0066	−0.0075	0.0052	−0.0178**	0.0015
	(0.0101)	(0.0066)	(0.0079)	(0.0063)	(0.0093)	(0.0071)	(0.0060)	(0.0051)	(0.0089)	(0.0084)
Cash Transfer	0.0065	−0.0163*	0.0078	−0.0162*	0.0066	−0.0163*	0.0086	−0.0088	0.0068	−0.0164*
	(0.0076)	(0.0091)	(0.0078)	(0.0094)	(0.0077)	(0.0095)	(0.0079)	(0.0100)	(0.0076)	(0.0094)
Accumulated Shock × Cash Transfer	−0.0128	0.1189***	−0.0192	0.1045**	−0.0124	0.0997**	−0.0164	0.0404	−0.0137	0.1059***
	(0.0129)	(0.0445)	(0.0130)	(0.0406)	(0.0132)	(0.0389)	(0.0112)	(0.0283)	(0.0126)	(0.0405)
Observations	21114	21114	21114	21114	21114	21114	21114	21114	21114	21114
Controls	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Time Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Household Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Sub-district Fixed Effects	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*

Standard errors in parentheses

Sub-district Clustered SE

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports coefficient estimates from the main model. The outcome variables are all moves and moves undertaken by part or whole households. Columns (1) and (2) present accumulated shocks calculated over two periods. Columns (3) and (4) present accumulated shocks calculated over three periods. Columns (5)–(10) present accumulated shocks calculated over four periods using alternative cutoffs. Controls include household size, the share of mobile household members, assets, Muslim household, and urban indicators.

The robustness checks in Table 2.19 support the main interpretation of accumulated rainfall shocks. Across the alternative definitions, the estimates continue to show that repeated rainfall exposure differs from a one-year rainfall shock. The relevant migration response is concentrated more strongly in moves undertaken by part or whole households, which is the type of move that requires more resources. This pattern is consistent with the idea that repeated shocks make migration a more serious adaptation option, while cash transfers make this option more feasible for poor households. The result is therefore not driven only by the baseline threshold used in Table 2.9.

2.8 Discussion and Policy Implications

The results have several implications for how to interpret migration as an adaptation strategy. Migration should not be viewed only as a sign of vulnerability or only as a sign of successful adaptation. In some cases, migration reflects distress: households move because local conditions become difficult to sustain. In other cases, migration reflects capacity: households move because they finally have the resources to finance relocation. The same observed outcome can therefore have different meanings depending on the constraints facing the household.

The findings show that cash transfers can affect both the need to migrate and the ability to migrate. After short-term rainfall shocks and natural disasters, transfers reduce migration, which is consistent with in-place adaptation. Households may use the transfer to repair losses, smooth consumption, or avoid costly displacement. Under accumulated rainfall exposure, the positive interaction is concentrated in family moves; for the temperature indicator, the interaction is positive but statistically imprecise. In this case, the transfer may help households choose migration when remaining becomes less attractive.

This dual role is important for policy. Cash transfers are usually discussed as poverty-alleviation tools, but in environmentally exposed settings they can also shape adaptation choices. A program that reduces migration after one type of shock may increase migration after another type of shock, and both responses can be consistent with welfare-improving adaptation. The policy implication is therefore not that social protection should aim to reduce migration or to promote migration in all cases. Rather, social protection can expand the set of feasible choices available to poor households.

The distinction between temporary and persistent environmental stress is also important for policy design. A one-year rainfall anomaly or a single disaster may lead households to prefer recovery in place. Repeated rainfall shocks or temperature anomalies may instead change expectations about future livelihoods. Policies that support adaptation should therefore account for the duration and perceived persistence of environmental stress. Short-run relief may be most useful after sudden shocks, while longer-term support may need to include mobility, retraining, or relocation assistance.

There are also important limitations. The migration measure captures moves lasting longer than six months. This means that short seasonal moves, temporary labor migration, or very short displacement episodes are not fully observed. If households respond to environmental shocks through short-term migration, the estimates may understate the full mobility response. The results should therefore be interpreted as effects on more durable migration behavior.

Another limitation is that cash-transfer receipt may be correlated with time-varying household conditions. Our analysis uses household fixed effects and focuses on poor households, which helps improve comparability, but the estimates should still be interpreted cautiously. The interaction between cash transfers and climate shocks is informative about how migration responses differ by access to resources, but it is not always equivalent to a fully randomized experiment.

Finally, the results are specific to poor households in Indonesia. Indonesia is a useful setting because of its exposure to rainfall variability, temperature anomalies, and natural disasters, and because of the availability of cash-transfer programs. However, the mechanisms may differ in settings with different migration costs, labor markets, social protection systems, or exposure to environmental risks. The broader lesson is that financial resources change the set of adaptation strategies available to poor households, but the direction of the migration response depends on the type of environmental shock.

2.9 Conclusion

This paper studies how access to financial resources changes the migration response of poor households to environmental shocks. The main question is whether cash transfers help households migrate after climate stress by relaxing liquidity constraints, or instead

help them remain in place by supporting local adaptation.

Using the Indonesian Family Life Survey, climate data, natural-disaster information, and variation in government cash-transfer receipt, the analysis shows that both mechanisms are present. Cash transfers reduce migration responses to current-year rainfall shocks and natural disasters, consistent with the climate-inhibitor mechanism. Under accumulated rainfall exposure, the positive interaction is concentrated in the family-move outcome, consistent with the climate-driver mechanism operating through a more costly relocation margin. The temperature interaction is positive but statistically imprecise.

The paper also shows that adaptation occurs along margins other than migration. Occupational results suggest that cash-transfer households are less likely to enter risky agricultural work and more likely to leave agricultural occupations in response to rainfall and temperature shocks. This indicates that households may use financial resources to reduce exposure to climate-sensitive livelihoods even when they do not migrate.

These findings contribute to the literature on climate migration by showing that financial resources do not have a single uniform effect on mobility. Resources can help households remain in place after temporary shocks, but they can also help households leave after more persistent environmental stress. This helps reconcile mixed findings in the existing literature and highlights the importance of distinguishing between different types of climate shocks.

Several extensions remain for future work. One extension is to study shorter-term and seasonal migration, which is not fully captured by the IFLS migration measure used here. It would also be useful to examine whether households that migrate after shocks are better off in terms of consumption, employment, or income stability. Finally, future work could study heterogeneous effects by gender, age, household composition, and the gender of the household head.

Overall, the paper shows that financial constraints are important for understanding environmental migration. Poor households do not respond to climate shocks only according to the severity of the shock; they respond according to the options available to them. Cash transfers change those options. Sometimes they allow households to stay, and sometimes they allow households to leave.

Chapter 3

Weather Shocks and Bride Kidnapping in Kyrgyzstan[1]

3.1 Introduction

Climate change is one of the most pressing global challenges of our time, with significant implications for the environment, the economy, and human well-being (IPCC, 2022). Developing countries are particularly vulnerable to the impacts of climate change, given their limited resources, weak infrastructure, and high dependence on natural resources (Hsiang & Kopp, 2018; IPCC, 2022). Therefore, understanding the implications of climate change for developing countries is crucial for policymakers, scholars, and practitioners alike.

Despite nature's impartiality, the consequences of climate change are not gender-neutral either. This is primarily attributed to discriminatory social norms, pre-existing gender inequalities in society, and unequal power relations between women and men (Rao et al., 2023; Reggers, 2019; Terry, 2009). Such factors make women more susceptible to the consequences of climate change (Oswald Spring, 2019; Reggers, 2019).

One of the ways women can be adversely affected by climate change is through in-

[1]**Co-authored with Ivan Trestcov (CERGE-EI)**

creased incidents of harmful marriage practices. Banerjee et al. (2013) argue that the degree to which a country's marriage institutions favor men, such as patrilocality and concern for girl's "purity," can help explain low female labor-force participation, reduced investment in female human capital, and a higher tolerance for gender-based violence. The economic literature already contains a lot of research on the effects of various climate shocks on one particular practice - child marriage (Pope et al., 2022). This literature shows that the effect of shocks on child marriage varies significantly depending on cultural contexts and norms. In societies where a bride price is customary (e.g. Vietnam and countries in Sub-Saharan Africa), such shocks can raise the likelihood of child marriage as families may consider selling their daughters to offset economic losses. Conversely, in societies where dowry is the norm (e.g. India), shocks can decrease the probability of child marriage as the girls' families may not have the resources to meet the requirements for the dowry (Corno et al., 2020; Tsaneva, 2020; Trinh & Zhang, 2021).

This paper studies another harmful marriage practice - bride kidnapping or marriage by capture, in which a man abducts the woman he wishes to marry. According to a report by the United Nations, several countries, including Kyrgyzstan, Kazakhstan, Armenia, Ethiopia, Georgia, Moldova, and Uzbekistan, among others, have documented cases of bride kidnapping (UN Women, 2018). Bride kidnapping poses significant dangers to the victims, including severe physical, psychological, and social consequences. It can result in forced marriage, sexual and physical violence, unwanted pregnancies, social isolation, and mental health issues. The United Nations considers bride kidnapping a form of violence against women and girls that infringes upon their human rights and can negatively affect their physical and mental well-being (UN Women, 2018). Despite being detrimental, bride kidnapping receives little to no attention in the larger gender economics literature.

Despite sharing the unfortunate label of harmful marriage practices, bride kidnapping presents a fundamentally different dynamic than child marriage. While child marriage often involves economic transactions where parents receive payments or benefits from the other side, bride kidnapping eliminates any such transaction. In the case of bride kidnapping, the abductor, the prospective husband, forcibly takes the woman he intends to marry without offering any form of compensation to her family. This absence of a transactional element accentuates the coercive nature of bride kidnapping, making it not just a financial arrangement but a stark violation of a woman's autonomy and consent. This distinction is essential as it influences societal attitudes and impacts the victims

differently. Understanding these unique features of bride kidnapping enriches the analysis of its correlation with climate shocks, thereby enhancing the depth of the discussion on gender-specific repercussions of climate change. This distinction also implies a different economic mechanism through which climate shocks may affect marriage practices. In the bride-price literature, negative shocks can increase child marriage because the bride's household receives a payment. The mechanism considered in this chapter operates on the other side of the marriage transaction. A dry winter can tighten household resources and make the costs associated with conventional marriage, including kalym and other marriage-related expenditures, more difficult to finance for the groom and his household. Because bride kidnapping can reduce some of these costs, it may become relatively more attractive under economic stress. We therefore focus on financial constraints and marriage costs as the main economic channel linking dry winters to bride kidnapping.

Our study aims to contribute to this gap by examining how climate shocks affect the composition of observed marriages and attitudes toward bride kidnapping in Kyrgyzstan - a landlocked country in Central Asia that is highly vulnerable to climate change and also has the highest rates of bride kidnapping in the world (United Nations Population Fund (UNFPA), 2018). Specifically, we examine whether dry winters increase the probability that an observed marriage was formed through bride kidnapping rather than through love or arranged marriage. We complement this analysis by examining whether dry winters are associated with changes in attitudes toward bride kidnapping relative to love marriage. This aspect of our study contributes to the literature, as previous studies did not study the question due to a lack of data.

Furthermore, our contribution extends beyond the realms of gender economics and environmental literature as it holds relevance for the general public and policymakers. The issue of bride kidnapping is particularly poignant in Kyrgyzstan, where it raises questions about whether it is a crime or a tradition to be respected. Recent incidents, such as the tragic murders of Burulai Turdaaly Kyzy in 2018[2] and Aizada Kanatbekova in 2021[3], who were both killed by their abductors due to neglect from police and government, have garnered significant attention from both national and international media and led

[2]retrieved from https://www.hrw.org/news/2018/05/31/young-womans-murder-kyrgyzstan-shows-cost-tradition

[3]retrieved from https://centralasiaprogram.org/bride-kidnapping-kyrgyzstan-reflection-conservative-values#_ftn1

to eruptions of protest in the largest cities of Kyrgyzstan. (see Figure 3.1). However, we must also consider the countless silent victims who suffer mentally and possibly physically due to the persistence of this harmful practice. By shedding light on these connections policymakers can incorporate this knowledge when formulating policies to mitigate the consequences of climate change and create a safer, more equitable environment for women.

We use monthly weather data on precipitation and a panel survey from Kyrgyzstan to answer our research question. The identification strategy employed in this study capitalizes on the exogenous and random variation in rainfall/snowfall patterns, specifically focusing on the occurrence of dry winters. In recent years, Kyrgyzstan's agriculture has been severely affected by climate shocks, particularly due to unprecedented droughts and a lack of irrigation water. One of the primary cause is insufficient snowfall during winter, which leads to reduced meltwater and impedes river flow in the spring-summer period, failing to fill regional reservoirs adequately (BIOM, 2021). Snowfall is crucial for replenishing groundwater, nourishing soil, and protecting winter crops from harsh conditions. The Ministry of Agriculture reports that inadequate snowfall has resulted in significant losses of 10% to 15% in winter crop yields. By leveraging variation in the timing of rainfall and snowfall shocks, we examine whether dry winters are associated with how observed marriages were formed and with changes in attitudes toward bride kidnapping. We find that dry winters increase the share of observed marriages formed through bride kidnapping and are associated with a relative shift in attitudes in favor of bride-kidnapping marriages compared with love marriages. Because the attitude measure is the difference between ratings of the two marriage types, it does not identify which underlying rating drives this change.

Our analysis reveals the intricate interplay between climate change and bride kidnapping in Kyrgyzstan, highlighting the nuanced nature of the consequences resulting from climate change. The relationship between dry winters, socio-economic factors, and attitudes towards bride kidnapping is complex and multi-faceted. One of the key findings is the crucial role of income as a determinant, as individuals from lower-income backgrounds display a higher susceptibility to attitude shifts. This suggests that economic vulnerability may exacerbate the impact of climate shocks on societal perceptions and behaviors.

Gender differences did not emerge as statistically significant; the study indicates a similar response to climate shocks among both males and females. We do not find strong or

systematic evidence that the attitudinal response differs by education level. This implies that the influence of climate-related factors transcends gender and educational boundaries and affects the attitudes of individuals across the spectrum. A higher daughters' share is associated with a smaller attitudinal shift following a dry winter, suggesting the potential influence of familial dynamics in perpetuating or challenging harmful practices.

The rest of the paper has the following structure. Section 3.2 describes the climate specifics of Kyrgyzstan and the bride kidnapping practice. Section 3.3 summarises the data. Section 3.4 justifies the empirical strategy of the paper. Section 3.5 presents the results of the analysis. Section 3.6 concludes.

3.2 Background

3.2.1 Climate in Kyrgyzstan

Kyrgyzstan, a mountainous country in Central Asia, faces various climate change challenges, including melting glaciers, changes in precipitation patterns, and increased frequency of extreme weather events such as floods and droughts. According to the National Communication of the Kyrgyz Republic to the UNFCCC (Ministry of Energy and Industry of the Kyrgyz Republic, 2016), "climate change is expected to have significant impacts on key sectors in Kyrgyzstan, such as agriculture, water resources, health, energy, and infrastructure." Climate change can significantly impact the country's economy, which relies heavily on agriculture and livestock farming. Moreover, it can exacerbate existing social and economic challenges, including poverty, food insecurity, and gender inequality.

According to the World Bank (2021), Kyrgyzstan faces several challenges related to water resources, including water scarcity, inadequate infrastructure for irrigation, and water pollution. These issues have significant implications for agriculture, a critical sector of the Kyrgyz economy. Agriculture accounts for approximately one-third of Kyrgyzstan's GDP and employs more than half of the country's workforce (World Bank, 2021). The Climate Risk Index 2022 report by Germanwatch (2022) states that Kyrgyzstan is ranked among the countries most affected by extreme weather events from 2001 to 2020. One of the main factors contributing to this status is changes in precipitation patterns, particularly the increasing frequency and intensity of droughts and floods. We focus on

precipitation shocks as the key variables of interest.

In recent years, agriculture in Kyrgyzstan has suffered a lot due to climate shocks. The country's agricultural lands faced an unprecedented drought and lack of irrigation water. Farmers attended numerous rallies and demanded that the authorities take action to preserve the crops.[4] The situation with insufficient water supply is majorly attributed to the lack of snow during the winter, according to BIOM (2021), a public non-profit voluntary organization unifying young specialists, scientists, and leaders who participate in addressing environmental problems of the Kyrgyz Republic and Central-Asian region. In Kyrgyzstan, water influx in river systems largely relies on seasonal snow accumulation. Recent observations reveal a trend of reduced snowfall during winter seasons. Consequently, insufficient meltwater from seasonal snow cover has impeded river flow during the spring-summer period, leading to a shortage of water volume required to fill regional reservoirs to desired levels (BIOM, 2021). In addition, sufficient precipitation plays a pivotal role in the agricultural sector by serving as a vital water resource for crop cultivation. Snowfall plays a significant role in nourishing the soil, replenishing groundwater reserves, and fostering optimal plant growth. Another crucial aspect of winter snowfall in Kyrgyzstan is its protective capacity, shielding winter crops from harsh cold temperatures and gusty winds while facilitating adequate soil aeration owing to its loose composition. Data from the Ministry of Agriculture in Kyrgyzstan reveals that inadequate snowfall in some years resulted in notable losses ranging from approximately 10% to 15% in winter crop yield. [5] Based on the information above, we use the variable that indicates insufficient snowfall in winter compared to the historical mean (or dry winter) as our climate shock variable.

3.2.2 Bride Kidnapping in Kyrgyzstan

One particular issue of gender inequality in Kyrgyzstan is bride kidnapping, a form of forced marriage involving the abduction of young women by men who wish to marry them without their consent. Bride kidnapping is known as "ala kachuu," which translates as "to take and run away." Bride kidnapping is a widespread practice in Kyrgyzstan, and it

[4]retrieved from https://economist.kg/novosti/2021/08/11/zasuha-i-potopy-chto-proishodit-s-klimatom-kyrgyzstana-i-pochemu-dalshe-budet-huzhe/?ysclid=ldcz4v41x93397284

[5]retrieved from https://ru.sputnik.kg/20190219/kyrgyzstan-zima-posledstviya-chnergetika-poliv-1043407465.html/

has serious social, economic, and psychological consequences for women and their families (UN Women, 2018). According to CSCE (2017) data, an estimated 12000 young women are kidnapped and forced to marry their abductors yearly in Kyrgyzstan. As many as one out of five are raped in the process. The National Statistical Committee in 2017 estimated that up to 20% of marriages in Kyrgyzstan result from bride kidnapping[6], and United Nations Population Fund (UNFPA) (2018) estimates this number to be as high as 35%.

Bride kidnapping became illegal in 1994, but the practice continues today, especially in rural areas (CSCE, 2017). The laws forbidding bride kidnapping in Kyrgyzstan have been largely ineffective due to cultural acceptance of the practice - Kleinbach et al. (2005) reported that 38% of the population considers bride kidnapping a "good traditional way to get a bride.". Despite the prevalence of this practice in Kyrgyzstan, only 79 incidents were officially reported to the Ministry of Internal Affairs between 2014 and 2016 (Bengard, 2017).

Traditionally, young women can reject marriage by abduction, but there are strong incentives against it. Becker et al. (2017) show that women who decline a proposal of kidnapping are less likely to find an appropriate partner with equal social and economic status afterward due to women's reputations being compromised by spending time at the abductor's house. It is also considered shameful for the woman's household if she refuses the proposal (Kleinbach et al., 2005; Naumova, 2016). The girl's family may pressure her to agree to the marriage and comply with social norms. While bride kidnapping occurs in various parts of the world, it has been reported to be more prevalent in Central Asia compared to other regions.

Although we do not know the exact reason why the practice is so persistent in this region, there are several factors that contribute to the higher occurrence of bride kidnapping. First, cultural and historical factors: bride kidnapping has deep cultural roots in Central Asia, particularly among certain ethnic groups. It has been practiced for centuries and is often associated with local traditions, customs, and notions of honor. Central Asia has a history of nomadic traditions, where male dominance and patriarchal norms were prevalent (Becker, 2019). These traditions often included bride kidnapping as a means for men to assert their power and control over women, perpetuating the prac-

[6]retrieved from https://24.kg/obschestvo/45809_kraji_nevest_vkyirgyizstane_kakim_issledovaniyam_verit/

tice over generations. The cultural acceptance and historical continuity contribute to its persistence in the region. Second, socioeconomic conditions, such as poverty and limited access to education and employment opportunities, can contribute to the prevalence of bride kidnapping. In some cases, men resort to kidnapping because they lack the financial means to afford a traditional marriage dowry or cannot meet the expectations set by the bride's family. Third, weak law enforcement and inadequate legal frameworks to address bride kidnapping contribute to its persistence. In some cases, local authorities and communities may turn a blind eye to the practice, viewing it as a part of tradition or considering it a private matter between families. Lastly, limited awareness about women's rights and gender equality can perpetuate the practice of bride kidnapping. Insufficient education and awareness campaigns addressing the negative consequences and human rights violations associated with this practice hinder progress in combating it effectively (UN Women, 2018).

3.3 Data

To answer the research question we use two data sources. The first is the Life in Kyrgyzstan Survey (LiK)[7], an ongoing longitudinal survey in Kyrgyzstan. The survey follows 3000 households and more than 8000 individuals across all seven regions (oblasts) of Kyrgyzstan, as well as the two administratively separate cities of Bishkek and Osh. This creates a nationally representative dataset encompassing both rural and urban areas and Northern and Southern regions. The sample was initially selected based on the 2009 national population census, and the survey was conducted six times: first in the autumn of 2010 and then repeated in 2011, 2012, 2013, 2016, and 2019. The LiK survey is suitable for our study since it contains information about how marriages happened (love marriage, arranged, kidnapping) and attitudes towards bride kidnapping. In addition to information about bride kidnapping, LiK also has data on the socio-economic and demographic characteristics of households and individuals in Kyrgyzstan, which allows us to investigate possible channels as well as other detrimental effects that climate shocks have on women.

Many datasets provide information about precipitation and temperature. Examples

[7]retrieved from https://datasets.iza.org/dataset/124/life-in-kyrgyzstan-study-2010-2019

include the dataset produced by the Climatic Research Unit at the University of East Anglia, weather data obtained from the Center for Climatic Research at the University of Delaware, and NASA MERRA-2. These datasets provide information about precipitation and temperature on grids of 0.5 by 0.5 degrees, corresponding to about 50 by 50 kilometers. We use weather data obtained from the Climatic Research Unit at the University of East Anglia[8], particularly data about temperature and precipitation, which we use to investigate the effect of insufficient precipitation on marriage practices.

After combining the data about precipitation and temperature with different waves from LiK, we now have the individual-year panel dataset, which contains information about individuals in Kyrgyzstan with several socio-economic characteristics and the history of climate shocks experienced each year. Table 3.1 below shows the summary statistics of our panel dataset.

In our analysis, we extract precipitation data on a monthly scale to account for droughts during the survey period. We define a dry winter based on precipitation during December, January, and February. The indicator equals one if average winter precipitation is more than 1.5 standard deviations below the district's historical winter mean. Put simply, our dry winter variable equals one for a given year in a given district if there was too little snow compared to what is the historically normal amount of snow in that district. Firstly, snowfall in a particular district by itself influences the amount of water in the ground in that district. Secondly, in Kyrgyzstan, the implementation of irrigation systems is primarily district-based rather than nationwide (World Bank, 2021). The country has a diverse landscape, with variations in water availability, climate, and agricultural practices across different regions. As a result, irrigation systems are developed and managed at the district level to cater to each area's specific needs and conditions. This characteristic of the irrigation system in Kyrgyzstan is one of the reasons it suffers from a lack of coordination in addressing challenges such as aging infrastructure and inadequate investments in irrigation infrastructure modernization. However, it allows us to build our identification by matching winter precipitation to districts and creating the Dry Winter variable for individuals living in these districts.

The LiK survey provides information on how respondents' marriages were formed: through love marriage, arranged marriage, or bride kidnapping. We define the bride-

[8]retrieved from https://www.uea.ac.uk/groups-and-centres/climatic-research-unit

Table 3.1: Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Max
HH Members	42,789	6.28	2.60	1	20
Dry Winter	42,789	0.11	0.31	0	1
Age	42,789	37.35	16.87	15	105
Female	42,789	0.51	0.50	0	1
Rural	42,789	0.77	0.42	0	1
Married	42,789	0.61	0.49	0	1
Single	42,789	0.27	0.44	0	1
Year	42,789	2,013.53	3.08	2,010	2,019
Year Married	2,017	1,998.27	9.82	1,981	2,013
Age at Marriage	2,017	20.79	2.87	14	36
Husband Age at Marriage	2,017	24.36	4.14	16	99
Bride Kidnapped	2,017	0.14	0.35	0	1
Lost Love	11,622	−0.88	1.46	−3	3
Domestic Violence	11,622	−0.68	1.22	−3	3
Burdensome Demands	11,622	−0.63	1.17	−3	3
Cut From Friends	11,622	−0.50	1.17	−3	3

Notes: The table outlines summary statistics for all important variables used in the regressions. Dry winter is a dummy variable indicating that the amount of precipitation in the winter of the year of marriage was lower than the historical mean by 1.5 standard deviations. The marriage-history statistics reported in this table are based on the narrower Women and Fertility subsample for which the detailed first-marriage variables are jointly available, yielding 2,017 women. This descriptive subsample differs from the harmonized marriage-history sample used in the main bride-kidnapping regression. Each woman enters the main regression at most once. The last four variables are measures of attitude toward bride kidnapping evaluated as self-reported likelihoods of events (Lost Love, Domestic Violence, Burdensome Demands, and Being Cut Off from Friends) happening in marriage initiated with bride kidnapping. The attitude outcome is the bride-kidnapping score minus the corresponding love-marriage score and ranges from -3 to 3.

kidnapping outcome for observed marriages as an indicator equal to one if the marriage occurred through bride kidnapping and zero if it was a love or arranged marriage. The analysis therefore estimates the probability that an observed marriage was formed through bride kidnapping, conditional on a marriage taking place, rather than the unconditional probability that a woman experiences bride kidnapping. We restrict the sample to women to avoid double-counting the same marriage.

Because marriage type is self-reported, we cannot rule out that economic distress af-

fects whether a consensual elopement is reported as bride kidnapping; if such reporting changes following dry winters, the estimated effect may partly reflect changes in classification.

Marriage-history information is not collected in an identical form in every LiK wave. The detailed Women and Fertility module is available in the 2011–2013 waves and is not re-administered when the relevant information was already provided in an earlier wave. Marriage information is therefore harmonized across the available survey waves for the main analysis. We retain one observation per woman, using the earliest available marriage record. The resulting marriage-history dataset contains no repeated individual identifiers. Table 3.2 uses the 3,744 observations from this dataset for which the variables required by the regression are available.

For historical marriages, we do not observe the woman's district of residence at the time of marriage. We therefore match weather in the reconstructed marriage year to the district recorded in her earliest available LiK observation. If a woman moved between marriage and the survey, this may lead to measurement error in historical weather exposure.

In the last two waves of LiK there is also a section that asks about attitudes toward bride kidnapping. Respondents have to rank the following from 1 (very possible) to 4 (impossible) for bride kidnapping and love marriage: Separation from someone who one really loved/hoped to marry; Domestic violence and abuse by a husband toward wife; Burdensome demands on wife, e.g., work/serve husband's family; Be cut off from friends. To account for a general shift in attitudes toward marriage, we define our measure of the acceptability of bride kidnapping as the difference, or "gap," between the scores assigned to bride kidnapping and love marriage within each category. This gap can range from -3 to 3, with a larger gap indicating a greater acceptance of bride kidnapping relative to love marriage. For example, a gap of 3 for the outcome of domestic violence means that a respondent ranked the likelihood of domestic violence as "impossible" in the context of marriage by bride kidnapping but as "very possible" for love marriage. This suggests that the respondent considers bride kidnapping to be safer than love marriage in terms of how likely a woman will experience domestic violence. On the other hand, a gap of -3 means that a respondent ranked the likelihood of domestic violence as "very possible" in the case of a marriage by bride kidnapping but as "impossible" for love marriage. We use these raw gap scores, ranging from -3 to 3, as the attitude outcomes. We link the

residence of respondents who answer attitudes questions to the dry-winter shocks in the year the survey was conducted.

3.4 Identification Strategy

The amount of rainfall/snow in a particular location is exogenous and varies over time. Therefore, the identifying variation in our study arises from the random timing of the shocks of dry winters. We estimate two models. First, using a sample of women, we estimate the probability of the Marriage by bride kidnapping BK_{idt} of a woman i living in location d (a district in Kyrgyzstan recorded in the woman's earliest available LiK observation) at a year when the marriage happened t as follows:

$$BK_{idt} = \alpha + \beta * DryWinter_{idt} + \gamma * X_{idt} + \mu_t + \mu_d + \epsilon_{idt}$$

where i indexes women, d indexes districts, and t indexes the year of marriage. X_{idt} is a vector of the control variables of the woman, μ_t and μ_d are time and district fixed effects. BK_{idt} is the bride kidnapping variable, which equals one if the marriage happened through kidnapping, and zero if it is a love/arranged marriage. For this model, we use a dataset of women that contains information on how the marriage happened. The estimation dataset retains a single marriage-history observation for each woman; responses from repeated LiK interviews are not treated as separate marriage observations. As marriages are one-time events, we use district fixed effects rather than individual ones in this analysis.

We match the dry-winter shock to the year of marriage because winter precipitation affects water availability and agricultural conditions during the subsequent spring and summer; the LiK data do not provide the month of marriage, so the exact within-year timing cannot be identified.

Our second model is the following:

$$ATT_{idt} = \alpha + \beta * DryWinter_{idt} + \gamma * X_{idt} + \mu_t + \mu_i + \epsilon_{idt}$$

where i indexes individuals, d indexes districts, and t indexes the year of the interview.

X_{idt} is a vector of the control variables of the individual, μ_t and μ_i are time and individual fixed effects. The outcome variable ATT_{idt} is a variable that captures attitudes

of individual i towards bride capture - the larger the ATT_{idt}, the more accepting is individual i of the bride capture tradition. This regression analysis is done on the set of all individuals that responded to the bride kidnapping attitudes questions in the survey. The variable $DryWinter_{idt}$ is a dummy that equals one if district d was hit by climate shock, zero if not - this is our variable of interest in both models. We cluster standard errors on the district level to allow serial correlation across individuals in the same area. We observe 112 districts in our datasets.

3.5 Results

In this section, we present the main results of our study. We first examine the overall impact of climate shocks, specifically inadequate winter precipitation, on incidents of bride kidnapping and associated attitudes. Subsequently, we delve into a more nuanced analysis by considering different sub-samples and interaction terms including gender, compositions of children, education, and income levels.

3.5.1 Occurrences of Bride Kidnapping

The dependent variable in this specification is defined among observed marriages. Accordingly, the estimated effect should be interpreted as a change in the probability that a marriage is formed through bride kidnapping rather than through love or arranged marriage. It does not identify the effect of a dry winter on the unconditional probability of marriage or bride kidnapping in the full population of women. Table 3.2 demonstrates the effect of dry winters on the probability that an observed marriage is a bride-kidnapping marriage in districts of Kyrgyzstan. Column 1 reports the estimated coefficient of the regression without control variables, and Column 2 does so with control variables. Since the data on the timing of marriage is recall data, we have limited information about the socioeconomic characteristics of the woman in the year when the marriage happened. The information we do have and include in the control variables is whether a woman was living in a rural or urban area, her age at marriage, and the age of her groom at marriage. We include district and time-fixed effects to control for unobserved heterogeneity across different time periods and districts. The standard errors are clustered at the district level.

Our findings suggest that insufficient snowfall during winter increases the incidents of

Table 3.2: The Impact of the Climate Shock on Bride-Kidnapping Rate

	Bride Capture	Bride Capture
	(1)	(2)
Dry Winter	0.104**	0.106**
	(0.051)	(0.051)
Controls	N	Y
Year FE	Y	Y
District FE	Y	Y
Observations	3,744	3,744

*p<0.1; **p<0.05; ***p<0.01

Notes: Bride Capture is a dummy variable equal to one if the observed marriage was formed through bride kidnapping and zero if it was a love or arranged marriage. Among marriages not exposed to a dry-winter shock, the average share formed through bride kidnapping is 16.5%. Dry Winter is a dummy variable equal to one if winter precipitation in the respondent's district in the year of marriage was more than 1.5 standard deviations below the district's historical winter mean. Year and district fixed effects are included. Standard errors are clustered at the district level. Controls include rural/urban residence and the bride's and groom's ages at marriage. The estimation sample contains one observation per woman.

bride kidnapping over the number of love/arranged marriages in Kyrgyzstan districts. If the district experiences the shock of dry winter, it is associated with a 10-11 percentage point increase in the probability of marriage through bride kidnapping for women. The effect we find is therefore substantial relative to the observed baseline share. Because the outcome is defined among observed marriages, the estimate should be interpreted as a change in the composition of marriages rather than as an equivalent increase in the unconditional incidence of bride kidnapping.

As a placebo test, we examine whether dry winters occurring one or two years after the marriage year predict how the marriage was formed. Neither future dry-winter measure is statistically significant, providing no evidence that future weather shocks predict past marriage outcomes (Table 3.A3).

We also note that according to the World Bank, the percentage of the rural population in Kyrgyzstan is 63% [9], while in our sample, it is 77%. Since the rural population is oversampled and bride capture is more attributed to the rural areas, this might suggest a lower average rate of bride kidnapping and a smaller effect of dry winter for the overall population of Kyrgyzstan.

The positive relationship is consistent with the broader literature showing that economic shocks can affect harmful marriage practices, but the mechanism in our setting differs from the bride-price channel emphasized in studies of child marriage. In our setting, the relevant economic channel is that financial stress can make conventional marriage costs more difficult to finance, while bride kidnapping can reduce some of these costs for the groom's household.

3.5.2 Attitudes toward Bride Kidnapping

The results above show that dry winters increase the share of observed marriages formed through bride kidnapping. We next examine whether dry winters are also associated with changes in attitudes toward bride kidnapping.

Table 3.3 examines the impact of climate shocks on attitudes towards bride kidnapping. The table shows the effect of a dry winter on attitudes toward various events that might occur in a marriage that begins with a bride kidnapping. These events include

[9] retrieved from https://data.worldbank.org/indicator/SP.RUR.TOTL.ZS?end=2021&locations=KG&start=2000

"separation from someone who one really loved" (Lost Love), "domestic violence and abuse" (Domestic Violence), "burdensome demands on wife" (Burdensome Demands), and "being cut off from friends" (Cut From Friends). For those categories, a higher score indicates a stronger acceptance of bride kidnapping relative to love marriage. Consequently, a positive coefficient indicates that the perceived likelihood of the adverse event shifts in favor of bride-kidnapping marriage relative to love marriage; the gap measure alone does not identify which of the two underlying scores drives this change.

Table 3.3: The Impact of the Climate Shock on Attitudes Toward Bride Kidnapping

	Lost Love	Domestic Violence	Burdensome Demands	Cut from Friends
	(1)	(2)	(3)	(4)
Dry Winter	1.108***	0.571***	0.498***	0.514***
	(0.184)	(0.161)	(0.152)	(0.149)
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Indiv FE	Y	Y	Y	Y
Observations	11,622	11,622	11,622	11,622

*p<0.1; **p<0.05; ***p<0.01

Notes: Each column represents a dimension of attitudes toward bride kidnapping. Respondents rate how likely the following outcomes are to occur in marriages initiated by bride kidnapping and in love marriages: separation from someone they love, domestic violence and abuse, burdensome demands on the wife, and being cut off from friends. Responses range from 1 ('very possible') to 4 ('impossible'). The attitude outcome is the bride-kidnapping score minus the corresponding love-marriage score and ranges from -3 to 3; higher values therefore indicate a relatively more favorable assessment of marriages initiated by bride kidnapping. Dry Winter is an indicator equal to one if winter precipitation in the respondent's district in the survey year is more than 1.5 standard deviations below the district's historical winter mean. Year and individual fixed effects are included. Standard errors are clustered at the district level. Controls include age, age squared, marital status, and residence status. The average values of Lost Love, Domestic Violence, Burdensome Demands, and Cut from Friends are -0.88, -0.68, -0.63, and -0.50, respectively.

The results in Table 3.3 indicate that experiencing a dry winter is associated with higher gap scores for all four attitude outcomes. The estimated coefficients range from 0.498 to 1.108 points on the -3 to 3 gap scale. This suggests that dry winters not only

increase the likelihood of bride capture but also change attitudes towards the potential negative outcomes of marriages that begin with bride kidnapping. Specifically, the relative gap shifts in favor of bride-kidnapping marriage across these outcomes, although the gap measure does not identify which underlying marriage-type rating drives the change. Therefore, our findings suggest that dry-winter shocks affect both real incidents of bride kidnapping and attitudes towards this practice.

People in affected areas exhibit a relative shift in attitudes in favor of bride kidnapping following climate shocks. As previously mentioned, although the laws of Kyrgyzstan prohibit bride kidnapping, it continues to exist due to societal acceptance, resulting in a very low number of incidents officially reported to the Ministry of Internal Affairs. Since the persistence of this tradition is primarily attributed to the cultural acceptance of bride kidnapping, a positive shift in attitudes towards bride kidnapping after climate shock could be a mechanism for higher incidents of marriages through bride kidnapping. We explore further nuances of this mechanism in the next section.

3.5.3 Heterogeneity Analysis

In this section, our aim is to investigate the potential variations in attitudes towards bride kidnapping across different sub-samples of our dataset in response to dry winters. Since we do not have much information about women in the years when marriage happens, we cannot perform a heterogeneity analysis on the actual incidents of bride kidnapping. However, the richness of the LiK survey allows us to do this analysis for the attitudes part of the results. We aim to explore the factors that contribute to either exacerbating or improving the issue of a positive shift in attitudes toward bride kidnapping amid climate shocks.

First, we explore the impact of climate shocks on attitudes toward bride kidnapping for the groups of respondents from high-, middle-, and low-income households by adding respective interaction terms. Table 3.4 presents an examination of the interaction between income levels and climate shocks on attitudes towards bride kidnapping. This analysis uncovers the differential impact of a dry winter on various income brackets. Income is measured in the same survey wave as the attitude outcome and the dry-winter shock and may itself respond to the shock; the interactions therefore capture contemporaneous income heterogeneity rather than heterogeneity by predetermined income status. The

results from Table 3.4 add another layer to the narrative by suggesting that income might serve as a significant mechanism in mediating the relationship between climate shocks and attitudes towards bride kidnapping. When climatic conditions like a dry winter occur, they can intensify economic stress, particularly among lower-income households. These households often rely heavily on weather-dependent sectors such as agriculture and livestock, making them more susceptible to the adverse effects of climate shocks.

The coefficients of the interaction terms Dry X Low and Dry X Mid indicate the differential effects of dry winters on attitudes towards bride kidnapping for different income groups. A dry winter, in combination with low income, has a more pronounced impact on attitudes towards bride kidnapping across all categories relative to the high-income group (our reference group). This effect is statistically significant at the 5% level for the likelihood of lost love, burdensome demands, and being cut off from friends and at the 1% level for the likelihood of domestic violence.

Furthermore, middle-income households also show some changes in their attitudes, albeit to a lesser extent. It seems that while the economic strains caused by climate shocks might affect them, the magnitude is comparatively lower, potentially due to their more robust financial buffers. However, dry winter seems to have a statistically significant effect only on the likelihood of being cut off from friends, as indicated by the significant coefficient on Dry X Mid in column 4.

For the reference high-income group, dry-winter shock actually has a negative effect in all the categories of the dependent variable, meaning that for this group, attitudes shift from favoring bride kidnapping.

The disparity of these effects highlights that the implications of climate shocks like dry winters are not uniform across income groups, underscoring the importance of considering socio-economic stratifications in the evaluation of climate impacts on societal attitudes. Based on our findings, it can be inferred that the shift in attitudes towards bride kidnapping following dry-winter shocks is more pronounced among lower-income individuals. The stronger attitudinal response among lower-income households is consistent with the marriage-cost mechanism described above. Dry winters can intensify financial constraints precisely among households with fewer economic resources, making the costs of conventional marriage more difficult to finance. Bride kidnapping can reduce some of these costs, including kalym and related marriage expenditures, and may therefore become relatively more acceptable under economic stress. The concentration of the attitudinal response

Table 3.4: The Impact of the Climate Shock on Attitudes Toward Bride Kidnapping, Interaction Term Income

	Lost Love	Domestic Violence	Burdensome Demands	Cut from Friends
	(1)	(2)	(3)	(4)
Dry Winter	−0.002	−0.369**	−0.421**	−0.734***
	(0.212)	(0.185)	(0.175)	(0.176)
Low Inc.	−0.093	−0.023	0.016	−0.027
	(0.081)	(0.070)	(0.067)	(0.067)
Mid Inc.	0.093	0.069	0.066	0.056
	(0.059)	(0.051)	(0.049)	(0.049)
Dry X Low	0.677*	0.953***	0.800**	0.790**
	(0.398)	(0.348)	(0.329)	(0.331)
Dry X Mid	0.266	0.330	0.159	0.443**
	(0.267)	(0.233)	(0.221)	(0.221)
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Indiv FE	Y	Y	Y	Y
Observations	10,178	10,178	10,178	10,178

*p<0.1; **p<0.05; ***p<0.01

Notes: Each column represents a type of attitude toward bride kidnapping. Respondents answer how likely different events are to occur in marriage starting with the bride's kidnapping. The events are 'separation from someone who one really loved', 'domestic violence and abuse', 'burdensome demands on wife', and 'cut off from friends'. Respondents can range the answer from 1 to 4 where 1 is 'very possible' and 4 is 'impossible'. The attitude outcome is the bride-kidnapping score minus the corresponding love-marriage score and ranges from -3 to 3. Dry Winter is an indicator equal to one if winter precipitation in the respondent's district in the survey year is more than 1.5 standard deviations below the district's historical winter mean. Year and individual fixed effects are included. Standard errors are clustered at the district level. Controls include age, age squared, marital status, and residence status. Average values for 'Lost Love', 'Domestic Violence', 'Burdensome Demands', and 'Cut from Friends', respectively, are -0.88, -0.68, -0.63, -0.5. 'Low Inc.' is a dummy variable indicating the individual's income is below the 25th percentile; 'Mid. Inc.' is between the 25th and 75th percentiles.

among lower-income respondents is consistent with this financial-constraint channel.

These findings suggest that income level potentially moderates the effect of dry winters on attitudes toward bride kidnapping, particularly influencing how individuals perceive the likelihood of adverse outcomes in marriages initiated by such practices. These observations deepen our understanding of how climate shocks intersect with economic conditions to shape societal attitudes towards harmful practices like bride kidnapping.

Table 3.5 explores gender differences in the impact of the dry winter on attitudes toward bride kidnapping. Specifically, we examine the interaction between gender (female) and the dry winter. The analysis explores whether there is a difference in the attitudinal response to climate shocks between male and female respondents.

The interaction term 'Dry X Female' is not statistically significant for any of the outcomes, suggesting that the effect of a dry winter on attitudes toward bride kidnapping does not differ significantly between male and female respondents.

To summarize, our analysis suggests that the effect of the dry winter on attitudes towards bride kidnapping is not significantly different between genders. Both males and females tend to undergo similar shifts in their attitudes after experiencing the shock. We next examine heterogeneity by household composition.

Table 3.5: The Impact of the Climate Shock on Attitudes Toward Bride Kidnapping: Gender

	Lost Love	Domestic Violence	Burdensome Demands	Cut from Friends
	(1)	(2)	(3)	(4)
Dry Winter	1.153***	0.426*	0.525**	0.387*
	(0.255)	(0.223)	(0.210)	(0.206)
Dry X Female	−0.094	0.298	−0.058	0.263
	(0.363)	(0.318)	(0.299)	(0.293)
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Indiv FE	Y	Y	Y	Y
Observations	11,424	11,424	11,424	11,424

*p<0.1; **p<0.05; ***p<0.01

Notes: Each column represents a type of attitude toward bride kidnapping. Respondents answer how likely different events are to occur in a marriage, starting with the bride's kidnapping. The events are 'separation from someone who one really loved', 'domestic violence and abuse', 'burdensome demands on wife', and 'cut off from friends'. Respondents can range the answer from 1 to 4 where 1 is 'very possible' and 4 is 'impossible'. The attitude outcome is the bride-kidnapping score minus the corresponding love-marriage score and ranges from -3 to 3. Dry Winter is an indicator equal to one if winter precipitation in the respondent's district in the survey year is more than 1.5 standard deviations below the district's historical winter mean. Year and individual fixed effects are included. Standard errors are clustered at the district level. Controls include age, age squared, marital status, and residence status.

Table 3.6 investigates the interaction between the dry winter and the daughters' share in shaping attitudes toward bride kidnapping across the following categories: lost love, domestic violence, burdensome demands, and being cut off from friends. This analysis examines whether the effect of dry winters varies with the share of daughters among household members under age 25. A higher daughters' share may be associated with greater empathy toward abducted women and therefore with a smaller attitudinal shift following a dry winter.

The interaction term 'Dry X Share' is negative and significant across all four categories, suggesting that households with a larger share of daughters might experience less of a shift in attitudes following a dry winter. Because the daughters' share ranges from zero to one, a one-unit increase represents the extreme comparison between a household with no daughters and one in which all household members under age 25 are daughters. The corresponding interaction coefficients for 'lost love', 'domestic violence', 'burdensome demands', and 'being cut off from friends' are -0.814, -0.743, -0.328, and -0.683, respectively. This suggests that the presence of daughters in a household may indeed invoke more empathetic views toward abducted women, attenuating the impact of climate shocks on attitudinal shifts. These findings underscore the complex interplay of familial composition and climate shocks in shaping societal attitudes toward bride kidnapping, highlighting the potential role of familial empathy as a mitigating factor.

Lastly, Table 3.7 examines the interaction between the dry winter and low education in shaping attitudes toward bride kidnapping. The variable low education is a dummy variable indicating whether an individual has less than eight years of education (primary education). The analysis again focuses on the self-reported likelihood of different outcomes associated with bride kidnapping, including lost love, violence, burdensome demands, and being cut off from friends.

The interaction between Dry Winter and Low Education is not statistically significant at the 5% level in any of the four specifications. The coefficient for burdensome demands is negative and marginally significant at the 10% level, while the remaining interaction estimates are statistically insignificant and vary in sign. Overall, the results do not provide strong or systematic evidence that the attitudinal response to dry winters differs by education level. People with a lower educational level tend to give higher scores across all categories. This suggests that they often believe adverse events are less likely to occur in marriages resulting from kidnappings compared to love marriages. However, the

Table 3.6: The Impact of the Climate Shock on Attitudes Toward Bride Kidnapping, Interaction Term Daughters' Share

	Lost Love	Domestic Violence	Burdensome Demands	Cut from Friends
	(1)	(2)	(3)	(4)
Dry Winter	0.290*	0.188	−0.126	0.251*
	(0.171)	(0.151)	(0.142)	(0.143)
Daughters Share	0.118	0.014	−0.116	0.127
	(0.104)	(0.092)	(0.087)	(0.087)
Dry X Share	−0.814***	−0.743***	−0.328*	−0.683***
	(0.212)	(0.187)	(0.177)	(0.178)
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Indiv FE	Y	Y	Y	Y
Observations	8,727	8,727	8,727	8,727

*p<0.1; **p<0.05; ***p<0.01

Notes: Each column represents a type of attitude toward bride kidnapping. Respondents answer how likely different events are to occur in marriage starting with the bride's kidnapping. The events are 'separation from someone who one really loved', 'domestic violence and abuse', 'burdensome demands on wife', and 'cut off from friends'. Respondents can range the answer from 1 to 4 where 1 is 'very possible' and 4 is 'impossible'. The attitude outcome is the bride-kidnapping score minus the corresponding love-marriage score and ranges from -3 to 3. Daughters Share is a share of the number of females under 25 in the total number of HH members under 25. Dry Winter is an indicator equal to one if winter precipitation in the respondent's district in the survey year is more than 1.5 standard deviations below the district's historical winter mean. Year and individual fixed effects are included. Standard errors are clustered at the district level. Controls include age, age squared, marital status, and residence status. Average values for 'Lost Love', 'Domestic Violence', 'Burdensome Demands', and 'Cut from Friends', respectively, are -0.88, -0.68, -0.63, -0.5.

Table 3.7: The Impact of the Climate Shock on Attitudes Toward Bride Kidnapping: Education Interaction

	Lost Love	Violence	Demands	Cut Friends
	(1)	(2)	(3)	(4)
Dry Winter	1.360***	0.689***	0.857***	0.571***
	(0.233)	(0.207)	(0.195)	(0.191)
Low Educ.	0.265	0.272	0.221	0.001
	(0.338)	(0.299)	(0.281)	(0.276)
Wint. X Low Educ	−0.693	0.742	−0.975*	0.070
	(0.703)	(0.623)	(0.586)	(0.574)
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Indiv FE	Y	Y	Y	Y
Observations	9,500	9,500	9,500	9,500

*p<0.1; **p<0.05; ***p<0.01

Notes: This table analyzes the interaction terms of low education. The dependent variable is the attitude toward bride kidnapping measured by the self-reported likelihood of lost love, domestic violence, burdensome demands, or being cut from friends in marriage initiated by bride kidnapping. The attitude outcome is the bride-kidnapping score minus the corresponding love-marriage score and ranges from -3 to 3. Low Education is a dummy variable indicating that the individual has less than 8 years of education. Dry Winter is an indicator equal to one if winter precipitation in the respondent's district in the survey year is more than 1.5 standard deviations below the district's historical winter mean. Year and individual fixed effects are included. Standard errors are clustered at the district level. Controls include age, age squared, marital status, and residence status.

coefficients are also statistically insignificant.

In summary, our investigation into the heterogeneity of attitudes toward bride kidnapping reveals several insights. Examining attitudes across income groups, we uncover pronounced disparities. Individuals from low-income households demonstrate a notably positive shift in perceiving bride-kidnapping marriages as less prone to domestic violence and lost love due to dry winters. This change in attitude somewhat reduces in intensity among middle-income households but is still discernible. High-income households, conversely, show a negative shift. Such patterns hint at an underlying mechanism where, in the face of economic challenges exacerbated by climatic adversities, traditional practices like bride kidnapping might be seen as economically viable methods to alleviate financial burdens, such as avoiding "kalym" payments.

Another mechanism relates to our results on gender differences. Climate shocks could potentially impinge on the perception of male masculinity. As communities transition into a "survival mode" due to the stresses of climate shocks, men might feel their traditional roles as providers and protectors are threatened, leading them to assert their masculinity in other domains, including upholding traditional practices like bride kidnapping. However, our findings indicate that both males and females demonstrate similar shifts in attitudes, suggesting that the potential threat to masculinity might not be the dominant factor influencing attitudes in the context of this study.

Turning to household composition, the estimates indicate that a higher daughters' share is associated with a smaller attitudinal shift following a dry winter. Households with a higher proportion of daughters lean towards a less favorable view of bride kidnapping, especially in the dimensions of domestic violence, post a dry-winter shock. This could be indicative of families being more empathetic or protective, possibly fearing the potential repercussions for their own daughters in a society where bride kidnapping has become more normalized.

Finally, we explore the interaction between the dry-winter shock and low education levels. One might hypothesize that in economically constrained settings exacerbated by climate shocks, families might prioritize immediate survival needs over long-term educational investments, thereby inadvertently fostering environments where traditional beliefs flourish over more contemporary or progressive ones. In such contexts, even individuals with prior education might find their attitudes aligning more with age-old customs as they seek solace or coherence in familiar societal structures amidst the chaos of environ-

mental unpredictability. Our results do not support this channel since the coefficients for the interaction term are neither consistent in direction nor statistically significant. This suggests that education might not act as a robust buffer against traditional norms, at least in the context of this study. Even amidst broader education, deep-seated cultural norms and practices might persist, especially when communities face external pressures or challenges. Our analysis predominantly focuses on the immediate aftermath of short-term climatic shocks; hence we do not rule out the possibility that the education channel can have profound impacts on shaping attitudes over the long term.

3.6 Conclusion

Our study investigates the relationship between climate shocks, specifically insufficient precipitation in winter, and the incidence and attitudes towards bride kidnapping in Kyrgyzstan. This harmful marriage practice has significant negative consequences for the victims, including forced marriage, sexual and physical violence, unwanted pregnancies, social isolation, and mental health issues. Our findings show that insufficient winter precipitation increases the probability that an observed marriage is formed through bride kidnapping rather than through love or arranged marriage, and shifts attitudes toward bride kidnapping in a more favorable direction.

The paper's results highlight the complex interplay between climate change and bride kidnapping in Kyrgyzstan and emphasize that the consequences of climate change are not homogeneous. Our analysis underscores the complexity of the relationship between dry winters, socio-economic factors, and attitudes toward bride kidnapping. Income appears to be a crucial determinant, with lower-income individuals exhibiting a more susceptible attitude shift. Gender differences, while not significant, indicate a similar response to the shock among males and females. The presence of daughters in households suggests a potential influence in shaping negative attitudes.

Our results support the economic rationale behind more positive attitudes towards the bride-kidnapping practice in light of climate stress: lower-income households, perhaps driven by financial constraints, are more likely to reconcile with bride kidnapping as an acceptable norm, possibly viewing it as a strategy to navigate economic hardships.

On the other hand, our findings challenge the hypothesis that gender dynamics, par-

ticularly male dominance or attempts to assert masculinity, especially in the face of environmental stressors, play a primary role in perpetuating bride kidnapping. Both men and women seem to react similarly to climate shocks, suggesting that the practice is not merely an expression of male control or an outcome of heightened masculinity during tough times.

Our research has important implications for policymakers, as it suggests that climate change adaptation strategies need to consider not only the environmental and economic impacts of climate change but also its social and gender dimensions. This includes addressing harmful marriage practices like bride kidnapping, which can be exacerbated by climate shocks, and understanding the diverse ways in which climate change affects different groups within society.

Furthermore, our study contributes to a growing body of literature that seeks to understand the gendered dimensions of climate change impacts and highlights the need for more research on the links between climate change, gender inequality, and social norms in developing countries. By examining the implications of climate change for bride kidnapping in Kyrgyzstan, we have shed light on an under-researched aspect of the relationship between climate change and harmful marriage practices.

3.A Appendix

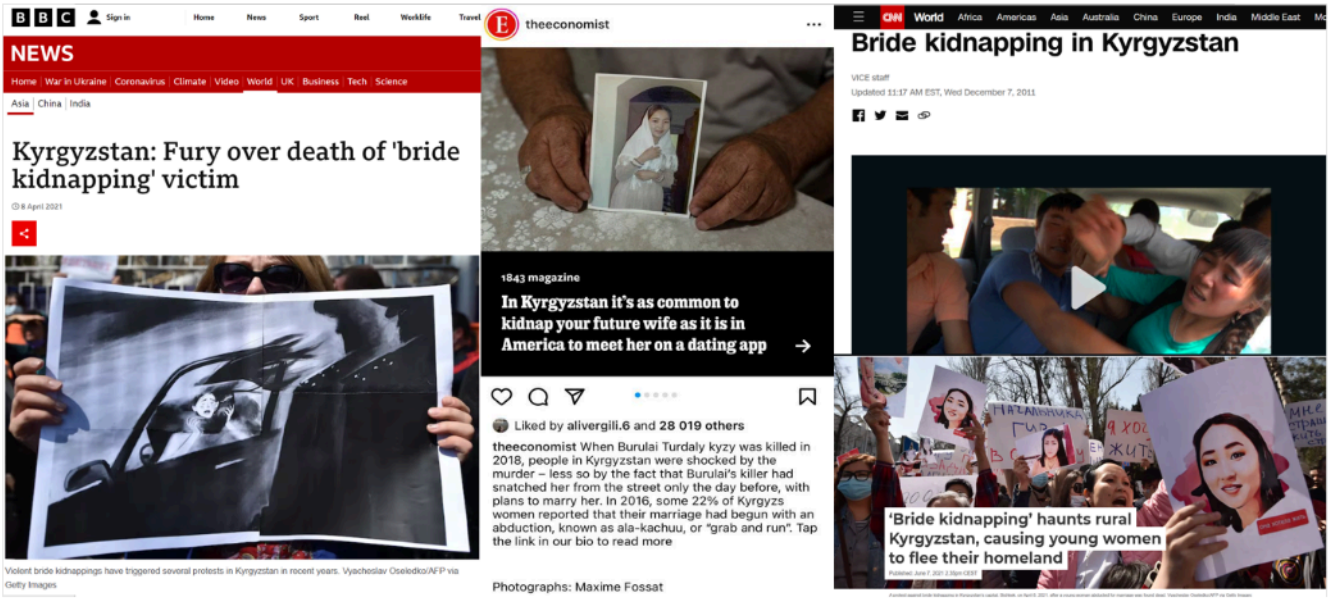


Figure 3.1: Media Coverage of the Protests against Bride Kidnapping in Kyrgyzstan

Table 3.A1: Summary Statistics by Marriage Practice

Variable	Levels	n	$\bar{x}$	Min	Max
Age	arranged marriage	571	34.1	18	58
	bride capture	281	35.3	18	54
	love marriage	1165	33.5	18	59
	all	2017	33.9	18	59
Age When Married	arranged marriage	571	20.7	16	35
	bride capture	281	20.2	15	29
	love marriage	1165	21.0	14	36
	all	2017	20.8	14	36
Husband Age When Married	arranged marriage	571	24.5	18	51
	bride capture	281	24.6	17	60
	love marriage	1165	24.2	16	99
	all	2017	24.4	16	99

Table 3.A2: Summary Statistics by Gender

Variable	**Levels**	**n**	$\bar{x}$	**Min**	**Max**
Age	Female	5636	41.58	17	95.00
	Male	5173	40.80	17	92.00
	all	10809	41.21	17	95.00
Empl. Decions: Myself	Female	5636	0.16	0	1.00
	Male	5173	0.17	0	1.00
	all	10809	0.17	0	1.00
Rural Residence	Female	5636	0.72	0	1.00
	Male	5173	0.75	0	1.00
	all	10809	0.73	0	1.00
Married	Female	5636	0.69	0	1.00
	Male	5173	0.73	0	1.00
	all	10809	0.70	0	1.00
BK Attitudes	Female	5636	14.03	6	24.00
	Male	5173	14.02	6	24.00
	all	10809	14.03	6	24.00
HH Income	Female	5636	129.55	0	2698.41
	Male	5173	137.47	0	2698.41
	all	10809	133.34	0	2698.41
Have Permanent Job	Female	5636	0.08	0	1.00
	Male	5173	0.15	0	1.00
	all	10809	0.11	0	1.00

Table 3.A3: Placebo Test: Future Dry Winters and Bride Kidnapping

	Bride Capture	
	(1)	(2)
	$t+1$	$t+2$
Future Dry Winter	0.075	0.070
	(0.080)	(0.058)
Controls	N	N
Year FE	Y	Y
District FE	Y	Y
Observations	1,472	1,472

Notes: Standard errors in parentheses. Bride Capture is a dummy variable equal to one if the observed marriage was formed through bride kidnapping and zero if it was a love or arranged marriage. Future Dry Winter is a dummy variable indicating that winter precipitation in the respondent's district was more than 1.5 standard deviations below the district's historical winter mean. Column (1) uses the dry-winter indicator one year after the marriage year, and Column (2) uses the indicator two years after the marriage year. Year and district fixed effects are included. Standard errors are clustered at the district level. Both columns use the same estimation sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Bibliography

Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2), 391–425.

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association*, 105(490), 493–505.

Adda, J., Dustmann, C., & Stevens, K. (2017). The career costs of children. *Journal of Political Economy*, 125(2), 293–337.

Ananat, E. O., Gruber, J., & Levine, P. B. (2007). Abortion legalization and life-cycle fertility. *Journal of Human Resources*, 42(2), 375–397.

Angelucci, M. (2015). Migration and financial constraints: Evidence from Mexico. *Review of Economics and Statistics*, 97(1), 224–228.

Angrist, J. D. & Evans, W. N. (1996). *Schooling and Labor Market Consequences of the 1970 State Abortion Reforms*. NBER Working Paper 5406, National Bureau of Economic Research.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12), 4088–4118.

Banerjee, A., Duflo, E., Ghatak, M., & Lafortune, J. (2013). Marry for what? caste and mate selection in modern India. *American Economic Journal: Microeconomics*, 5(2), 33–72.

Baser, O., Lu, Y., Sepulveda, F., Alemzadeh, A., & Endrizal, A. (2025). The effect of *Dobbs v. Jackson Women's Health Organization* on clinical diagnosis of postpartum depression. *Journal of Health Economics and Outcomes Research*, 12(1), 86–96.

Becker, A. (2019). *On the Economic Origins of Restrictions on Women's Sexuality*. Working Paper 7770, CESifo.

Becker, C. M., Mirkasimov, B., & Steiner, S. (2017). Forced marriage and birth outcomes. *Demography*, 54(4), 1401–1423.

Bengard, A. (2017). Krazhi nevest v Kyrgyzstane. kakim issledovaniyam verit'? [bride kidnapping in Kyrgyzstan. what research to believe in?].

Biermann, F. & Boas, I. (2010). Preparing for a warmer world: Towards a global governance system to protect climate refugees. *Global Environmental Politics*, 10(1), 60–88.

BIOM (2021). Droughts and floods: What is happening with Kyrgyzstan's climate and why it will be even worse.

Black, R., Arnell, N. W., Adger, W. N., Thomas, D., & Geddes, A. (2013). Migration, immobility and displacement outcomes following extreme events. *Environmental Science & Policy*, 27, S32–S43.

Bryan, G., Chowdhury, S., & Mobarak, A. M. (2014). Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh. *Econometrica*, 82(5), 1671–1748.

Carvajal, L. & Pereira, I. (2010). Evidence on the link between migration, climate shocks, and adaptive capacity. In R. Fuentes-Nieva & P. A. Seck (Eds.), *Risk, Shocks, and Human Development*. London: Palgrave Macmillan.

Cattaneo, C., Beine, M., Fröhlich, C. J., Kniveton, D., Martinez-Zarzoso, I., Mastrorillo, M., Millock, K., Piguet, E., & Schraven, B. (2019). Human migration in the era of climate change. *Review of Environmental Economics and Policy*, 13(2), 189–206.

Cattaneo, C. & Peri, G. (2016). The migration response to increasing temperatures. *Journal of Development Economics*, 122, 127–146.

Cheng, P., Lin, Z., & Liu, Y. (2011). Do women pay more for mortgages? *Journal of Real Estate Finance and Economics*, 43(4), 423–440.

Clarke, D. & Mühlrad, H. (2021). Abortion laws and women's health. *Journal of Health Economics*, 76, 102413.

Corno, L., Hildebrandt, N., & Voena, A. (2020). Age of marriage, weather shocks, and the direction of marriage payments. *Econometrica*, 88(3), 879–915.

CSCE (2017). Bride kidnapping in the Kyrgyz Republic.

Dench, D., Pineda-Torres, M., & Myers, C. (2024). The effects of post-Dobbs abortion bans on fertility. *Journal of Public Economics*, 234, 105124.

Dillon, A., Mueller, V., & Salau, S. (2011). Migratory responses to agricultural risk in Northern Nigeria. *American Journal of Agricultural Economics*, 93(4), 1048–1061.

Durrance, C. P., Wang, Y., & Wolfe, B. (2024). Do mandatory waiting periods for abortion increase intimate partner violence? *Journal of Health Economics*, 98, 102939.

Fang, L. & Munneke, H. J. (2020). Gender equality in mortgage lending. *Real Estate Economics*, 48(4), 957–1003.

Farin, S. M., Hochn-Velasco, L., & Pesko, M. F. (2024). The impact of legal abortion on maternal mortality. *American Economic Journal: Economic Policy*, 16(3), 174–216.

Gardner, G. (2024). The maternal and infant health consequences of restricted access to abortion in the United States. *Journal of Health Economics*, 98, 102938.

Germanwatch (2022). Global climate risk index 2022.

Goldin, C. & Katz, L. F. (2002). The power of the pill: Oral contraceptives and women's career and marriage decisions. *Journal of Political Economy*, 110(4), 730–770.

Gröschl, J. & Steinwachs, T. (2017). Do natural hazards cause international migration? *CESifo Economic Studies*, 63, 445–480.

Guttmacher Institute (2026). Interactive map: US abortion policies and access after Roe. Retrieved May 6, 2026.

Halliday, T. (2006). Migration, risk, and liquidity constraints in El Salvador. *Economic Development and Cultural Change*, 54(4), 893–925.

Howe, A. (2022). Court has voted to overturn Roe, according to draft opinion published by Politico. `https://www.scotusblog.com/2022/05/court-has-voted-to-overturn-roe-according-to-draft-opinion-published-by-politico/`.

Hsiang, S. & Kopp, R. (2018). An economist's guide to climate change science. *Journal of Economic Perspectives*, 32(4), 3–32.

IPCC (2013). *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.

IPCC (2014). *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.

IPCC (2022). *Climate Change 2022: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, UK and New York, NY, USA: Cambridge University Press. 3056 pp.

Jones, B. & O'Neill, B. (2016). Spatially explicit global population scenarios consistent with the Shared Socioeconomic Pathways. *Environmental Research Letters*, 11(8), 084003.

Jones, K. M. & Pineda-Torres, M. (2024). TRAP'd teens: Impacts of abortion provider regulations on fertility & education. *Journal of Public Economics*, 234, 105112.

Kleinbach, R., Ablezova, M., & Aitieva, M. (2005). Kidnapping for marriage (Ala Kachuu) in a Kyrgyz village. *Central Asian Survey*, 24(2), 191–202.

Kleven, H., Landais, C., & Søgaard, J. E. (2019). Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4), 181–209.

Ladd, H. F. (1998). Evidence on discrimination in mortgage lending. *Journal of Economic Perspectives*, 12(2), 41–62.

Liu, Z. & Liang, H. (2025). Are credit scores gender-neutral? evidence of mis-calibration from alternative and traditional borrowing data. *Journal of Behavioral and Experimental Finance*, 47, 101081.

Londoño-Vélez, J. & Saravia, E. (2025). The impact of being denied a wanted abortion on women and their children. *Quarterly Journal of Economics*, 140(2), 1061–1110.

Lundborg, P., Plug, E., & Rasmussen, A. W. (2017). Can women have children and a career? IV evidence from IVF treatments. *American Economic Review*, 107(6), 1611–1637.

Miller, S., Wherry, L. R., & Foster, D. G. (2023). The economic consequences of being denied an abortion. *American Economic Journal: Economic Policy*, 15(1), 394–437.

Ministry of Energy and Industry of the Kyrgyz Republic (2016). National communication of the Kyrgyz Republic to the United Nations Framework Convention on Climate Change.

Ministry of Foreign Affairs of the Netherlands (2018). Climate change profile: Indonesia.

Mueller, V., Gray, C., & Kosec, K. (2014). Heat stress increases long-term human migration in rural Pakistan. *Nature Climate Change*, 4(3), 182–185.

Munnell, A. H., Tootell, G. M. B., Browne, L. E., & McEneaney, J. (1996). Mortgage lending in Boston: Interpreting HMDA data. *American Economic Review*, 86(1), 25–53.

Myers, C. K. (2017). The power of abortion policy: Reexamining the effects of young women's access to reproductive control. *Journal of Political Economy*, 125(6), 2178–2224.

Myers, N. (2002). Environmental refugees: A growing phenomenon of the 21st century. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 357(1420), 609–613.

Naumova, O. (2016). Povedenie devushek v tradicii umykaniya (po materialam Kyrgyzstana) [the behavior of girls in the tradition of abduction (based on Kyrgyzstan)]. *Vestnik Antropol.*, 3(35), 74–91.

Nawrotzki, R. J. & Bakhtsiyarava, M. (2017). International climate migration: Evidence for the climate inhibitor mechanism and the agricultural pathway. *Population, Space and Place*, 23(4).

Oswald Spring, U. (2019). Gender, climate change and sustainable development goals. In O. Spring (Ed.), *Úrsula Oswald Spring: Pioneer on gender, peace, development, environment, food and water: With a foreword by Birgit Dechmann* (pp. 225–241). Springer International Publishing.

Oyez (2022). Dobbs v. Jackson Women's Health Organization. https://www.oyez.org/cases/2021/19-1392.

Poggi, C. (2019). Credit availability and internal migration: Evidence from Thailand. *Journal of Development Studies*, 55(5), 861–875.

Pope, D. H., McMullen, H., Baschieri, A., Philipose, A., Udeh, C., Diallo, J., & McCoy, D. (2022). What is the current evidence for the relationship between the climate and environmental crises and child marriage? a scoping review. *Global Public Health*.

Rao, N., Singh, C., Solomon, D., Camfield, L., Sidiki, R., Angula, M., Poonacha, P., Sidibe, A., & Lawson, E. T. (2023). Managing risk, changing aspirations and household dynamics: Implications for wellbeing and adaptation in semi-arid Africa and India. *World Development*, 137, 105170.

Reggers, A. (2019). Climate change is not gender neutral: Gender inequality, rights and vulnerabilities in Bangladesh: Policy strategies for adaptation and resilience. In *Gender Inequality, Climate Change and Disaster Risk Reduction* (pp. 107–124).

Renaud, F., Bogardi, J. J., Dun, O., & Warner, K. (2007). *Control, Adapt or Flee: How to Face Environmental Migration?* Unu-ehs intersections, UNU-EHS.

Saldana-Zorrilla, S. O. & Sandberg, K. (2009). Spatial econometric model of natural disaster impacts on human migration in vulnerable regions of Mexico. *Disasters*, 33(4), 591–607.

SMERU Research Institute (2013). The use of the Social Protection Card (KPS) and the implementation of the 2013 unconditional cash transfer program (BLSM). https://smeru.or.id/en/publication/use-social-protection-card-kps-and-implementation-2013-unconditional-cash-transfer.

Stecklov, G., Winters, P., Stampini, M., & Davis, B. (2005). Do conditional cash transfers influence migration? a study using experimental data from the Mexican Progresa programme. *Demography*, 42, 769–790.

Stern, N. (2007). *The Economics of Climate Change: The Stern Review*. Cambridge University Press.

Strauss, J., Witoelar, F., & Bondan, S. (2016). *The Fifth Wave of the Indonesia Family Life Survey: Overview and Field Report*. Technical report, RAND.

Supreme Court of the United States (2022). Docket for no. 19-1392, Thomas E. Dobbs, state health officer of the Mississippi Department of Health, et al. v. Jackson Women's Health Organization, et al. https://www.supremecourt.gov/search.aspx?filename=/docket/docketfiles/html/public/19-1392.html.

Terry, G. (2009). No climate justice without gender justice: an overview of the issues. *Gender and Development*, 17(1), 5–18.

Texas Real Estate Research Center (2022a). Texas housing insight: December 2022. Texas A&M University.

Texas Real Estate Research Center (2022b). Texas housing insight: September 2022. Texas A&M University.

Thiede, B. & Gray, C. (2017). Heterogeneous climate effects on human migration in Indonesia. *Population and Environment*, 39.

Trinh, R. A. & Zhang, Q. (2021). Adverse shocks, household expenditure and child marriage: evidence from India and Vietnam. *Empirical Economics*, 61(3), 1617–1639.

Tsaneva, M. (2020). The effect of weather variability on child marriage in Bangladesh. *Journal of International Development*, 32(8), 1346–1359.

UN Women (2018). Global study on legal measures to address violence against women: Regional review of laws on violence against women in the Arab States, Eastern and Southern Africa, South Asia and Southeast Asia.

United Nations Population Fund (UNFPA) (2018). Against my will: defying the practices that harm women and girls and undermine equality.

U.S. Supreme Court (2022). *Dobbs v. Jackson Women's Health Organization*. 597 U.S. 215.

World Bank (2010). Building a disaster-resilient future. Indonesia Rising: Policy Priorities for 2010 and Beyond. World Bank Office Jakarta. Retrieved from `https://documents1.worldbank.org/curated/en/674521468259730398/pdf/534720BRI041di100Box345611B1PUBLIC1.pdf`.

World Bank (2012a). Bantuan Langsung Tunai (BLT): Temporary unconditional cash transfer. `https://openknowledge.worldbank.org/entities/publication/8398e3e7-7b6f-5eaa-acb8-ffe69ece10a9`.

World Bank (2012b). PKH conditional cash transfer. `https://openknowledge.worldbank.org/entities/publication/bba7bd78-a716-573a-b72f-0e73699aac22`.

World Bank (2012c). *PKH Conditional Cash Transfer Social Assistance Program and Public Expenditure Review 6*. Technical report, World Bank, Jakarta, Indonesia.

World Bank (2017). Indonesia social assistance public expenditure review update. `https://documents.worldbank.org/curated/en/535721509957076661/pdf/120905-REVISED-PUBLIC-Screen-English-1211-update.pdf`.

World Bank (2021). Kyrgyz Republic agriculture sector review.

Zickgraf, C. & Perrin, N. (2016). Immobile and trapped populations. In F. Gemenne, D. Ionesco, & D. Mokhnacheva (Eds.), *Atlas Der Umweltmigration* (pp. 44–46). München: oekom.