Title
Essays in labor economics

Author
Chou, Tiffany

Publication Date
2011

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Essays in Labor Economics

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Economics

by

Tiffany Chou

Committee in charge:

Professor Eli Berman, Chair
Professor Julie Cullen
Professor Gordon Dahl
Professor Jesse Driscoll
Professor John Evans

2011

The Dissertation of Tiffany Chou is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2011

DEDICATION

To Mom and Dad.

TABLE OF CONTENTS

Signature Page ... iii

Dedication ... iv

Table of Contents ... v

List of Figures ... viii

List of Tables ... ix

Acknowledgements ... xi

Vita ... xiii

Abstract of the Dissertation ... xiv

Chapter 1: Is Fertility Decline Contagious? Trans-Atlantic norm shifts and fertility in the post-Vatican II period ... 1

- 1.1 Introduction ... 2
- 1.2 Religion in Europe ... 6
- 1.3 Empirical Estimation ... 9
- 1.4 Data ... 17
- 1.5 Results ... 20
- 1.6 Conclusion ... 31
- 1.7 Acknowledgement ... 33
- 1.8 Appendix ... 52
- 1.9 References ... 59

Chapter 2: Does Factor-Biased Technological Change Stifle International Convergence? Evidence from Manufacturing ... 64

2.1 Introduction 65

2.2 TFP Growth and Factor Accumulation in Manufacturing 68

2.3 Factor-Biased Technological Change in Production: A Framework for Estimation 72

2.4 Data 81

2.5 Results for 1980s 92

2.6 Results for 1990s 105

2.7 Implications 113

2.8 Conclusion 115

2.9 Acknowledgement 119

2.10 Appendix 146

2.11 References 149

Chapter 3: Does Development Assistance Reduce Violence? Evidence from Afghanistan 154

3.1 Introduction 155

3.2 Insurgency and Development Aid as COIN 158

3.3 Theoretical Background 160

3.4 Data 169

3.5 Results 177

3.6 Conclusion 191

3.7 Acknowledgement 194

3.8 Appendix 209

3.9 References 214

LIST OF FIGURES

Figure 1.1: Second-Generation vs. Europeans, First Differences 49

Figure 1.2: European Church Attendance 50

Figure 1.3: First Differences, All Countries 51

Figure 2.1: Non-convergence of Growth Rates 135

Figure 2.2: Non-Convergence of Manufacturing Value Added (per worker) 136

Figure 2.3: Decomposition of Increases in Manufacturing Value Added (per worker), 1980s 137

Figure 2.4: Decomposition of Increases in Manufacturing Value Added (per worker), 1990s 138

Figure 2.5: Technological Change with Relative Skill Bias 139

Figure 2.6: Leverage Plots, GIS 140

Figure 2.7: Conditional Convergence or Factor Bias? 141

Figure 2.8: Leverage Plots for Output Specification, INDSTAT 142

Figure 2.9: Total Manufacturing Employment, By Country 143

Figure 2.10: Total Manufacturing Value Added, By Country 144

Figure 2.11: Italian Employment, By Industry 145

Figure 3.1: Violence in Afghanistan 207

Figure 3.2: Nonlinear Relationship between Insurgent Violence and Afghan Stability Perceptions 208

LIST OF TABLES

Table 1.1: Summary Statistics, Condensed 34

Table 1.2: Trans-Atlantic Fertility Correlation 35

Table 1.3: Fertility, Religion, and Church Attendance 37

Table 1.4: Including Individual Marital Status and Labor Force Participation........ 38

Table 1.5: Additional Countries 39

Table 1.6: Intergenerational Fertility Patterns........ 40

Table 1.7: Untruncated TFR........ 41

Table 1.8: Pre-Post 1965 Immigrant Cohort Characteristics........ 42

Table 1.9: Summary Statistics, Second-Generation Women age 18-35, 1970 43

Table 1.10: Summary Statistics, Second-Generation Women age 18-35, 2000-2006 . 44

Table 1.11: Summary Statistics, Country-Level Covariates 45

Table 1.12: Including Country-Level Work and Marriage Norms........ 46

Table 1.13: Labor Force Participation and Marital Status as Outcomes........ 47

Table 1.14: Married Women 48

Table 2.1: Factor Shares in Value Added........ 120

Table 2.2: Summary of Data Sources........ 121

Table 2.3: Sample Descriptive Statistics - 1980........ 122

Table 2.4: Descriptive Statistics, GIS........ 123

Table 2.5: Descriptive Statistics, INDSTAT3 124

Table 2.6: Descriptive Statistics, INDSTAT4........ 125

Table 2.7: Factor Bias Estimates from Production Function 126

Table 2.8: Factor Bias Estimates from Production Function - Specification Checks 127

Table 2.9: Factor Bias Estimates from TFP Specification 128

Table 2.10: Factor Bias Estimates – High vs. Middle Income Countries (assuming URS) 129

Table 2.11: Total Employment vs. Skilled-Unskilled 130

Table 2.12: Factor Bias Estimates, INDSTAT3 132

Table 2.13: Factor Bias Estimates, INDSTAT4 133

Table 2.14: Industry Coverage 134

Table 3.1: Summary Statistics 195

Table 3.2: Temporal Patterns of Violence 196

Table 3.3: Development Spending and Violence 197

Table 3.4: Small vs. Large LGCD projects 198

Table 3.5: Small vs. Large CERP projects, Oct 2009 – Apr 2010 199

Table 3.6: Development Spending and Violence, by Stability Index 200

Table 3.7: NSP Spending by Project Sector 201

Table 3.8: LGCD Spending by Project Sector 202

Table 3.9: CERP counts by Project Category 203

Table 3.10: NSP Spending by Project Activity 204

Table 3.11: NSP and LGCD spending, full results 205

Table 3.12: Small vs. Large CERP projects, Oct 2009 – Apr 2010 206

ACKNOWLEDGEMENTS

I would like to thank Professors Eli Berman, Julie Cullen, and Gordon Dahl for excellent feedback, continued encouragement, and abundant patience.

Chapter 1 also benefited from discussions with John H. Evans, Daniel Egel, L. Choon Wang, Sarada, Youjin Hahn, Oana Hirakawa, Philip Neary, and Hee-Seung Yang. Valuable comments were also provided by seminar participants at UC San Diego, Monash University, the Office of Economic Policy at the Treasury Department, and attendees of the 2010 ASREC meeting and the 2010 All-California Labor Conference.

Chapter 2 is currently being prepared for submission for publication of the material. Berman, Eli; Tiffany Chou. The dissertation author was a primary investigator and author of this material. We acknowledge the helpful comments of Robert Barro, Francesco Caselli, Simon Gilchrist, Samuel Kortum, Kevin Lang, Dilip Mookherjee and Andrew Weiss; seminar participants at Boston University, Case Western Reserve University, University of Pittsburgh, Rice University, UC San Diego, NBER Labor Studies and Productivity sessions, NEUDC and SOLE conferences. This work continues research with Stephen Machin during which a version of these data were prepared.

Chapter 3 is currently being prepared for submission for publication of the material. Berman, Eli; Tiffany Chou. The dissertation author was a primary investigator and author of this material. Thanks to Michael J. Callen, Luke Nayef Condra, Michael J. Callen, Joseph Felter and Jacob Shapiro for excellent discussion and comments. We acknowledge grant #2007-ST-061-000001 by the United States Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events and by U.S. Department of Defense's Minerva Initiative (AFOSR Grant FA9550-09-1-0314). Any opinions, findings, conclusions or recommendations in this document are those of the authors and do not reflect views of the United States Department of Homeland Security or United States Department of Defense.

VITA

2004	Bachelor of Arts, University of California, Berkeley
2006	Master of Arts, University of California, San Diego
2011	Doctor of Philosophy, University of California, San Diego

ABSTRACT OF THE DISSERTATION

Essays in Labor Economics

by

Tiffany Chou

Doctor of Philosophy

University of California, San Diego, 2011

Professor Eli Berman, Chair

This dissertation consists of three unrelated papers in labor economics. The first chapter documents the role of norms, both cultural and religious, in the fertility decisions of second-generation women in the US. Using two cohorts of immigrants (1970 and 2000s), I find that fertility declines among second-generation immigrants in the US are highly correlated with contemporaneous falls in total fertility rates (TFR) in

Europe, implying that changes in the origin countries after parental emigration are still mirrored among current immigrants. This cross-country correlation is stronger for women from predominantly Catholic countries, which is consistent with immigrants from Catholic Europe sharing the Church's pro-natalist theology.

The second chapter estimates the extent to which factor bias within manufacturing affects productivity growth across countries in the last two decades of the 20th Century. Skill-biased technological change (SBTC) implies that countries with more skilled labor and capital experience higher growth in total factor productivity (TFP), which is the case in both developed and developing countries in the 1980s. Labor-biased technological change is especially strong among the "newly industrializing countries" in the 1990s. These results are consistent with the empirical literature on skill-biased technological change, and may explain why "conditional convergence" of per capita income across countries is so slow.

The final chapter examines the violence-reducing effect of development spending in Afghanistan. Using data from three distinct reconstruction programs and military records of insurgent-initiated events, the analysis finds that overall spending has no clear effects on the frequency of rebel attacks. Moreover, the types of development program most effective at reducing violence in Iraq –small CERP projects—does not appear to do so in Afghanistan.

Chapter 1: Is Fertility Decline Contagious? Trans-Atlantic norm shifts and fertility in the post-Vatican II period

Abstract

This paper finds that fertility declines among second-generation immigrants in the US are highly correlated with contemporaneous drops in total fertility rates (TFR) in European countries of origins between 1970 and 2000. Since second-generation immigrants are all born in the US and share US-specific prices, technology, and institutions, this trans-Atlantic correlation implies that changes in norms are being shared between immigrant women and their European "cousins". The current literature has shown that *past* outcomes in the originating countries predict current outcomes among immigrants; this paper demonstrates that even changes in the country of origin after emigration are still mirrored by immigrants in the US. The strength of this correlation appears to be linked to religion; immigrant cohorts from strongly Catholic nations exhibit larger declines in the number of children, which is consistent with immigrants and "cousins" sharing the Catholic Church's pro-natalist theology or social services. In contrast, European TFR has little power in predicting the fertility of women from majority Protestant nations once controls for marital status and labor force participation are included.

1.1 Introduction

Fertility is one of the most significant household decisions, altering not just the consumption patterns and behavior of individual families, but also the broader social, economic, and political landscape. Since Becker and Schultz's seminal insight, economists have been concerned with how prices and institutions affect both the monetary and time cost of raising children (Becker, 1981; Schultz, 1981). In particular, the economic and institutional environments in Southern Europe are thought to have caused sharp declines in fertility in areas that were previously characterized by large families (Kohler, Billari, & Ortega, 2006). Employment frictions in European labor markets, coupled with weak social insurance programs and limited parental leave opportunities, create an environment where the cost of having children is quite high.[1,2] While these studies explore how differences in incentives affect fertility, this paper fits into the recent literature that examines and attempts to quantify the role of preferences, norms, and culture in determining economic outcomes (Bisin & Verdier, 2010).

Using Census data on two cohorts of Western European immigrants in the US, I find that the fertility decisions of second-generation women are highly correlated with total fertility rates (TFR) in the original sending country: women whose fathers

---

[1] Adsera (2004a) and del Boca (2002) look at whether higher costs of transitioning into and out of the labor force explain lower fertility among Southern European women. Youth unemployment is also relatively high in Mediterranean Europe, making it more difficult for young men and women to start families (Bettio & Villa, 1998).

[2] The generosity of social assistance programs is potentially endogenous; places where the family is expected to provide, for example, unemployment aid or childcare, may be unlikely to have the government provide these services. This is exactly the situation in Southern Europe where strong family ties coincide with conservative social policy (Bertola, Jimeno, Marimon, & Pissarides, 2001; Alesina, Glaeser, & Sacerdote, 2001).

are from high-fertility countries tend to have larger families.[3] Like other empirical studies of culture, the identification of norms rests on the argument that since second-generation women all live in the US, they are likely to share similar economic and institutional environments with other Americans while remaining relatively unaffected by policy shifts or economic shocks in Europe.[4] However, they are still likely to share some unobserved norms with the European "cousins" across the Atlantic, which appears as the residual correlation in outcomes between immigrants in the US and European natives. The results here suggest that a one-child decrease in European fertility predicts 0.3 fewer children per woman, even after controlling for an individual's observable characteristics.[5] The magnitude of this correlation is quite large as this is comparable to the fertility differential between a college-educated woman and one with only a high school diploma.

This trans-Atlantic fertility correlation is concisely illustrated in Figure 1.1, which plots the changes in immigrant fertility against declines in source-country fertility. The x-axis is simply the change in TFR between 1970 and 2000 in each of the European countries. The y-axis is best conceptualized as the result of a two-step process. First, for each of the immigrant cohorts, estimate the number of children born

[3] The immigrant assimilation literature (e.g., LaLonde & Topel, 1997; Blau, 1992) is primarily interested in how the outcomes of US immigrants compare to those of "native" Americans. Rather than asking whether second-generation immigrants "catch-up" to other Americans, this paper looks at whether those persistent fertility differences are predicted by outcomes in the originating country.

[4] This assumption is key to identification, but I am unaware of any work that evaluates its veracity in this particular context. The current literature on intergenerational assimilation (e.g. Card, DiNardo, & Estes, 2000; Blau et al., 2008) seems to find that even second-generation immigrants have remarkably similar outcomes to other US natives.

[5] For comparison, the TFR in the US in 1970 is 2.46 children per woman. Compared to this base fertility, an additional 0.3 children associated with coming from a country with large fertility declines appears quite large.

to women from each of the individual sending countries. Then take the 1970-2000 difference between estimated second-generation fertility for each country. The vertical axis plots this difference in “cohort" fertility against the change in TFR in the originating country; the correlation being estimated between immigrant and “native” fertility is the slope of the line through Figure 1.1.[6]

As Europe-specific economic institutions like rigid labor markets or legislative controls on contraception are unlikely to be driving this trans-Atlantic fertility correlation between Europeans and Americans, the pattern of countries in Figure 1.1 suggests shared religion as a plausible alternative. The three countries in the lower-left corner (Ireland, Portugal, and Spain) that experienced large declines in fertility both in Europe and in the US are all predominantly Roman Catholic. The Catholic Church is particularly salient because of its international reach and its consistent, pro-fertility teachings. Using the percentage of Catholics in the source country as a proxy for immigrant religion,[7] I find that the trans-Atlantic correlation is significantly higher for women from highly-Catholic countries, consistent with the idea that shared “Catholic”-ness as a mechanism through which shifts in norms could transmit from Europe to the US.

In addition, I find that fertility declines among immigrants are also weakly correlated with declining religiosity in the wake of the Second Vatican Council (Vatican II). By the time the conference closed in 1965, Vatican II had issued multiple

[6] While this is exposited as a two-step process, the empirical strategy estimates this slope in one step using pooled Census cross-sections and an appropriately specified estimating equation.
[7] The US Census does not collect data on religious affiliation. There are alternate surveys that do ask individuals’ religion and religiosity, they lack sufficient information to specifically identify second-generation immigrants.

progressive decrees, surprising both Church insiders and laypeople who had expected the conference to reaffirm Church traditions in a time of great social and political change (Wilde, 2004). Berman, Iannaccone, and Ragusa (2006) argue that Vatican II and the consequent decline in Church social services played a key role in the recent fertility declines in Southern Europe. As nuns are the primary providers of Catholic social, rather than theological, services for the Church, a sharp decline in nuns per capita post-Vatican II substantially lowers the level and generosity of (child-friendly) Church social services, resulting in lower fertility across multiple Catholic countries. Here, I find that the post-Vatican II declines in European church attendance also predict declining fertility among European-Americans. This suggests that declining adherence to Church fertility norms, as opposed to Church-provided social services, is an underlying cause of the religion-specific fertility similarities between Europeans and Americans documented here.

This paper makes two major contributions and in doing so, is an early attempt to tie together the disparate economics literatures on culture and religion. First, it adds to the small but growing body of work on the empirical role of preferences, as opposed to prices, in economic decisions.[8] In particular, past authors have examined the temporal and spatial persistence of culture, and this paper documents empirical evidence of what appears to be contagion in preferences. Even though large changes in TFR in the sending countries are occurring after the parents have already emigrated,

---

[8] Unlike other social scientists, economists are traditionally skeptical of attributing changes in behavior to changes in norms (Stigler & Becker, 1977). However, sociologists have debated the effect of religion, and Catholicism in particular, on behavior since the early 20th Century (Weber, 1946; Durkheim, 1965). In that sense, one contribution of this paper is that it brings conventional economist perspective and methodology to a traditionally non-economist topic.

these outcomes still significantly predict American fertility. Second, it argues that religion, or more specifically Catholicism, is a potential source of norms and transmitter of these cultural shifts. Religious affiliation is already known to be significantly correlated with a host of outcomes ranging from education (Lehrer, 2009) to happiness (Clark & Lelkes, 2009), but there is little emphasis on how religion might shape preferences, in addition to shifting the "price" of specific behaviors.[9]

The remainder of this paper is organized as follows. The next section provides a brief discussion of religion in Western Europe, with focus on the contextually relevant differences between the Catholic Church and Protestant denominations. Section 1.3 motivates and discusses the estimating equations while Section 1.4 briefly details the data used. Section 1.5 presents and discusses the main results and some robustness tests, and Section 1.6 concludes.

1.2 Religion in Europe

Christianity has historically dominated throughout Western Europe, but since the Protestant Reformation, individual countries have tended to be either highly Catholic (Ireland, Spain, Italy, Portugal) or highly Protestant (UK, Scandinavia).[10] Though the Catholic Church and individual Protestant denominations vary along

[9] Almond, Edlund, and Milligan (2009) find that second-generation Asian-Canadians who self-report as Christian or Muslim still demonstrate a cultural preference for sons but are unlikely to exercise this preference using sex-selective abortion. However, they do not quantitatively link immigrant differentials to skewed sex ratios in the originating countries.

[10] The majority of Greeks and Eastern Europeans are Orthodox and hence neither Protestant nor Catholic. France is technically a highly-Catholic country but has had a rocky relationship with the Vatican, culminating in 1905's Separation Law, which mandated very strict separation between church and state (Warner, 2000).

multiple dimensions, the most relevant to this paper are the differences in teachings related to fertility and family life. The Catholic Church teaches that contraception is intrinsically wrong since it directly conflicts with Nature,[11] and theologically emphasizes the importance of marriage and the traditional two-parent family. In addition, the Church provides a number of social services like hospitals, schools, and charities.[12] The hierarchical structure of the Catholic Church also implies that Catholic teachings on family and contraception should be quite consistent regardless of geography.

The Protestant denominations, in contrast, tend to be more varied in their views on birth control and gender norms, emphasizing the choice and appropriateness of such decisions to the individual over adherence to the official theological stance.[13] In addition, European Protestants tend to belong to "mainline" denominations (Anglican, Lutheran), not the more conservative evangelical churches in North America. Together, these characteristics suggest that American and European Catholics should share similar pro-fertility norms, ones that tend to be quite different from those of other Christians or non-Catholics. The Catholic Church's fertility-related teachings should affect how individuals perceive the costs and benefits of raising children. Hence, practicing Catholics should exhibit higher fertility compared

[11] As stated by a 1968 Papal encyclical, "…the direct interruption of the generative process already begun and, above all, all direct abortion, even for therapeutic reasons, are to be absolutely excluded as lawful means of regulating the number of children" (*Humanae Vitae* 14).
[12] In the 1950s, 11% of all American students were enrolled in Catholic schools, and the Church operated one-fifth of all hospital beds (Fialka, 2003).
[13] For example, the Church of England states that "the responsibility for deciding upon the number and frequency of children was laid by God upon the consciences of parents 'in such ways as are acceptable to husband and wife'" (www.cofe.anglican.org).

to otherwise similar households, which does appear to be the case (Westoff & Jones, 1979).[14]

Second Vatican Council (Vatican II)

Called by Pope John XXIII to "renew" the Church in a time of sociopolitical and technological change, the Second Vatican Council resulted in many changes to official Church doctrine and practices (Alberigo, 2006). Whereas the previous Vatican council (1869-1870) had reemphasized conservative Church teachings, Vatican II's decrees tended to be liberal updates to those longstanding policies. Whereas Mass used to be in Latin, Vatican II allowed religious services to be conducted in the local vernacular and incorporate local customs. The Church also relaxed dietary restrictions and relinquished its claim as the one true church (Wilde, 2004; Hout & Greeley, 1987).[15] Whether or not it was successful in its stated aim of updating the Church, the religiosity of Catholics sharply declined post-Vatican II. Fewer people were entering or remaining in the priesthood or other religious orders. While the 1970s were a time of widespread secularization, declines in church attendance were much sharper for Catholic countries than for Protestant ones (see Figure 1.2).

Declining religiosity after Vatican II increased individual willingness to openly question Catholic doctrine, making it more difficult for the Catholic Church to elicit

[14] As discussed by Stolzenberg, Blair-Loy, and Waite (1995), this could be a case of reverse causality where families that choose to have many children also choose to remain attached to the Church. This actually appears to be the case in Spain (Adsera, 2004b).
[15] Amidst all these changes, however, the Church re-emphasized its bans on abortion and contraception and did not change its policy on the ordination of women.

individual behavior in line with its theological stances.[16] Consistent with this hypothesis, empirical research has found that declining religiosity is correlated with declining fertility (Frejka & Westoff, 2008; Norris & Inglehart, 2004). This post-Vatican religiosity-fertility connection can also be indirectly examined in the context of shared norms by including church attendance as an additional control in the regression specification.

There are two substantial caveats to this test that bear mentioning. First, as with religion, immigrant church attendance is not directly observed and has to be proxied with European church attendance. Secondly, I cannot actually identify the mechanism through which religiosity affects fertility. While the above argument is inherently about how tightly individuals adhere to church teachings, religiosity could also be indicative of declines in the institutional effect of the Church (Berman, Iannaccone, & Ragusa, 2006). In particular, falling church attendance reflects not just decreasing adherence to the Church, but also decreases in Church-provided social services. Whether church attendance is interpreted as changes in norms or changes in institutions, religiosity and immigrant fertility are only weakly correlated, as Section 1.5 will demonstrate.

1.3 Empirical Estimation

A simple model of the individual's utility maximization problem is:[17]

---

[16] Scholars today still debate whether Vatican II caused declining religiosity because its decrees were too conservative (did not do enough to align Church with contemporary attitudes) or too liberal.
[17] Model simplified from Berman et al. (2006).

$$\max U(C, f, -\pi|f - \phi|)$$

$$\text{s.t. } C = w(T - \lambda f),$$

where C is consumption, w is wage, and T is total hours available. The decision variable f is total fertility, measured in number of children, and λ is the time-cost of raising each child.

The final term in the utility function is interpreted as a "theological premium" and consists of two parameters; ϕ is some "ideal" number of children, and π is the disutility incurred for being away from the ideal.[18] For example, Catholics could have a higher ideal fertility ϕ or place more emphasis on their religious identity with a larger π than non-Catholics. Shifts in Catholic norms could appear as either adjustments in Catholic fertility ϕ or as changes in attachment to Catholic identity π.[19]

Institutional effects differ from shared norms in that they affect prices and only appear in the budget constraint. Like secular institutions, religion can also affect the price of children, either by lowering child-care costs λ (e.g. providing child-friendly social services like daycare or school), affecting labor market opportunities, or lowering the value of those wages by limiting consumption of certain goods.[20]

---

[18] This utility function can be thought of as a specific case of Akerlof and Kranton's (2000) social identity model where agents are assigned to categories by religion and identity prescriptions are the ideal fertility.

[19] While these two channels are described separately in the utility function, they are not individually identified in the empirical work that follows. Hence, "changes in norms" in the remainder of this discussion refer to changes in the overall theological premium and not specifically to either parameter.

[20] Wages could also be affected by, for example, how a woman feels about being a working mother when young children are present or getting more education, which would actually be a norm rather than an institutional effect even though it appears in the budget constraint. As labor force participation, education, marriage, and fertility are all jointly determined and affected by religious teachings, they all fall under the umbrella of "norms" even though they are not explicitly modeled as such. A fully specified model would treat all of these endogenous decisions as choices, each with their own identity weight and idealized outcome.

From this utility-maximization problem, the resulting "demand" for children is a function of market prices (including wages), technology, institutions, and norms:

$$fert_{ict} = f(p_i, tech_t, inst_{ic}, norms_{ic}),$$

where *i* denotes individual, *c* is country of origin, *t* is time, and *p* is a vector of prices. Note the subscripts assigned to the arguments on the right-hand side: only the two final terms have country-of-origin subscripts. The key feature of using second-generation immigrants in the US is that they all share US prices, conditional on observable characteristics, and have access to the same set of US technology. Furthermore, the *only* institutions that could be shared between immigrants and their cousins abroad are those that are international in scope (like religion); in the US, immigrants all experience the same set of (US) institutions. Finally, norms or culture could be shared with their cousins abroad.

A simple estimable version of this demand function is:

$$kids_{ict} = \gamma \cdot TFR_{ct} + \rho_t Z_{ict} + \alpha_t + \lambda_c + \varepsilon_{ict}, \tag{1}$$

where *Z* is a vector of individual characteristics intended to control for prices, α is a cohort fixed effect, and λ are country-of-origin indicators. Ignoring the TFR variable, this is a relatively standard regression equation that relates individual fertility (*kids*) to observable characteristics like age and education (*Z*) with the cohort effects absorbing any secular changes in fertility across time, including those due to advances in technology. The country fixed effect should account for any residual institutional or cultural effects linked to the originating country, provided they are time-invariant.

Bias in the migration decision is always a concern in studies of immigrants. The use of second-generation, rather than new, immigrants minimizes this particular issue since any differences in migration selection or costs would need to have intergenerational effects in order to bias estimates of Eq. (1). The country and cohort fixed effects also help in this regard since this migration selection would also have to change across over time within individual countries. Another potential source of bias with using immigrants is that labor market discrimination in the US might make the included characteristics poor controls for market prices. This particular issue seems unlikely since the estimation sample consists of Western Europeans in the late 20th Century.

Given the inclusion of fixed effects and the use of second-generation women, the coefficient on TFR is identifying how *changes* in fertility in the source-country after emigration predict fertility outcomes today. The literature on immigrant assimilation argues that the longer an immigrant as been in the US, the more her outcomes should look like those of a US native (Chiswick, 1978). As women born in the US have, presumably, been raised here, it follows that any historical cultural links that were not absorbed by the country indicators should be quite weak: $\gamma = 0$.[21] Graphically, γ is interpreted as the slope of the OLS regression line through Figure 1.1, though the estimation is actually conducted by pooling the individual cross-sections. A nonzero γ would be evidence that shifts in norms are being shared between immigrants in the US and European natives.

[21] Cultural effects may depend on reinforcement from the social environment, which would bias this regression toward zero.

As specified, the lone source country covariate is TFR since this preserves the simple interpretation of γ as how observed child-bearing behavior in Europe is, as a whole, correlated with that of immigrants.[22] Eq. (1) can be expanded by including other country-level factors, like female labor force participation or marriage rates, that are jointly determined with fertility, but this lends γ a "partial effect" interpretation. Since fertility is the primary outcome of interest, the inclusion of additional country covariates is left as a robustness check.

The regression specification in Eq. (1) is easily modified to look for religion-specific differences in the cross-country fertility correlation. In the first-difference specification, this would appear as a differential slope for highly-Catholic countries:

$$\Delta y_c = \gamma \cdot \Delta TFR_c + \delta \cdot cath_c \cdot \Delta TFR_c + \beta \cdot cath_c + \Delta\alpha + \Delta u_c , \qquad (2)$$

where $cath_c$ is the time-invariant percentage of Catholics in country c. Note that religion in this equation is a continuous variable, not just an indicator for countries with a Catholic majority.[23] The intercept β gives the religion-specific differential in fertility decline between "Catholic" and non-Catholic second-generation women.

The coefficient of interest here is δ, the coefficient on the Catholic-TFR interaction. A shared religious fertility norm between Catholics in Europe and the United States would appear in the regression as a positive estimate for δ. A shared, pan-Atlantic norm for non-Catholics, evidenced by a positive γ, would be somewhat

---

[22] A more practical reason for this parsimonious specification is that it conserves the already-low number of degrees of freedom.

[23] Since $cath_c$ is a scalar rather than binary, the usual difference-in-difference interpretation of the interaction coefficient is not technically correct. Rather than estimating two regression lines, Eq. (2) estimates a continuum of lines where the slopes and intercepts are constrained to change continuously as the percentage of Catholics increases.

surprising since they do not have explicit pro-fertility theological stances nor the institutionalized social service provision of the Catholic Church. Though the estimating equation is motivated and interpreted in first differences, the estimation is again conducted in a single step using the individual-level data:

$$\begin{aligned} kids_{ict} &= \gamma \cdot TFR_{ct} + \delta \cdot cath_c \cdot TFR_{ct} \\ &\quad + \beta \cdot cath_c \times t + \rho_t Z_{ict} + \alpha_t + \lambda_c + \varepsilon_{ict}. \end{aligned} \tag{3}$$

Note that there are only two cohorts, *t=0* and *t=1*, and the percentage of Catholics is interacted with the cohort indicator since it does not vary over time.[24]

As discussed in the previous section, Vatican II changed many aspects of Roman Catholicism theology and was followed by large declines in religiosity. As individuals attended religious services less frequently, the Vatican became less able to elicit fertility behavior aligned with its pro-fertility teachings In terms of the utility maximization problem, decreasing religiosity appears as a decrease in the theological premium among Catholics. Using church attendance as the measure of religiosity and attachment to the Church, augmenting Eq. (2) with church attendance results in:

$$\begin{aligned} \Delta y_c &= \gamma \cdot \Delta TFR_c + \delta \cdot cath_c \cdot \Delta TFR_c \\ &\quad + \theta \cdot \Delta attend_c + \eta \cdot cath_c \cdot \Delta attend_c \\ &\quad + \beta \cdot cath_c + \Delta\alpha + \Delta u_c, \end{aligned} \tag{4}$$

where $attend_{ct}$ is the fraction of individuals from country *c* and year *t* who said they attended religious services weekly when they were young.[25] Among Catholic countries

---

[24] It might be more natural to model fertility as a discrete count variable and estimate Eq. (3) as a Poisson regression. Estimates from this count model provided similar marginal effects as the OLS presented here.

[25] Childhood, rather than adult, attendance is a more appropriate measure of religious norms since norm formation is more likely to happen during childhood and adolescence than adulthood. Adult attendance is also endogenous; families are more likely to attend religious services when they have children.

with the same change in TFR, ones with large declines in church attendance might also exhibit larger drops in immigrant fertility, signified by a positive η. In this specification, η is interpreted as reflecting changes in shared norms specifically while δ includes all other, non-norm channels that could potentially be shared between Catholics in the US and in Europe (e.g. social service provision). As before, estimation occurs at the individual-level:

$$\begin{aligned} kids_{ict} &= \gamma \cdot TFR_{ct} + \delta \cdot cath_c \cdot TFR_{ct} \\ &+ \theta \cdot attend_{ct} + \eta \cdot cath_c \cdot attend_{ct} \\ &+ \beta \cdot cath_c \times t + \rho_t Z_{ict} + \alpha_t + \lambda_c + \varepsilon_{ict}. \end{aligned} \tag{5}$$

Two details about the interpretation of η are worth mentioning. First, since prior studies have shown the relationship between declining religiosity and fertility in Europe, those declines should *already* be reflected as variation in the source country TFR measure and it does not make sense to include attendance separately. The inclusion of attendance directly into the regression implies that attendance has additional power in predicting immigrant outcomes on top of its effect on European fertility, which would be true if it is a better proxy for shared changes in norms than TFR. In other words, we can think of including both TFR and attendance as running a "horserace" to see which variable is a better predictor of immigrant fertility. Second, the delineation between norms (η) and other, institutional factors (δ) is actually somewhat fuzzy since religiosity could also affect service provision by the Church (the "nuns effect"). I cannot explicitly rule out this particular channel, though it seems somewhat unlikely since it would have to be that church attendance in, for example, Ireland, decreased the number of service-providing nuns in Ireland *and* among Irish-Americans.

One econometric issue that has not been discussed yet is the small sample size. While it is true that the regressions are estimated using thousands of person-level observations, the main variables of interest will only vary by country-of-origin. Rather than having approximately 6000 degrees of freedom, there are really only 20 (10 countries x 2 time periods), and failing to account for this could result in standard errors that are too small (Moulton, 1970). To address this issue, the simplest adjustment would be to cluster all standard errors at the country-of-origin level, which is what is done in the following tables. More conservatively, I could aggregate up to 20 country-time observations, estimate regressions at this level, and use small-sample t-statistics (Donald & Lang, 2007). Given that there are only 10 countries, any one country could be particularly influential in the coefficient estimates. This can easily be address by systematically omitting individual countries from the regression and seeing if the results change dramatically. Both the Donald and Lang aggregation and leave-one-out estimation strategies provide the same results as just simple standard error clustering, so this is all that is reported in the following tables. The alternative is to increase the number of source countries, which is left as a robustness check. Those results suggest that this norm sharing is primarily a European phenomenon, though there are many potential reasons for why Latin America and Asia may not fit the simple model specified here.

1.4 Data

Second-Generation Immigrants

I use two cohorts of immigrants to estimate the above regression. The individual data on second-generation Europeans come from the 1970 US Census (1% Form 2 State sample) and the 2000-2006 March CPS, both retrieved from IPUMS (Ruggles et al., 2008; King et al., 2009).[26] Due to data limitations, the country of origin for second-generation women is identified by the father's place of birth.[27] The use of census data limits the individual-level covariates that are available; the only variables used are a quadratic in age, indicators for educational attainment, metropolitan status, and household income per person (in 1995 dollars). Labor force participation and marital status are both purposely omitted from the main specifications since they are jointly determined by fertility.[28] The sample of countries is limited to only Western Europe in order to minimize concerns about differential migration mechanisms between Westerners and those leaving the Eastern bloc during the 1940s.[29] This has the additional effect of limiting labor market discrimination in the US as these women are predominantly white Europeans.

---

[26] The 1970 Census was the last year that the Census Bureau specifically asks about parental birthplace; from 1980 onward, the question was replaced by one asking about ancestry, which removes the ability to identify immigrant generations. Mothertongue is another potential method of identifying immigrants, but it is not available in the CPS samples.

[27] Since childbearing is a distinctly female phenomenon, it might be more appropriate to use the mother's country of birth. However, the 1970 PUMS only makes maternal country-of-birth available if the father was born in the United States.

[28] The main results are robust to dropping all individual characteristics other than the quadratic in age, which is necessary to account for the fact that census respondents are at different points in their lifecycle.

[29] In particular, the US repealed immigration quotas in 1965 in favor of the current family reunification system. While this could result in differential migration selection patterns among the first generation, I

The dependent variable, the measure of second-generation fertility, is just the number of children currently at home. To minimize undercounting among women whose adult children have since moved out, the estimation sample is limited to women up to age 35. Since this fertility measure has the same interpretation as truncating total lifetime fertility at age 35 rather than 45, country-specific TFRs are also calculated using only age-specific fertility rates up to 35. For details about construction the total fertility rate used, see the Appendix.

Country-Level Covariates

The total fertility rates for countries of ancestry for both 1970 and 2000 are taken from the United Nations' *Demographic Yearbook* series.[30] Because only women up to age 35 are included in the estimation, the measure of European fertility is analogously truncated.

The data on the percentage of Catholics in each country are provided by the Vatican's *Annuarium Statisticum Ecclesiae*, which is the Church's official statistical publication. Within countries, there is little variation in the reported fraction of Catholics, so only the average percentage is used in regression.[31]

---

do not see significant country-specific responses due to changes in the US immigration scheme (see Appendix for details).

[30] Live birth rates from 1970 are taken from Table 24 of the 1975 *Demographic Yearbook, Special Topics Edition*. Birth rates for 2000 are from Table 11 of the 2000 *Demographic Yearbook*.

[31] The Vatican does not document its data gathering methodology, but these statistics are probably calculated using the number of baptisms and deaths in each diocese, rather than from a full census of congregations. Naturally, this raises questions about how accurate the Vatican's counts are since individuals who they consider Catholic may not self-identify as such.

Church attendance rates are from Iannaccone (2003), which were calculated from retrospective questions on church-going behavior in the 1991 and 1998 waves of the International Social Survey Program. The specific measure used here is childhood church attendance, which captures religiosity when individuals are being raised rather than as adults. The one Western European country that is missing attendance data is Greece.

Summary Statistics

Summary statistics for the main variables of interest are in Table 1.1; immigrant fertility is on the left while country-level fertility, percentage of Catholics, and church attendance are on the right. Complete tables of means for all individual- and country-level variables are provided in the Appendix.

Looking first at the left half of the table, the most obvious trend is the sharp secular decline in the number of children borne by second-generation women. In 1970, Western European immigrants had anywhere between 1.2 and 1.8 children per woman; three decades later, only Scandinavians have close to 1 child per woman. After adjusting for age, the differences in fertility between the two cohorts of immigrants is quite striking; immigrant families had 0.8 fewer children per woman, and the countries with the largest declines are Ireland, Spain, Portugal, and Greece.

Total fertility rates in the home country demonstrate a similar trend; TFR has fallen by about 0.7 children per woman; the countries with the largest declines are Ireland, Spain, Portugul, Italy, and Greece. While increased labor force participation and public policy undoubtedly had a role in these fertility decisions, such economic incentives are probably Europe-specific and do not affect American women. That we still see fertility decline among the same set of countries in the US is quite surprising.

1.5 Results

Estimates of Eq. (1), the trans-Atlantic fertility correlation, are displayed in Table 1.2. The first column uses the source country's TFR as the lone explanatory variable. Without accounting for individual-level characteristics or time effects, this simple pooled regression predicts that women with fathers from high-fertility countries are expected to have more children than those from low-fertility countries. The coefficient of 0.616 implies that a one child increase in source country TFR predicts an additional 0.6 children born to second-generation immigrants.

The second column includes the individual-level characteristics Z_{ict}, which all appear with their anticipated signs. Education is negatively correlated with fertility; women without a high school diploma tend to have more children while college graduates tend to have less.[32] Women residing in metropolitan areas have 0.08 fewer children, though this is only marginally significant. Women in higher income households also tend to have fewer children but the effect appears to be quite small; an additional $1000 per family member decreases fertility by 0.03 children. The coefficients for the observable person characteristics change very little even with the inclusion of cohort and country-of-origin indicators in Columns (3) and (4).

[32] The omitted education category is "high school graduate".

Even though the usual Becker-Schultz-Mincer variables are included in these columns, the coefficient on home country TFR is still positive and significant. As emphasized in the previous sections, since the sample consists of second-generation immigrants, it is very unlikely that European social programs and economic conditions are directly driving this positive correlation. Furthermore, the inclusion of country fixed effects should account for unobservables that vary by country of origin but do not evolve much over time, including historical cultural attitudes.

The last column contains estimates of the full regression in Eq. (1), which includes both time and country fixed effects and allows the coefficients on the individual characteristics to change across time. Even in this specification, TFR in the home country is still significantly different from zero. A woman whose father was from a country with, on average, one more child per woman is expected to have 0.3 more children herself even after controlling for the usual socioeconomic variables. A one child increase in source country TFR is quite large; in the 1970s, it is the difference between coming from the lowest fertility country in the sample (Scandinavia) and the highest (Ireland). Given that the average decline in European TFR between 1970 and 2000 was about 0.7 child per woman, this estimate implies that immigrant fertility also fell by about 0.21 (= 0.3 x 0.7) children per woman above and beyond the demographic changes that occurred over that period.

Compared to the coefficients on the individual characteristics, this estimate for TFR is economically quite large. Among the included person characteristics, the previous columns imply education is the biggest determinant of fertility. Compared to a high school graduate, a woman holding a college degree has 0.4 fewer children. The results in the final column suggest that simply coming from a country with high fertility decline has almost as large an effect on individual fertility as increasing educational attainment.

Is this trans-Atlantic fertility correlation explained by a difference in the dominant religion across Western Europe? As Catholics tend to have higher fertility than Protestants and religion is highly correlated with country of origin, it could be that religion is the underlying cause of this immigrant-native fertility correlation. Table 1.3 includes the two religious indicators (the percentage of Catholics and historical church attendance) to the full specification. For convenience, Column (1) reproduces the estimates from Column (5) of the previous table.

Column (3) contains the estimates of Eq. (3). The two religion coefficients, β and δ, indicate that the immigrant-native fertility relation appears to be different for Catholic and non-Catholic countries-of-origin. The fertility of a second-generation woman is predicted to be anywhere from 0.28 to 0.73 (= 0.278 + 0.454) higher depending on whether her father came from a very Catholic country. In the context of the first difference equation in Eq. (2), this estimate implies that the post-1970s fertility decline among "Catholic" immigrants was much larger than that for non-Catholics.

Are these shared changes in fertility among Catholics related to the institutional shock of Vatican II? The estimates in Column (4) imply that declining church attendance predicts declining fertility among both Protestants and Catholics, though the connection between religiosity and fertility outcomes is much higher for Catholics. Church attendance declined by 17% on average between 1970 and 2000, so evaluated at the mean, the results imply that fertility among non-Catholics second-generation women fell by 0.07 (= 0.386 x -0.17) children per woman. This same decline in church attendance among Catholics predicts a much larger drop in fertility of 0.13 children per woman.

If we interpret religiosity as reflecting shared norms and TFR as reflecting all other shared religious factors, these results are consistent Vatican II affecting both immigrants and European natives by reducing the weight of Church theology in individual preferences. However, shifts in religiosity and, by extension, shared norms, explain only part of the post-Vatican II fertility decline in Europe and the US. The Catholic-TFR interaction term is still positive and significant even with the inclusion of religiosity, so there is still the possibility for institutional changes, in particular social service provision, to be shared across the Atlantic.

The coefficient on TFR in Columns (3) and (4) are still non-zero, suggesting that fertility declines are also shared between Protestant Europe and "Protestant" immigrants. As both Protestant shared norms and institutional services should be much weaker than those for Catholics, this estimate is somewhat surprising. As will be shown in the next section, this coefficient drops to zero once individual marital status and labor force participation are included.

Marital Status and LFP

The set of included individual-level covariates in Eq. (1) and its extensions is quite small; in particular, marital status and labor force participation are entirely omitted. It could be that the observed fertility correlation is actually due to, for example, falling marriage rates among both Europeans and second-generation immigrants. Including individual marital status as a regressor would partial out this particular pathway connecting trans-Atlantic fertility behavior, so it is purposely excluded from regression. Table 1.4 repeats the analysis with the inclusion of observed marital status and LFP.

The first three columns correspond to estimates of Eq. (1) where total fertility rate is the only country-level regressor. Both marital status and labor force participation have large coefficients, are highly significant and of the anticipated sign. According to Column (3), second-generation women who work tend to have 0.4 fewer children; married women are predicted to have 0.8 more. Both of these estimates are larger in magnitude than the coefficient on TFR, but home country fertility is still significant in predicting immigrant fertility. Nor did the magnitude of this correlation change much; a one child increase in source country fertility predicts an additional 0.303 children born to second-generation immigrants, which is very close to the original estimate of 0.296 from Column (1) of Table 1.3. While individual marital and labor force status are important determinants of immigrant fertility, they cannot fully account for why second-generation immigrants exhibit a fertility trend that is similar to that of Europeans.

The final 3 columns of Table 1.4 include both church attendance and the percentage of Catholics. In this specification, uninteracted TFR is no longer significant; the similarity in fertility between Protestants across countries shared norm can be almost entirely explained by marriage and work characteristics.[33] The change in magnitude on this slope implies that immigrants from high fertility countries tend to have high fertility characteristics (i.e. they are more likely to be married and less likely to work, both of which predict more children). It could be that among mainline Protestants, family size is not very important to religious identity since these denominations tend to emphasize the individual (including the individual's choice to use contraception) rather than strict obedience to the church.

The Catholic-TFR interaction in Columns (4)-(6), however, remains significant and is actually much larger in magnitude than the original estimate ($\delta = 1$ instead of 0.7). The increase in magnitude implies that immigrants from high-fertility Catholic Europe are more likely to be active in the labor force and less likely to be married. The directionality is consistent with the migrant selection literature that finds immigrants to be positively selected in earnings potential. The fact that the correlation among Catholics survives the addition of these person characteristics suggests that family size itself, rather than just marriage or labor force participation, is an important portion of Catholic norms and identity.

Church attendance still matters but its estimate is no longer significantly different between non-Catholics and Catholics. As the coefficients on attendance were

[33] There are two extensions that can be explored here. One, I could include European marriage and LFP as additional cultural proxies. Two, the outcome of interest for both immigrants and Europeans could both be changed to marriage or LFP. Both of these are provided in the Appendix.

only weakly significant in the previous table, it is not too surprising that including two very relevant regressors (and absorbing additional variation in the dependent variable) results in estimates that are not significantly different from zero. The estimate of 0.769, however, is very close to that on the Catholic-attendance interaction in the previous table (η = 0.770).

Expanded sample of countries

The positive fertility correlation between second-generation women and non-immigrant cousins is not only limited to Western Europeans, but it is more difficult to distinguish a religious differential. Figure 1.3 expands the plot of changes in immigrant fertility against source country fertility to a total of 24 countries. Whereas the original 10 countries were exclusively from Western Europe, this sample includes women from Eastern Europe, Asia, Latin America, and Canada. The graph shows that Eastern Europeans also demonstrate this pan-Atlantic fertility connection; the six new data points lie almost exclusively between the pre-existing Western European observations.[34] Austria is the one outlier, but its position pulls the coefficient in the estimated OLS results toward zero. Latin America, however, exhibits much larger declines in home fertility than any European country, likely due to economic growth and increasing per capita incomes rather than any religious factors. Rapid development and increases in overall schooling could also change the dynamics of migration selection, further biasing the results.

[34] These are Czechoslovakia, Austria, Hungary, Poland, Russia, and Ukraine. Of these, Russia and Ukraine are predominantly Orthodox Christian. Poland is highly Catholic (95%) while Austria and Hungary have Catholic majorities.

These observations are borne out by the estimates in Table 1.5. Including the six Eastern European countries does not noticeably change any of the coefficient estimates, either in magnitude or significance. The coefficient on source country fertility rate decreases slightly from 0.278 to 0.254 while the Catholic-TFR interaction is larger in magnitude (up to 0.509 from 0.455) but also less precisely estimated.

In contrast, the inclusion of Latin America causes the slope on total fertility rate to drop to zero (in Column (5)). Latin American immigrants likely differ from Europeans in unobservables, both in terms of migrant selection and labor market opportunities, so the small set of included individual-level characteristics included here could result in a much larger omitted variable bias for these women. While this was less of a concern for European descendants, labor market discrimination is potentially a much more significant issue since it is likely to affect the types of jobs held and the earnings power of Hispanics in the US. In addition, immigration patterns in the US have shifted since 1970; before, entrants were mostly from Europe, but contemporary cohorts are predominantly from Mexico and Puerto Rico.[35]

The case of Latin America presents two issues worth future consideration. First, on paper, Latin America is also predominantly Roman Catholic, but the trajectory of the Catholic Church after Vatican II was quite different in this region. Rather than the sharp declines in institutional strength that hit Europe, the number of priests, nuns, and seminarians all rose in the 1980s and 1990s, allowing the Church to maintain its network of social services (Hagopian, 2009). There is also the question of

[35] In the 2000 sample, Mexico alone provides 3,743 observations, which is more than all the Western European nations combined. In the 1970 sample, Mexico was not even the largest source country.

whether rapid economic development, and increases in female educational attainment in particular, in the 1980s and 1990s in these countries also caused source-country fertility to decline dramatically. These arguments suggest that institutional decline of the Catholic Church was, if anything, muted in Latin America compared to Europe.

To test if Catholic norms transmission is unique to Western Europe, the final column of the table interacts each of the three fertility variables with an indicator for if the source country is Western European (e.g. is used in Columns (1) and (2)). Most notable in these estimates is the fact that the Catholic-TFR interaction is only significant for Western European immigrants, suggesting that the trans-Atlantic norms shift posited here appears to be specifically a European phenomenon. While Latin America may not fit the pattern demonstrated by the European countries, it does substantially increase the number of countries, which should alleviate some of the concern about the low number of clusters.

Intergenerational Transmission

The discussion so far has been about contemporaneous shared norms, though there are other channels through which immigrant outcomes could be related to those abroad. Since the economic literature has examined norms as the correlation between *current* immigrant outcomes and *past* source-country TFR, this specification is a logical starting point that contextualizes this paper's results with other studies. To explore the role of current and past behavior on outcomes, I augment the original regression in Eq. (1) with lagged values of fertility:

$$kids_{ict} = \gamma_1 TFR_{ct} + \gamma_0 TFR_{c,t-1} + \phi_0 kids_{c,t-1} + \rho_t Z_{ict} + \varepsilon_{ict} \,. \qquad (6)$$

The dependent variable, $kids_{ict}$, remains the same as before, as does TFR_{ct}. Since this regression involves lags, I can only estimate Eq. (6) using the 2000 CPS immigrant cohort.

The slopes γ_0 and ϕ_0 correspond to intergenerational links to second-generation fertility. For woman *i* observed in year *t*, the coefficient γ_0 relates the previous generation's TFR (i.e. home country fertility at *t-1*) to current fertility while ϕ_0 relates (second-generation) immigrant fertility today to (second-generation) immigrant fertility last period. Fernandez and Fogli specifically estimate the specification of (6) that only includes γ_0 and find a coefficient of about 0.22-0.25.[36] In contrast, the key coefficient here is going to be γ_1: how do *current* fertility outcomes predict to *current* immigrant behavior? Both γ_0 and ϕ_0 demonstrate persistence in unobserved norms, but γ_1 signals that norms even post-emigration are also being shared.

Table 1.6 presents various specifications of Eq. (6), all of which indicate significant intergenerational fertility patterns. In Column (1), the coefficient on TFR last period is negative but not significant. Fertility of the parents' generation in Europe does not significantly predict the childbearing behavior of second-generation women in 2000, at least for the ten countries used here. While it is a statistical zero, it is very surprising to see that this coefficient has the opposite sign of Fernandez and Fogli, though they have a much larger number of countries and a different time period (1970 immigrants, 1950 source-country TFR). In contrast, the lagged immigrant fertility

---

[36] I am not aware of any work that focuses specifically on intergenerational persistence as captured by ϕ_0.

positively predicts current immigrant fertility today. Second-generation women from countries with historically high second-generation fertility also tend to have more children today.[37]

The last column presents a surprise; γ_1 is positive and significant even when both intergenerational fertility measures are included. Shared norms between cousins and second-generation women are evident even after controlling for historical fertility both abroad and in the US. Historical TFR should control for the institutional and policy environment in the source-country at the time of emigration; if it is the case that policy changes in Europe cannot affect Americans, then γ_1 truly is purged of any lingering intergenerational spillover effects between immigrants and natives.

The regression in Eq. (6) is somewhat cumbersome to interpret due to the lags and individual-level regressors. A more straightforward interpretation uses first-differences:

$$y_{ct} = \gamma \cdot \Delta TFR_{ct} + y_{c,t-1} + u_{ct}\,, \tag{7}$$

which implicitly assumes that $\phi_0 = 1$ and that $\gamma_0 = -\gamma_1$.[38] The slope γ from Eq. (7) is actually just the slope coefficient originally discussed at length in Section 1.3. The last two rows of Table 1.6 test the two assumptions implicit in the first-difference specification in Eq. (7). The data available fail to reject that the two γ coefficients are equal but opposite in sign (p-value = 0.54), but do reject that $\phi_0 = 1$, though not strongly. This slight misspecification suggests caution in reading the original

[37] Using fertility among 1970 first, rather than second, generation immigrants does not substantially change these results.
[38] Note the change in dependent variable from kids to y; this is analogous to regression-adjusting the raw measure of immigrant fertility (kids) for shifts in a set of demographic variables, denoted Z.

regression in Eq. (1), though it does benefit from a simpler interpretation and the additional 5,000 Census observations.

1.6 Conclusion

Economic explanations for fertility have traditionally been concerned with prices and institutions, with a recent literature examining the role of norms. Like similar studies, this paper leverages the portability of culture across geographic areas to isolate an effect for shared norms, as opposed to prices or institutions, on fertility. The key findings of this paper are: (1) there is a significant positive correlation in fertility declines between immigrants and Europeans that appears to be driven by shared norms across the ocean, and (2) the strength of this correlation is higher among "Catholic" immigrants and persists even after controlling for marital status and labor force participation. Shared religion, which is unobserved in US census data, is a potential mechanism through which *changes* in norms might be transmitted between immigrants and Europeans. While the Catholic Church is also a potent provider of social services in both the US and Europe, the shared trans-Atlantic fertility trends examined here appear to be driven by declining adherence to Church fertility norms rather than simultaneous declines in the institutional strength of the Church.

Given the economic consequences of low fertility levels, many developed countries have spent considerable energy confronting the challenge of declining birthrates. However, fertility has proven to be a remarkably difficult behavior to alter without resorting to drastic legislative actions like China's "one-child" policy or

Romania's bans on birth control and abortion in the 1970s and 1980s (Pop-Eleches, 2010). While the discussion here does not provide much in the way of feasible population interventions, it strongly suggests that cultural/religious preferences play an integral role in child-bearing decisions of Americans, a feature that policymakers should bear in mind when deciding how to address stagnant population growth. More generally, future work should continue to consider how heterogeneity in preferences due to norms and religion affect culturally relevant outcomes of individuals and households.

1.7 Acknowledgement

Chapter 1 also benefited from discussions with John H. Evans, Daniel Egel, L. Choon Wang, Sarada, Youjin Hahn, Oana Hirakawa, Philip Neary, and Hee-Seung Yang. Valuable comments were also provided by seminar participants at UC San Diego, Monash University, the Office of Economic Policy at the Treasury Department, and attendees of the 2010 ASREC meeting and the 2010 All-California Labor Conference.

Table 1.1: Summary Statistics, Condensed

	Second-Generation Immigrants: Average Number of Children			Source Country: Total Fertility Rate			% Catholics	Church Attendance	
Country	1970	2000	Δ_{adj}	1970	2000	Δ		1970	2000
UK	1.517	0.771	-0.878	2.165	1.576	-0.589	0.085	0.56	0.28
Ireland	1.684	0.806	-0.971	2.952	1.775	-1.177	0.761	0.97	0.92
Scandinavia	1.569	1.024	-0.708	1.835	1.699	-0.136	0.010	0.19	0.12
Germany	1.379	0.819	-0.713	1.756	1.265	-0.491	0.318	0.5	0.36
France	1.219	0.674	-0.620	2.171	1.767	-0.404	0.755	0.62	0.33
Spain	1.451	0.423	-1.104	2.289	1.221	-1.068	0.878	0.77	0.55
Portugal	1.824	0.669	-0.951	2.325	1.266	-1.059	0.904	0.72	0.61
Italy	1.522	0.711	-0.748	2.02	1.182	-0.838	0.966	0.86	0.77
Greece	1.194	0.482	-0.909	2.061	1.193	-0.868	0.010	0.64	0.37
Netherlands	1.538	0.824	-0.816	2.263	1.640	-0.623	0.313	0.56	0.28
Overall	1.506	0.733	-0.801	2.184	1.458	-0.725	0.500	0.639	0.459

Notes: Second-generation immigrants are identified as women born in the US to foreign-born fathers. For immigrants, the 1970 and 2000 columns report unadjusted means while the Δ_{adj} column reports differences after adjustiong reported fertility for age. Data on source country fertility are from the United Nations' Demographic Yearbook series and calculated as the sum of age-specific fertility rates for women 15-35. Data on the percentage of Catholics in the source country are from the Vatican's *Annuarium Statisticum Ecclesiae* while church attendance comes from Iannaccone (2003).

Table 1.2: Trans-Atlantic Fertility Correlation

	(1)	(2)	(3)	(4)	(5)
TFR in home country	0.616^{**}	0.255^{**}	0.169^{**}	0.394^{**}	0.296^{**}
	(0.148)	(0.058)	(0.042)	(0.106)	(0.090)
Age	-	0.306^{**}	0.298^{**}	0.299^{**}	0.400^{**}
		(0.029)	(0.032)	(0.030)	(0.026)
Age^2	-	-2.967^{**}	-2.823^{**}	-2.829^{**}	-4.552^{**}
		(0.586)	(0.638)	(0.610)	(0.594)
Less than HS grad	-	0.070	0.069	0.071	-0.026
		(0.046)	(0.046)	(0.047)	(0.042)
Some college	-	-0.228^{**}	-0.207^{**}	-0.214^{**}	-0.133^{**}
		(0.021)	(0.026)	(0.029)	(0.035)
College graduate	-	-0.415^{**}	-0.391^{**}	-0.399^{**}	-0.376^{**}
		(0.045)	(0.039)	(0.038)	(0.043)
Lives in metro area	-	$-0.077^{\dagger}$	$-0.072^{\dagger}$	-0.043^{**}	0.022
		(0.033)	(0.033)	(0.032)	(0.032)
Income per person	-	-0.033^{**}	-0.033^{**}	-0.033	-0.055^{**}
		(0.003)	(0.003)	(0.003)	(0.004)
$t \times$ Age	-	-	-	-	-0.320^{**}
					(0.066)
$t \times Age^2$	-	-	-	-	4.933^{**}
					(1.367)
$t \times$ Less than HS	-	-	-	-	$0.315^{\dagger}$
					(0.149)

Table 1.2: Trans-Atlantic Fertility Correlation, Continued

	(1)	(2)	(3)	(4)	(5)
t × Some college	-	-	-	-	0.080
					(0.059)
t × College graduate	-	-	-	-	0.229*
					(0.101)
t × Metro	-	-	-	-	-0.160*
					(0.068)
t × Income per person	-	-	-	-	0.041**
					(0.004)
Includes cohort fixed effect?			X	X	X
Includes country fixed effects?				X	X
R^2	0.036	0.418	0.419	0.422	0.464
Number of observations	6198	6198	6198	6198	6198

Notes: Regression equation is Eq. (1) in the text. The dependent variable is the number of children currently at home under age 18. The omitted education category is high school graduates. Estimates of country and time fixed effects are omitted for brevity. Robust SEs, clustered by country of origin, in parentheses. Symbols denote coefficients significantly different from zero at 10%(†), 5%(*), and 1%(**).

Table 1.3: Fertility, Religion, and Church Attendance

	(1)	(2)	(3)	(4)
Total fertility rate (γ)	0.296**	0.495**	0.278**	0.172†
	(0.090)	(0.080)	(0.053)	(0.079)
%Catholics $\times$ TFR (δ)	-	-	0.454*	0.738**
			(0.161)	(0.127)
%Catholics $\times$ t (β)	-	0.213**	0.556**	0.808**
		(0.059)	(0.133)	(0.060)
attendance (θ)	-	-	-	0.386†
				(0.177)
attendance $\times$ %Catholics (η)	-	-	-	0.770*
				(0.291)
R^2	0.464	0.464	0.464	0.464
Number of observations	6198	6198	6198	5902

Notes: Data on second-generation immigrants from ten source countries from the 1970 Census and 2000-2006 March CPS. TFR from the *Demographic Yearbook* (various years), percentage of Catholics from *Annuarium Statisticum Ecclesiae* (various years), and attendance from Iannaccone (2003). The dependent variable is the number of children currently at home under age 18. All regressions include individual-level characteristics and both country and time fixed effects. Robust SEs, clustered by country of origin, in parentheses. Symbols denote coefficients significantly different from zero at 10%(†), 5%(*), and 1%(**).

Table 1.4: Including Individual Marital Status and Labor Force Participation

	(1)	(2)	(3)	(4)	(5)	(6)
Total fertility rate (γ)	0.324*	0.285*	0.303*	0.047	0.051	-0.012
	(0.102)	(0.106)	(0.112)	(0.083)	(0.071)	(0.081)
%Catholics × TFR (δ)	-	-	-	0.987**	0.941**	1.069**
				(0.138)	(0.115)	(0.128)
%Catholics × t (β)	-	-	-	0.911**	0.876**	0.931**
				(0.056)	(0.067)	(0.062)
attendance (θ)	-	-	-	0.671**	0.623**	0.769**
				(0.185)	(0.138)	(0.162)
attendance × %Catholics (η)	-	-	-	0.357	0.137	-0.040
				(0.279)	(0.246)	(0.263)
Currently married?	-	0.966**	0.847**	-	0.969**	0.852**
		(0.055)	(0.046)		(0.057)	(0.048)
Currently in labor force?	-0.589**	-	-0.365**	-0.585**	-	-0.359**
	(0.056)		(0.046)	(0.059)		(0.048)
R^2	0.498	0.538	0.549	0.497	0.537	0.549
Number of observations	6198	6198	6198	5902	5902	5902

Notes: Data on second-generation immigrants from ten source countries from the 1970 Census and 2000-2006 March CPS. TFR from the *Demographic Yearbook* (various years), percentage of Catholics from *Annuarium Statisticum Ecclesiae* (various years), and attendance from Iannaccone (2003). The dependent variable is the number of children currently at home under age 18. All regressions include individual-level characteristics and both country and time fixed effects. Robust SEs, clustered by country of origin, in parentheses. Symbols denote coefficients significantly different from zero at 10%(†), 5%(*), and 1%(**).

Table 1.5: Additional Countries

	Western Europe		+ Eastern Europe		+ Latin America/Asia	
	(1)	(2)	(3)	(4)	(5)	(6)
Total fertility rate (γ)	0.296**	0.278**	0.335**	0.254**	0.054	-0.263
	(0.090)	(0.053)	(0.085)	(0.066)	(0.042)	(0.262)
TFR × I{W. Europe}	-	-	-	-	-	0.035
						(0.134)
%Catholics × TFR (δ)	-	0.455*	-	0.509*	-	0.370
		(0.161)		(0.197)		(0.358)
%Catholics × TFR × I{W. Europe}	-	-	-	-	-	0.895**
						(0.268)
%Catholics × t (β)	-	0.556**	-	0.558**	-	0.309
		(0.133)		(0.173)		(0.359)
%Catholics × t × I{W. Europe}	-	-	-	-	-	0.852**
						(0.206)
R^2	0.464	0.464	0.452	0.452	0.447	0.448
Number of countries	10	10	16	16	24	24
Number of observations	6198	6198	8373	8373	18645	18645

Notes: Data on second-generation immigrants from ten source countries from the 1970 Census and 2000-2006 March CPS. TFR from the *Demographic Yearbook* (various years), percentage of Catholics from *Annuarium Statisticum Ecclesiae* (various years), and attendance from Iannaccone (2003). The dependent variable is the number of children currently at home under age 18. In Column (6), Western Europe is defined as the 10 countries that appear in Columns (1) and (2). All regressions include individual-level characteristics and both country and time fixed effects. Robust SEs, clustered by country of origin, in parentheses. Symbols denote coefficients significantly different from zero at 10%(†), 5%(*), and 1%(**).

Table 1.6: Intergenerational Fertility Patterns

	(1)	(2)	(3)	(4)
Total fertility rate (γ_1)	-	-	-	0.373**
				(0.070)
TFR, previous generation (γ_0)	-0.110	-	-0.264**	-0.330**
	(0.082)		(0.080)	(0.072)
Immigrant fertility, previous gen. (ϕ_0)	-	0.536*	0.934**	0.432
		(0.236)	(0.270)	(0.247)
R^2	0.308	0.310	0.314	0.316
Number of observations	1350	1350	1350	1350
p-value for test: $\phi_0 = 1$	-	0.081	0.810	0.047
p-value for test: $\gamma_1 = -\gamma_0$	-	-	-	0.538

Notes: Estimates of Eq. (6) using observations from 10 countries of origin, pooled from 2000-2006 March CPS. All regressions include individual-level covariates. Immigrant fertility last period is the regression-adjusted number of children born to second-generation by age 35, calculated from the 1970 Census. TFR data from the *Demographic Yearbook* (various years). The dependent variable is the number of children currently at home under age 18. Robust SEs, clustered by country of origin, in parentheses. Symbols denote coefficients significantly different from zero at 10%(†), 5%(*), and 1%(**).

Table 1.7: Untruncated TFR

	(1)	(2)	(3)	(4)
Total fertility rate (γ)	0.186**	0.294**	0.191**	-0.136
	(0.047)	(0.062)	(0.054)	(0.093)
%Catholics $\times$ TFR (δ)	-	-	0.164	0.646**
			(0.115)	(0.138)
%Catholics $\times$ t (β)	-	0.193**	0.347*	0.609**
		(0.040)	(0.127)	(0.117)
attendance (θ)	-	-	-	1.017**
				(0.171)
attendance $\times$ %Catholics (η)	-	-	-	-0.240
				(0.528)
R^2	0.464	0.464	0.464	0.464
Number of observations	6198	6198	6198	5902

Notes: Data on second-generation immigrants from ten source countries from the 1970 Census and 2000-2006 March CPS. TFR from the *Demographic Yearbook* (various years), percentage of Catholics from *Annuarium Statisticum Ecclesiae* (various years), and attendance from Iannaccone (2003). The dependent variable is the number of children currently at home under age 18. All regressions include individual-level characteristics and both country and time fixed effects. Robust SEs, clustered by country of origin, in parentheses. Symbols denote coefficients significantly different from zero at 10%(†), 5%(*), and 1%(**).

Table 1.8: Pre-Post 1965 Immigrant Cohort Characteristics

A: Men

		UK		Italy		Germany		Total	
	Entry cohort:	40s	70s	40s	70s	40s	70s	40s	70s
	Age	54.830	44.180	54.769	44.249	55.827	44.987	55.276	44.340
Education:	Less than HS	0.170	0.022	0.738	0.425	0.235	0.077	0.376	0.171
	HS grad	0.426	0.164	0.154	0.365	0.367	0.167	0.314	0.234
	Some college	0.170	0.291	0.062	0.120	0.173	0.286	0.138	0.231
	BA+	0.234	0.523	0.046	0.089	0.224	0.470	0.171	0.364
	In labor force	0.851	0.970	0.938	0.936	0.980	0.927	0.938	0.951
Marital status:	Single	0.106	0.052	0.046	0.039	0.031	0.085	0.052	0.053
	Married	0.830	0.806	0.923	0.909	0.908	0.748	0.895	0.832
	Entry post-1945	0.723		0.462		0.316		0.452	
	# observations	47	677	65	482	98	234	210	1393

B: Women

		UK		Italy		Germany		Total	
	Entry cohort:	40s	70s	40s	70s	40s	70s	40s	70s
	Age	54.230	44.073	54.654	44.299	54.743	44.429	54.550	44.237
Education:	Less than HS	0.339	0.034	0.942	0.580	0.395	0.088	0.466	0.175
	HS grad	0.422	0.361	0.058	0.249	0.362	0.392	0.332	0.346
	Some college	0.138	0.371	0	0.111	0.171	0.343	0.131	0.303
	BA+	0.101	0.234	0	0.059	0.072	0.178	0.070	0.176
	In labor force	0.569	0.742	0.423	0.595	0.559	0.712	0.540	0.699
Marital status:	Single	0.064	0.032	0	0.037	0.026	0.018	0.035	0.028
	Married	0.725	0.780	0.865	0.894	0.783	0.754	0.776	0.797
	# Children	2	1.907	2.327	2.363	1.428	1.921	1.776	2.015
	Entry post-1945	0.679		0.423		0.434		0.518	
	# observations	109	817	52	405	152	569	313	1791

Notes: 1940's entry cohort identified as individual born abroad and were aged 50-60 in the 1970 Census, Form 1 State sample. 1970's cohort consists of individuals born abroad and were aged 40-50 in the 1990 Census, 5% sample. The "Total" column combines immigrants from these three countries.

Table 1.9: Summary Statistics, Second-Generation Women age 18-35, 1970

				Education								
Country	Obs.	# kids at home	Age	≤ HS	HS grad	Some college	BA+	Not in LF	Never married	Living at home	In metro area	HH income
US (comparison)	225363	1.353	25.44	0.275	0.459	0.179	0.087	0.508	0.277	0.199	0.719	11.967
UK	634	1.517	27.62	0.177	0.457	0.216	0.150	0.506	0.241	0.186	0.855	15.587
Ireland	553	1.684	28.33	0.170	0.570	0.150	0.110	0.562	0.271	0.203	0.926	15.138
Scandinavia	445	1.569	28.07	0.137	0.488	0.200	0.175	0.524	0.178	0.115	0.787	15.403
Germany	845	1.379	27.47	0.161	0.496	0.192	0.151	0.553	0.243	0.166	0.846	15.114
France	73	1.219	27.08	0.219	0.398	0.164	0.219	0.548	0.301	0.178	0.863	15.964
Spain	71	1.451	27.52	0.254	0.436	0.141	0.169	0.521	0.225	0.239	0.831	13.121
Portugal	102	1.824	28.18	0.382	0.568	0.137	0.049	0.539	0.216	0.225	0.824	11.245
Italy	1778	1.522	28.22	0.210	0.582	0.124	0.084	0.566	0.238	0.215	0.898	13.468
Greece	211	1.194	27.75	0.109	0.459	0.223	0.209	0.469	0.261	0.223	0.929	16.487
Netherlands	136	1.538	27.12	0.154	0.559	0.199	0.088	0.537	0.257	0.199	0.787	13.318
average, immigrants	4848	1.506	27.93	0.184	0.526	0.165	0.124	0.545	0.239	0.192	0.871	14.513
SD, immigrants		1.523	5.013	0.388	0.500	0.392	0.330	0.498	0.427	0.394	0.335	10.680

Notes: Data are from the 1970 Decennial Census. For comparison, the row labeled "US" calculates means for women who have been in the United States for at least two generations. Foreign-born women are omitted. Income is measured in thousands of 1995 dollars, per family member.

Table 1.10: Summary Statistics, Second-Generation Women age 18-35, 2000-2006

Country	Obs.	# kids at home	Age	Education: ≤ HS	Education: HS grad	Education: Some college	Education: BA+	Not in LF	Never married	Living at home	In metro area	HH income
US (comparison)	117906	0.988	26.78	0.104	0.311	0.363	0.223	0.260	0.479	0.229	0.752	16.407
UK	170	0.771	26.83	0.088	0.247	0.335	0.329	0.247	0.500	0.294	0.894	22.842
Ireland	93	0.806	28.94	0	0.086	0.398	0.516	0.269	0.430	0.172	0.968	23.841
Scandinavia	41	1.024	28.220	0.049	0.122	0.488	0.341	0.244	0.366	0.073	0.829	23.068
Germany	304	0.819	26.80	0.076	0.188	0.398	0.339	0.227	0.500	0.247	0.822	20.003
France	46	0.674	24.70	0.043	0.435	0.326	0.196	0.239	0.652	0.435	0.783	20.566
Spain	52	0.423	27.27	0.0769	0.250	0.423	0.250	0.192	0.558	0.212	0.827	20.685
Portugal	142	0.669	26.32	0.085	0.246	0.493	0.176	0.204	0.507	0.317	0.951	17.514
Italy	349	0.711	28.09	0.043	0.186	0.318	0.453	0.221	0.490	0.266	0.968	24.917
Greece	85	0.482	26.60	0.059	0.165	0.518	0.259	0.247	0.612	0.294	0.918	19.416
Netherlands	68	0.824	28.06	0.044	0.176	0.338	0.441	0.206	0.412	0.118	0.794	28.735
average, immigrants	1350	0.733	27.27	0.060	0.201	0.385	0.354	0.228	0.499	0.256	0.896	22.175
SD, immigrants		1.021	5.362	0.238	0.401	0.487	0.478	0.420	0.500	0.437	0.305	21.291

Notes: Data are from the 2000-2006 March Current Population Surveys. For comparison, the row labeled "US" calculates means for women have been in the United States for at least two generations. Foreign-born women are omitted. Income is measured in thousands of 1995 dollars, per family member.

Table 1.11: Summary Statistics, Country-Level Covariates

			1970		2000		Δ	
		# obs	Mean	SD	Mean	SD	Mean	SD
Fertility:	Total fertility rate	10	2.563	0.530	1.507	0.265	-1.056	0.515
	TFR, age ≤ 35	10	2.184	0.329	1.458	0.253	-0.726	0.333
Labor force participation:	LFP, all women	10	0.234	0.068	0.456	0.085	0.222	0.087
	LFP, young adults (age ≤ 35)	10	0.604	0.073	0.625	0.099	0.021	0.091
	LFP, young women	10	0.453	0.089	0.574	0.104	0.121	0.107
% never married:	never married, young adults	10	0.554	0.063	0.715	0.061	0.160	0.072
	never married, young women	10	0.480	0.069	0.650	0.072	0.170	0.090
Religion:	% of Catholics	10	0.500	0.391	-	-	-	-
	% attended Church when young	9	0.648	0.226	0.479	0.254	-0.169	0.097

Notes: Statistics are unweighted; each country is treated as one observation. The United States is omitted. Data on fertility and marital status are from the United Nations' Demographic Yearbook series. Labor force participation statistics are from the ILO's LABORSTA database. Data on the percentage of Catholics are from the Vatican's *Annuarium Statisticum Ecclesiae* while church attendance comes from Iannaccone (2003).

Table 1.12: Including Country-Level Work and Marriage Norms

	(1)	(2)	(3)	(4)	(5)	(6)
Total fertility rate (γ)	0.248^{**}	0.253^{*}	0.262^{**}	0.052	$0.176^{\dagger}$	0.043
	(0.067)	(0.096)	(0.084)	(0.069)	(0.084)	(0.078)
LFP, young women	-0.611^{*}	-	-0.620^{**}	0.173	-	0.175
	(0.226)		(0.239)	(0.099)		(0.103)
% never married, young women	-	0.202	-0.069	-	0.012	-0.028
		(0.407)	(0.292)		(0.224)	(0.230)
%Catholics × TFR (δ)	-	-	-	0.979^{**}	0.729^{**}	1.002^{**}
				(0.134)	(0.160)	(0.178)
%Catholics × t (β)	-	-	-	0.926^{**}	0.802^{**}	0.940^{**}
				(0.080)	(0.100)	(0.117)
attendance (θ)	-	-	-	0.716^{**}	$0.384^{\dagger}$	0.724^{**}
				(0.166)	(0.173)	(0.169)
attendance × %Catholics (η)	-	-	-	0.565^{*}	$0.764^{\dagger}$	$0.577^{\dagger}$
				(0.210)	(0.342)	(0.258)
R^2	0.463	0.464	0.464	0.464	0.464	0.464
Number of observations	6198	6198	6198	5902	5902	5902

Notes: Data on second-generation immigrants from ten source countries from the 1970 Census and 2000-2006 March CPS. TFR and rates of singlehood from the *Demographic Yearbook* (various years), percentage of Catholics from *Annuarium Statisticum Ecclesiae* (various years), and attendance from Iannaccone (2003). Country-level data on labor force participation are from the ILO's LABORSTA database. The dependent variable is the number of children currently at home under age 18. All regressions include individual-level characteristics and both country and time fixed effects. Robust SEs, clustered by country of origin, in parentheses. Symbols denote coefficients significantly different from zero at 10%(†), 5%(*), and 1%(**).

Table 1.13: Labor Force Participation and Marital Status as Outcomes

	Y = I{in labor force}		Y = I{ever married}	
	(1)	(2)	(3)	(4)
LFP, young women	-0.175 (0.110)	-0.354** (0.141)	-	-
%Catholics x LFP	-	0.632* (0.279)	-	-
Fraction of singles, young women	-	-	0.129 (0.167)	0.128 (0.350)
%Catholics × singlehood	-	-	-	0.121 (0.463)
%Catholics × *t*	-	0.005 (0.030)	-	0.003 (0.096)
R^2	0.065	0.065	0.053	0.053
Number of observations	6787	6787	6787	6787

Notes: Data on second-generation immigrants from ten source countries from the 1970 Census and 2000-2006 March CPS. TFR and rates of singlehood from the *Demographic Yearbook* (various years), percentage of Catholics from *Annuarium Statisticum Ecclesiae* (various years), and attendance from Iannaccone (2003). Country-level data on labor force participation are from the ILO's LABORSTA database. The dependent variable is either an indicator for being in the labor force or for being married. All regressions include individual-level characteristics and both country and time fixed effects. Robust SEs, clustered by country of origin, in parentheses. Symbols denote coefficients significantly different from zero at 10%(†), 5%(*), and 1%(**).

Table 1.14: Married Women

	(1)	(2)	(3)	(4)
Total fertility rate (γ)	0.212	1.118**	0.595*	1.275**
	(0.194)	(0.098)	(0.179)	(0.234)
%Catholics × TFR (δ)	1.436**	0.174	0.681*	-0.264
	(0.271)	(0.202)	(0.257)	(0.403)
%Catholics × t (β)	1.534**	1.245**	0.951**	0.747*
	(0.175)	(0.173)	(0.155)	(0.298)
attendance (θ)	0.814†	-1.535**	0.452	-1.508**
	(0.426)	(0.223)	(0.254)	(0.371)
attendance × %Catholics (η)	1.642†	2.566†	1.407†	2.695†
	(0.874)	(0.821)	(0.633)	(0.844)
Spouse's income (in thousands)	-	0.028**	-	0.033**
		(0.002)		(0.002)
t × Spouse's income	-	-	-	-0.015**
				(0.003)
Spouse in labor force?	-	-0.152†	-	-0.222†
		(0.078)		(0.101)
t × Spouse in LF	-	-	-	0.699
				(0.545)
Spouse is foreign-born?	-	-0.203*	-	-0.197*
		(0.064)		(0.072)
t × Spouse is foreign-born	-	-	-	0.122
				(0.205)
Spouse has college education?	-	-0.159**	-	-0.123†
		(0.044)		(0.064)
t × spouse is college educated	-	-	-	-0.037
				(0.105)
R^2	0.426	0.589	0.464	0.636
Number of observations	3700	3700	3700	3700
Time-varying coefficients on Z?	No	No	Yes	Yes

Notes: Data on second-generation immigrants from ten source countries from the 1970 Census and 2000-2006 March CPS. TFR from the *Demographic Yearbook* (various years), percentage of Catholics from *Annuarium Statisticum Ecclesiae* (various years), and attendance from Iannaccone (2003). The dependent variable is the number of children currently at home under age 18. All regressions include individual-level characteristics and both country and time fixed effects. Robust SEs, clustered by country of origin, in parentheses. Symbols denote coefficients significantly different from zero at 10%(†), 5%(*), and 1%(**).

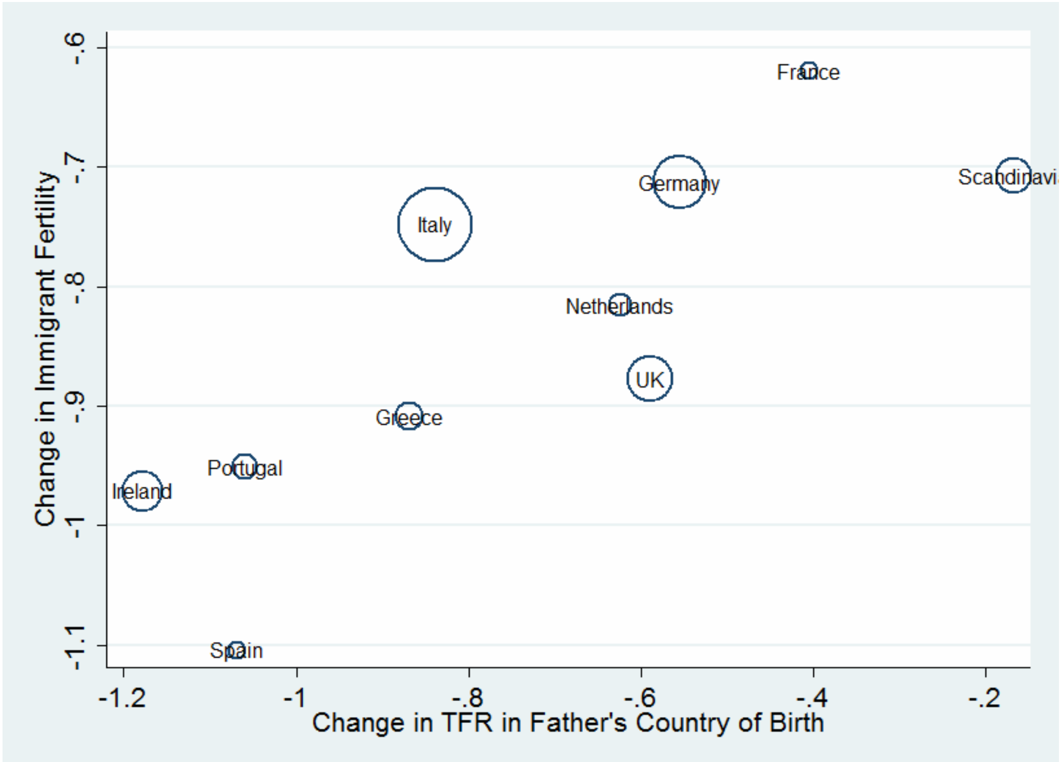

Figure 1.1: Second-Generation vs. Europeans, First Differences

Notes: The x-axis is the change in TFR between 1970 and 2000 in Europe. The y-axis is the change in fertility among second-generation women over that same period, adjusted by a quadratic in age. The size of each marker corresponds to the number of observations used in regression.

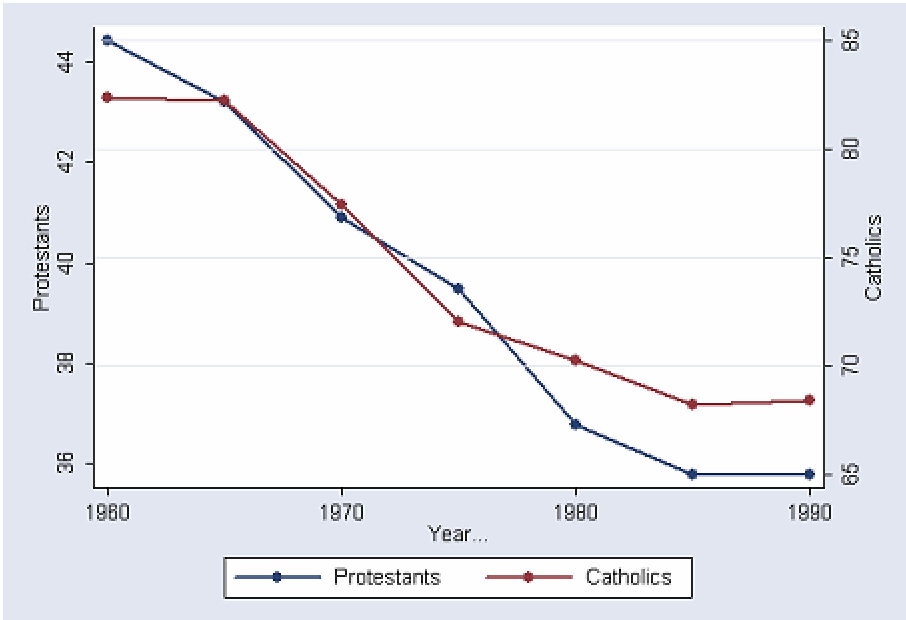

Figure 1.2: European Church Attendance

Notes: The Catholic countries in this figure are Ireland, Italy, Portugal, and Spain. The Protestant countries are Denmark, Norway, Sweden, and the UK. Church attendance data from Iannaccone (2003).

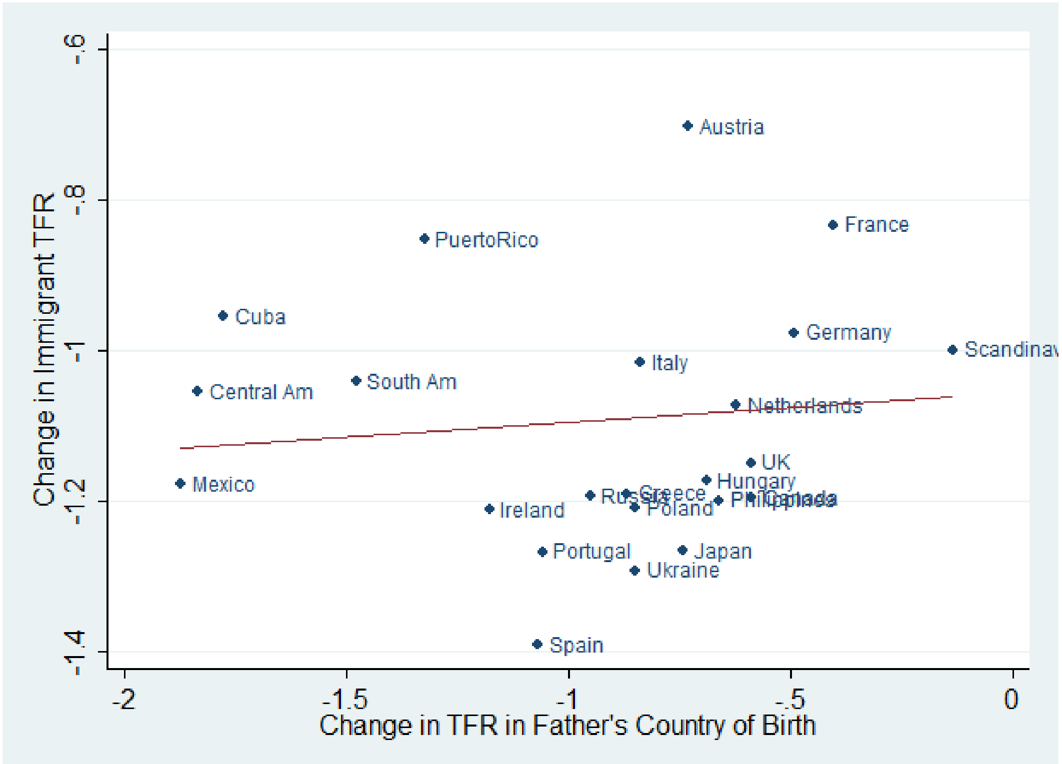

Figure 1.3: First Differences, All Countries

Notes: Total fertility rates are from the *Demographic Yearbook*. Immigrant fertility is the predicted number of children born to second generation women by age 35, predicted using either the 1970 Census (for 1970 data points) or from pooled 2000-20006 March CPS (for 2000). The OLS regression line, weighted by the number of observations in the census, has slope = 0.040 (se = 0.046).

1.8 Appendix

Adjusting Country-Level TFR

Fertility is generally interpreted as the expected number of children born to a woman at the end of her childbearing years. The most frequently used measure of fertility, total fertility rate, is interpreted in this manner though not actually constructed as such.[39] TFR is actually calculated as the sum of current age-specific birth rates (the annual number of births per woman in a given age bin), which may (or may not) be accurate prediction of completed fertility by the time a woman is finished having children. Since the explanatory variable of interest is total fertility rate, the ideal left-hand side would be constructed analogously as the sum of age-specific birth rates. The data requirements for this, namely birth counts for women of all ages for each parental country-of-birth, make it infeasible.

Rather than a technical analog, an alternative fertility measure for immigrants could have the same interpretation, if not the same construction, as TFR. Completed fertility, the total number of children ever born to a woman who has finished her childbearing years, fits this criterion but is not unavailable in the CPS samples. The only other fertility question that appears in both types of census data is the number of children currently at home under age 18. This “children-so-far” measure underreports true fertility for women with adult children, so the estimation sample is limited to women up to age 35. Since this fertility measure amounts to truncating total lifetime

[39] Cohort fertility rate is one such measure of completed family size, but can only be calculated for cohorts that are already past their childbearing years.

fertility at age 35 rather than 45, country-specific TFRs are also calculated using only age-specific fertility rates up to 35.

To account for women being at different points in their lifecycles, all regressions have a quadratic in age. This is intended to extrapolate the "children-so-far" measure of fertility to age 35, which can then be accurately compared to TFR.

The main findings of this paper, however, are not solely due to the constructed fertility measure. Table 1.7 presents results using the full sum of age-specific fertility rates rather than only those up to age 35. In Column (1), the coefficient on TFR has shrunk to 0.186 from 0.296 but still remains significant. All three coefficients reported in Column (3) are smaller than their Table 1.3 counterpart, and the Catholic-TFR interaction is no longer statistically significant, though it is still positive. However, this coefficient is positive and significant once church attendance is included. Attendance itself does not differentially matter for Catholics.

Differential Migration and Changes in Immigration Policy

There is a potential selection issue in comparing 1970 and 2000 immigrant cohorts because the parents of the two cohorts entered the US under very different immigration regimes. Between 1924 and 1965, entry into the US was regulated by national origins quotas, which were intended to preserve the ethnic composition of the population by admitting only those whose language, traditions, and political systems were similar to Americans (Eckerson, 1966). These quotas were repealed in 1656 in favor of the current family reunification system. This shift in immigration policy is problematic because the quotas specifically favored those from Northern Europe and sharply limited the number of entrants from other parts of the world. If the characteristics of Southern Europeans who managed to enter during the quota period were systematically different from those who entered after 1965, then this change in cohort composition could be the true factor driving the post-1970s fertility decline.

Table 1.8 addresses this possibility by listing average characteristics for immigrant cohorts from one country sharply affected by quotas (Italy) and two that were not (the UK and Germany).[40] The final column implies that immigrant cohorts do observably differ by entry year, especially in terms of education. Nearly 17% of men who came to the US in the 1940s had college degrees, but over 35% did for the more recent cohort. The fraction of college educated women has also doubled from 7% to over 18% by the 1970 cohort. While immigrant characteristics do appear to differ between those when entered the US in the 1940s and those who entered in the 1970s, this does not appear to be country-specific. In particular, the post-1965 entrants appear to be almost twice as likely to have college degrees, but this is true of all three countries, not just Italy. The lack of country-specific shifts in immigrant demographics suggests that the 1965 changes in immigration law does not seem to be observably selecting higher quality immigrants from quota-constrained Italy compared to unconstrained UK and Germany.

[40] Italy was limited to 4,000 entrants annually. The UK and Germany were allowed 34,000 and 57,000 respectively.

Alternate Outcomes: Labor Force Participation and Marital Status

European fertility is the obvious proxy for fertility outcomes among immigrants but is not the only one. Table 1.12 includes European LFP and rates of singlehood as potential measures of shared norms, though it turns out that TFR remains the most relevant.[41] In the first three columns, high European labor force participation implies lower immigrant fertility, but TFR still remains significant. The coefficient on European labor force participation is more than twice the size of the one on TFR, but it is not as precisely estimated. Increasing rates of singlehood predict lower immigrant fertility, but unlike LFP, its coefficient is never statistically different from zero. Once religion is included, however, the coefficient on labor force participation becomes positive but no longer significant. The final three columns of this table display similar results as Table 1.3: large positive coefficients on both the Catholic-TFR interaction and church attendance.[42] Taken together, Table 1.4 and Table 1.12 imply that the earlier finding of trans-Atlantic fertility correlation among Catholics cannot be explained by underlying shifts in work or marriage norms among second-generation immigrants.

Since a woman's decision to work or attitudes toward marriage may be cultural or rooted in religious teachings, the same analysis can be run using labor force participation or marital status as outcome variables. Table 1.13 does exactly this, and

[41] To be consistent, these additional variables could also be interacted with the fraction of Catholics to distinguish religion-specific norms regarding marriage and/or women working. However, there are too few countries in the sample to estimate this regression.
[42] An argument can be made that labor force participation and marriage norms should be treated like fertility and interacted with the percentage of Catholics, but there are not enough degrees of freedom to estimate the fully interacted regression.

neither LFP nor marriage rates in Europe appear to be correlated with those for immigrants.

The initial columns use an indicator for individual labor force participation as the dependent variable while the final two use marital status. The first difference to note between this table and the fertility one is that the simple specifications here do a very poor job of predicting work or marriage. The R^2 for the fertility regressions were around 0.46, the included regressors explain around 6% of the observed variation in outcomes here.

For labor force participation, the estimated norm in Column (1) is actually negative, though not different from zero. However, the interaction between the fraction of Catholics and source country LFP is positive and significant. Using the same interpretation as the fertility regressions, this coefficient implies that large gains in female labor force participation in Europe were also evident in employment and differential from the LFP trend of non-Catholic nations. While this looks suggestive of norms, the two slopes added together are not significantly different from zero (p-value = 0.192); gains in employment in Catholic nations are not statistically mirrored in the actions of second-generation women. It should be noted that Southern Europe is demographically puzzling not only for its large declines in fertility, but also for continued low rates of female labor force participation. The data do not suggest any sort of cross-country correlation regarding marriage or rates of singlehood.

Married women and husband's characteristics

As childbearing is traditionally a decision made by couples, most studies of fertility examine just married women. However, since the 1970s, more and more women are delaying marriage or choosing not to get married at all. Whereas 75% of second-generation women were married in the 1970 sample, only half are married in the pooled CPS.

Restricting the sample to only married women provides room to expand the regressors in Z_i to include the characteristics of the husband, but amounts to inducing positive selection bias into the estimates since unmarried women tend to have no children and marriage norms almost assuredly vary between Catholics and non-Catholics. Table 1.14 contains estimates of Eq. (7) estimated using only the sample of married women and augmented with husband's characteristics (income in thousands, his labor force participation, whether he was foreign-born, and whether he has a college degree). Column (3) restricts the estimates of Eq. (7) to only married women but is otherwise comparable to the main result in Column (4) of Table 1.3. The sample selection results in overestimates on the coefficients of interest; a one child increase in TFR predicts an additional 0.9 children among immigrants from wholly Catholic countries in Table 1.3, but an additional 1.3 (= 0.595 + 0.681) children for married women. Similarly, the 17 percentage point decline in church attendance implies a 0.2 (= 1.156*0.17) fall in Catholic immigrant fertility when estimated using the entire sample but a 0.3 drop among married women. The magnitude of these coefficients is somewhat smaller in Column (4) when spousal characteristics are included. Instead of

1.3, the coefficient on TFR for Catholics is 1.011; the slope on attendance has dropped to 1.187 from 1.859.

1.9 References

Adsera, Alicia. (2004a). Changing fertility rates in developed countries. The impact of labor market institutions. *Journal of Population Economics*, 17(1), 17-43.

Adsera, Alicia. (2004b). Marital Fertility and Religion: Recent Changes in Spain. IZA Discussion Papers 1399, Institute for the Study of Labor (IZA).

Ahn, Namkee, & Pedro Mira. (2002). A note on the Changing relationship between fertility and female employment rates in developed countries. *Journal of Population Economics*, 15(4), 667-682.

Akerlof, George A. & Rachel E. Kranton. (2000). Economics and Identity. *Quarterly Journal of Economics*, 115(3), 715-753.

Albanesi, Stephania & Claudia Olivetti. (2010). Maternal Health and the Baby Boom. NBER Working Paper No. 16146, National Bureau of Economic Research, Inc. July.

Alberigo, Guiseppe. (2006). *A Brief History of Vatican II*. Translated by Matthew Sherry. New York: Orbis Books.

Alesina, Alberto, Edward Glaeser, & Bruce Sacerdote (2001). Why Doesn't the US Have a European-Style Welfare System? NBER Working Papers 8524, National Bureau of Economic Research, Inc.

Almond, Douglas, Lena Edlund, & Kevin Milligan. (2009). Son Preference and the Persistence of Culture: Evidence from Asian Immigrants to Canada. NBER Working Papers 15391, National Bureau of Economic Research, Inc. Oct.

Antecol, Heather. (2000). An examination of cross-country differences in the gender gap in labor force participation rates. *Labour Economics*, 7(4), 409 – 426.

Becker, Gary S. (1981). *A Treatise on the Family*, Cambridge MA: Harvard University Press.

Berman, Eli, Laurence R. Iannacone, & Guiseppe Ragusa. (2006). From Empty Pews to Empty Cradles: Fertility Decline Among European Catholics. June. Working Paper.

Bertola, Guiseppe, Juan Francisco Jimeno, Ramon Marimon, & Christopher Pissarides. (2001). EU Welfare Systems and Labor Markets: Diverse in the Past, Integrated in the Future? In Bertola, Boeri and Nicoletti (ed.), *Welfare and Employment in a United Europe: A Study for the Fondazione Rdolofo Debenedetti*, Cambridge, MA: MIT Press, 23-122.

Bettio, Francesca & Paola Villa. (1998). A Mediterranean perspective on the breakdown of the relationship between participation and fertility. *Cambridge Journal of Economics*, 22(2), 137–171.

Bisin, Alberto & Thierry Verdier. (2010). The economics of cultural transmission and socialization. May. Working Paper.

Blau, Francine D. (1992). The Fertility of Immigrant Women: Evidence from High-Fertility Source Countries. In *Immigration and the Workforce: Economic Consequences for the United States and Source Areas*, National Bureau of Economic Research, Inc., 93-134.

Blau, Francine D., Lawrence M. Kahn, Albert Yung-Hsu Liu, & Kerry L. Papps. (2008). The Transmission of Women's Fertility, Human Capital and Work Orientation Across Immigrant Generations. NBER Working Papers 14388, National Bureau of Economic Research, Inc.

Bongaarts, John. (1999). Fertility Decline in the Developed World: Where will it end? *American Economic Review, Papers and Proceedings*, 89(2), 256-260.

Borjas, George J. (1993). The Intergenerational Mobility of Immigrants. *Journal of Labor Economics*, 11(1), 113-35.

Card, David, John DiNardo & Eugena Estes. (2000). The More Things Change: Immigrants and the Children of Immigrants in the 1940s, the 1970s, and the 1990s. In *Issues in the Economics of Immigration*, National Bureau of Economic Research, Inc., 227-270. .

Carroll, Christopher D., Byung-Kun Rhee, and Changyong Rhee. (1999). Does Cultural Origin Affect Saving Behavior? Evidence From Immigrants. *Economic Development and Cultural Change*, 48(1), 33 – 50.

Chiswick, Barry R. (1978). The Effect of Americanization on the Earnings of Foreign-born Men. *Journal of Political Economy* 86(5), 897-921.

Clark, Andrew E. & Orsolya Lelkes (2009). Let us pray: religious interactions in life satisfaction. PSE Working Paper 2009-01.

de Laat, Joost & Aludena Sevilla Sanz. (2007). Working Women, Men's Home Time and Lowest Low Fertility. Economics Series Working Papers 308, University of Oxford, Department of Economics.

del Boca, Daniela. (2002). The effect of child care and part time opportunities on participation and fertility decisions in Italy. *Journal of Population Economics*, 15(3), 549–573.

Donald, Stephen G. and Kevin Lang. (2007). Inference with Difference-In-Differences and Other Panel Data. *Review of Economics and Statistics*, 89(2), 221–233.

Durkheim, Emile. [1915] (1965). *The Elementary Forms of the Religious Life*. Translated by J. W. Swain. New York, NY: Free Press.

Eckerson, Helen F. (1966). Immigration and National Origins. *Annals of the American Academy of Political and Social Science*, 367(1), 4-14.

Fernandez, Raquel. (2007). Women, Work, and Culture. *Journal of the European Economic Association*, 5(2-3), 305-332.

Fernandez, Raquel & Alessandra Fogli. (2009). Culture: An Empirical Investigation of Beliefs, Work, and Fertility. American Economic Journal: Macroeconomics, 1(1), 146-77.

Fialka, John J. (2003). *Sisters: Catholic Nuns and the Making of America*. New York: St. Martin's Press.

Fortin, Nicole M. (2005). Gender Role Attitudes and the Labour-market Outcomes of Women across OECD Countries. *Oxford Review of Economic Policy*, 21(3), 416–438.

Frejka, Tomas & Charles F. Westoff. (2008). Religion, Religiousness and Fertility in the US and in Europe. *European Journal of Population*, 24(1), 5–31.

Giuliano, Paola. (2007). Living Arrangements in Western Europe: Does Cultural Origin Matter? *Journal of the European Economic Association*, 5(5), 927–952.

Greenwood, Jeremey, Ananth Seshadri, Guillaume Vandenbroucke. (2005). The Baby Boom and Baby Bust. *American Economic Review*, 95(1), 183-207.

Guinnane, Timothy W., Carolyn M. Moehling, & Cormac O Grada. (2006). The fertility of the Irish in the United States in 1910. *Explorations of Economic History*, 43(3), 465-485.

Hagopian, Frances. (2009). *Religious Pluraism, Democracy, and the Catholic Church in Latin America*. Notre Dame, IN: University of Notre Dame Press.

Hout, Michael & Andrew M. Greeley. (1987). The Center Doesn't Hold: Church Attendance in the United States, 1940-1984. *American Sociological Review*, 52(3), 325-345.

Iannaccone, Laurence R. (2003). Looking Backward: A Cross-National Study of Religious Trends. June. Working Paper.

International Labour Office. *LABORSTA Internet* [Data file]. Available from http://laborsta.ilo.org.

King, Miriam, Steven Ruggles, Trent Alexander, Donna Leicach, & Matthew Sobek. (2009). *Integrated Public Use Microdata Series, Current Population Survey: Version 2.0*. [Machine-readable database]. Minneapolis, MN: Minnesota Population Center [producer and distributor]. Available from http://cps.ipums.org/cps/.

Kohler, Hans-Peter, Francesco C. Billari, & Jose Antonio Ortega. (2006). Low Fertility in Europe: Causes, Implications and Policy. in F. R. Harris (ed.), *The Baby Bust: Who will do the Work? Who Will Pay the Taxes?* Lanham, MD: Rowman & Littlefield Publishers, 48-109.

Lalonde, Robert J. & Robert H. Topel. (1997). Economic impact of international migration and the economic performance of migrants. In M. R. Rosenzweig & O. Stark (ed.), *Handbook of Population and Family Economics*, vol. 1, ch. 14, 799-850.

Lehrer, Evelyn L. (2009). Religion, Human Capital Investments and the Family in the United States. In R. M. McCleary (ed.), *The Oxford Handbook of the Economics of Religion*.

Lino, Mark. (2010). *Expenditures on Children by Families, 2009*. U.S. Department of Agriculture, Center for Nutrition Policy and Promotion. Miscellaneous Publication No. 1528-2009.

Moulton, Brent R. (1990). An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units. *Review of Economics and Statistics*, 72(2), 334–338.

Norris, Pippa & Ronald Inglehart. (2004). *Sacred and Secular: Religion and Politics Worldwide*. Cambridge University Press.

Pop-Eleches, Cristian. (2010). The Supply of Birth Control Methods, Education and Fertility: Evidence from Romania. *Journal of Human Resources*, 45(4), 971-997.

Reimers, Cordelia W. (1985). Cultural Differences in Labor Force Participation Among Married Women. *American Economic Review*, 75(2), 251-255.

Reher, David Sven. (1998). Family Ties in Western Europe: Persistent Contrasts. *Population and Development Review*, 24(2), 203-234.

Ruggles, Steven et al. (2008). *Integrated Public Use Microdata Series: Version 4.0* [Machine-readable database]. Minneapolis, MN: Minnesota Population Center [producer and distributor]. Available from http://usa.ipums.org/usa/.

Schultz, T. Paul. (1981). *Economics of Population*, Reading MA: Addison-Wesley.

Secretaria Status: Rationarium Generale Ecclesiae. (Various years). *Annuarium Statisticum Ecclesiae*, Vatican City: Tipografia Poliglotta Vaticana.

Stigler, George J. & Gary S. Becker. (1977). De Gustibus Non Est Disputandum. *American Economic Review*, 67(2), 76-90.

Stolzenberg, Ross M., Mary Blair-Loy, & Linda J. Waite. (1995). Religious Participation Over the Early Life Course: Age and Family Life Cycle Effects on Church Membership. *American Sociological Review*, 60(1), 84-103.

United Nations. (1976, 2002). *Demographic Yearbook*. New York: United Nations Publishing Service.

Warner, Carolyn M. (2000). *Confessions of an Interest Group: The Catholic Church and Political Parties in Europe*, Princeton, N.J.: Princeton University Press.

Weber, Max. (1946). The Protestant Sects and the Spirit of Capitalism. In H.H. Gerth & C. Wright Mills (eds.), *Max Weber: Essays in Sociology*. New York: Oxford University Press.

Westoff, Charles F. & Elise F. Jones. (1979). The End of "Catholic" Fertility. *Demography*, 16(2), 209-217.

Wilde, Melissa J. (2004). How Culure Mattered at Vatican II: Collegiality Trumps Authority in the Council's Social Movement Organizations. *American Sociological Review*, 69(4), 576-602.

World Bank. *HNPStats* [Data file]. Available from http://go.worldbank.org/N2N84RDV00.

Chapter 2: Does Factor-Biased Technological Change Stifle International Convergence? Evidence from Manufacturing

Abstract

Factor-biased technological change implies divergent productivity growth across countries with different factor intensities. We estimate that factor bias within manufacturing industries across both developing and developed countries in the 1980s and 1990s. Technological change is strongly biased against less-skilled workers and toward both skilled workers and capital in both decades. Labor-saving technological change is especially strong among a set of newly industrializing, low-income countries in the 1990s, possibly due to technology transfer from high- and middle-income countries. Our results are consistent with the empirical literature on skill-biased technological change, and can explain why "conditional convergence" of GDP/capita across countries is slow.

2.1 Introduction

Why do some countries remain so much poorer than others? The two basic approaches to income convergence yield quite hopeful conclusions. The factor accumulation approach (Solow, 1956) predicts that low productivity is the result of low ratios of skill and capital to labor. In the presence of diminishing returns, countries with low skill and capital intensity have highly productive skill and capital. This implies relatively rapid accumulation of capital and skill per worker in poor countries, resulting in eventual convergence in factor intensities and labor productivity. A second approach argues that low labor productivity is the result of using inferior technologies.[43] As replicating technology must be less costly than inventing new technologies, technology use should converge, leading to eventual convergence of total factor productivity.

The evidence, on the other hand, is not hopeful at all. Since the early 1960s, growth rates of GDP per capita have not been higher in countries with low GDP per capita (Barro, 1991). Most studies find convergence only after conditioning on available measures of international differences in institutions and preferences that explain the slow accumulation of skill and capital in poorer countries (Barro & Sala-I-

[43] Solow (1957) measured total factor productivity growth in the US. Eaton and Kortum (1996, 1999) offered evidence of international technology transfer using R&D and patent statistics. Technology transfer models fall into two broad categories. The "appropriate" technology model (Schumacher, 1973; Basu & Weil, 1998; Acemoglu & Zilibotti, 2001; Caselli & Coleman, 2006) posits that new technologies are not absorbed immediately in developing countries because of a lack of human or physical capital, differences in production technologies in use, or differences in factor prices. In contrast, the conventional assumption in growth theory is of pervasive technology in use everywhere. A weaker assumption is that technologies differ, but recent innovations are so efficient that they are adopted across a wide range of industries, factor price combinations and local technological capabilities. That concept is related to research on "General Purpose Technologies" (Bresnahan & Trajchtenberg, 1995; Helpman, 1998), such as electrification and information technology, which increase productivity in a wide range of industries.

Martin, 1995; Mankiw, Romer & Weil, 1992). Even then, this “conditional convergence” is quite slow.

This paper suggests an alternative explanation for slow productivity convergence: factor-biased technological change. If technological innovations favor the skilled over the unskilled, then industries with more skilled workers should have faster total factor productivity (TFP) growth (Klenow, 1998; Kahn & Lim, 1998; Ciccone & Papaioannou, 2009). Similarly, it would not be surprising if countries with a high proportion of less-skilled workers had slower growth rates of income per capita. Substantial evidence now exists demonstrating that technological change has favored skilled (i.e. more educated) workers over unskilled (less educated) workers at least for the past few decades in many other parts of the world.[44] Table 2.1 provides a sampling of that evidence, showing the declining wagebill shares of production workers in the manufacturing industries of both high and middle income countries. Across countries, the importance of high-skilled labor in developed countries has only continued to increase through the 1990s, especially in the service sector (Jorgenson & Timmer, 2011).

This paper estimates the factor bias of technological change in manufacturing and applies the estimates to the puzzle of slow productivity convergence. Our data are

[44] For evidence of recent skill-biased technological change in the US, see Bound and Johnson (1992), Katz and Murphy (1992), Lawrence and Slaughter (1993), or Berman, Bound, and Griliches (1994). Historical evidence is offered by Goldin and Katz (1996, 1998), as far back as the beginning of the century. Evidence from other OECD countries is available in Freeman (1988), Freeman and Katz (1994), Katz and Revenga (1989), Katz, Loveman and Blanchflower (1995), Davis (1992), and Berman, Bound, and Machin (1998). Skilled labor experienced increased relative wages in several developing countries despite widespread trade liberalization in the 1980s (Feliciano, 2001; Hanson & Harrison, 1995; Robbins, 1995; Berman, Bound, & Machin 1998; Berman & Machin 2000).

an international panel of manufacturing industries through the 1980s, 1990s, and early 2000s. Factor-bias parameters are estimated twice: once using a production function and again using a TFP specification. Both approaches yield consistent, strong evidence that technological change over this period favored capital and was biased against labor. In the 1980s, labor-saving can be further narrowed down to bias against unskilled labor specifically; technology appears to be at least weakly skill-biased. In the 1990s we find evidence of labor-saving technological change even in the set of low income rapidly developing countries. Productivity growth was much higher among middle-income countries in the 1980s and among low-income ones in the 1990s, consistent with the "technology transfer hypothesis" (Berman & Machin, 2000; Griffith, Redding, & Van Reenen, 2004).

Factor bias estimates for the 1980s indicate that an industry or a country with twice the ratio of skills and capital to less-skilled labor enjoys a 1.4%-1.8% faster annual rate of TFP growth. Results for the 1990s suggest slightly higher overall rates of labor-saving (1.7%-2.4%), but this obscures a very strong factor-bias among the set of low income countries (2.3%-4.9%). These estimates are net of any country (and industry) effects, making this finding orthogonal to that of "conditional" convergence across countries.

The next section of this paper provides background on the lack of international productivity convergence, demonstrating that non-convergence is evident across the countries in our sample, even if we limit the analysis to only the manufacturing sector. Section 2.3 develops a production function framework for estimation. Section 2.4

describes the data, including potential estimation problems. Section 2.5 presents estimates from 1980s General Industrial Statistics data and discusses their plausibility in the context of a world with accelerated technology-transfer. The sixth section provides productivity estimates for the 1990s using Industry Statistics (INDSTAT) data. Section 2.7 examines the implications of estimated factor-bias for productivity convergence. Section 2.8 concludes.

2.2 TFP Growth and Factor Accumulation in Manufacturing

Figure 2.1 examines the manufacturing sample used here in the context of global non-convergence. The top panel reproduces the standard finding that income per capita growth rates between 1960 and 2000 are uncorrelated with income levels (Barro, 1991).[45] Poorer countries do not appear to grow faster than wealthier ones. This pattern of international non-convergence is quite stable, as demonstrated by the two lower panels. Both the 1980s (left panel) and the 1990s (right panel) display the same pattern of non-convergence: a triangle pointing right.

The sample of manufacturing data used in this paper consists of the labeled countries in Panel B of Figure 2.1.[46] For the 1980s, the relationship between growth rates and GDP levels among the included countries roughly mimics the pattern seen in the larger sample. The cross-country variance of growth rates declines with income and the average growth rate shows a slight reduction as income increases. For the

[45] Data are drawn from the Penn World Table, version 6.2 (Heston, Summers, & Aten, 2006).
[46] Only countries with usable measures of capital at the beginning and end of each of decade are included in the estimation sample. This results in a disproportionate number of high income countries.

1990s manufacturing sample, both the variance and mean of growth rates appear to be constant across the income range.

National growth rates are quite persistent; the correlation between the 1960-90 growth rate and the 1980-90 growth rate is 0.88 (α=0.00) for the nineteen labeled countries in the 1980s sample. For the 26 countries in the right panel, the correlation between forty-year and ten-year growth rates 0.56 (α=0.003).

Non-convergence in GDP growth rates is evident not just within decades, but also in just the manufacturing sector. The two panels of Figure 2.2, which plots growth rates of manufacturing value added instead of GDP against income, show the same triangle. Growth rates do not decline with income and have higher variance at lower income levels. Countries with high growth rates in GDP per capita generally have high growth rates in manufacturing value added per worker. The correlation between the thirty-year GDP per capita growth rate and the value added per worker rate in the 1980s is 0.22 (α=0.39), but rises to 0.42 (α=0.09) without Chile. In the 1990s, the correlation between 1960-2000 GDP growth rate and manufacturing value added per worker is 0.44 (α=0.03). Overall, manufacturing value added per worker in this figure mimics the pattern of non-convergence in international GDP per capita. This is consistent with the conventional view that successful NICS, such as Korea, have grown by rapidly expanding manufacturing.

Is it factor accumulation or TFP that is not converging in manufacturing? Within our sample, growth rates can be decomposed into TFP growth on the one hand and factor accumulation (skill and capital intensity) on the other. Suppose the production function is $Y = A \cdot F(L, S, K)$ where the inputs are unskilled labor, skilled labor and capital respectively. Assuming constant returns and competitive markets, a standard definition of TFP growth is:

$$\Delta TFP = \Delta y - \left[\psi_L \Delta l + \psi_S \Delta s + (1 - \psi_L - \psi_S)\Delta k\right],$$

where lowercase letters denote logarithms and ψ's indicate factor shares. Letting $E = S + L$ be total employment, the growth rate of value added per worker is:

$$\begin{aligned}\Delta y - \Delta e = {} & \psi_L(\Delta l - \Delta e) + \psi_S(\Delta s - \Delta e) \\ & + (1 - \psi_L - \psi_S)(\Delta k - \Delta e) + \Delta TFP\end{aligned}$$

$$= \text{“factor accumulation”} + \Delta TFP.$$

Figure 2.3 illustrates this decomposition for the 1980s data, with the left panel plotting factor accumulation against GDP per capita, and the right panel plotting TFP growth against GDP per capita. The left panel makes it clear that little Solow convergence occurred in the form of capital or skill accumulation. This pattern is analogous to the results of Mankiw, Romer and Weil (1992), who found convergence only once they conditioned on accumulation rates. Manufacturing TFP growth is not contributing to convergence either; it shows the same triangular pattern seen in Figure 2.2. The right panel clearly illustrates that most (86%) of the variance in the growth of manufacturing value added per worker is in TFP growth. It is worth stressing that since TFP growth rates are calculated as a residual, improved measurement might reallocate growth from TFP to factor accumulation (Griliches & Jorgenson, 1967;

Young, 1995).[47] On the other hand, these TFP growth rates for manufacturing are not simply measurement error. The correlation between the 1960-1990 GDP/capita growth rate and the manufacturing TFP growth rate in the 1980s is 0.17 (α=0.50), but rises to 0.42 (α=0.095) without Chile. This correlation is remarkably tight, considering the difference in data sources and the fact that the two growth rates have only ten of thirty years in common. It is safe to conclude that a large component of growth in manufacturing output per worker is TFP growth. Furthermore, factor accumulation would have to be understated by an order of magnitude, and disproportionately so, in the lower income economies for the two panels of Figure 2.3 to be reversed.

The accumulation-TFP decomposition is somewhat more ambiguous in the 1990s. The left panel of Figure 2.4 shows that much of the variance in the growth rates of value added comes from factor accumulation rather than TFP growth. There is one major caveat: skill accumulation is not one of the components of factor accumulation. The 1990s data do not distinguish between skilled and unskilled workers, so "factor accumulation" is just growth in capital intensity (K/L).[48] Even with this limitation, the 1990s manufacturing panel still demonstrates a lack of convergence in factor accumulation and TFP growth within income groups.

To sum up, the manufacturing data reproduce the pattern of non-convergence evident in GDP per capita. For the 1980s, most of the non-convergence is in TFP growth rates, but both TFP growth and factor accumulation are quite varied in the

[47] Young (1995) addresses a debate as to whether the rapid growth of several East Asian economies is due to TFP growth or to factor accumulation. One of the messages of this paper is that the dichotomy is false, since factor bias translates current factor accumulation into future TFP growth.
[48] In addition, the employment data for the 1990s exhibit some severe employment reporting issues which further confounds the calculation of TFP growth.

1990s. If replication is less costly than invention, why is TFP growth not contributing to convergence in value added per worker?

2.3 Factor-Biased Technological Change in Production: A Framework for Estimation

This section develops a framework to explain how factor-biased technological change can yield divergent TFP growth rates. This framework also generates estimating equations, allowing the data to report the magnitude of TFP divergence due to factor bias.

A Cobb-Douglas production function with exponents that change over time is a flexible way to represent factor-biased technological change:

$$Y = e^{\alpha+\rho t} L^{\beta_L(t)} S^{\beta_S(t)} K^{\beta_K(t)}, \tag{8}$$

where Y is product, L is unskilled labor, S is skilled labor and K is capital. Time is indexed by t. Using lowercase letters to indicate logarithms, Eq. (8) can be rewritten as:

$$y = \alpha + \rho t + \beta_L(t) \cdot l + \beta_S(t) \cdot s + \beta_K(t) \cdot k \,. \tag{9}$$

Output elasticities are given by:

$$\frac{\partial y}{\partial f} = \beta_f(t) \text{ for } f \in \{L, S, K\}.$$

The rate at which $\beta_f(t)$ changes the bias of technological change towards factor f is:

$$\frac{\partial^2 y}{\partial f \partial t} = \frac{\partial \beta_f(t)}{\partial t} = \beta_f^{'}(t) .$$

Constant returns to scale (CRS) require that the exponents $\beta_f(t)$ sum to one.[49] A weaker assumption that will prove useful is that returns to scale, constant or otherwise, remain unchanged by technological progress (i.e. the $\beta_f^{'}$ terms sum to zero). The implications of "unchanging returns to scale" (URS) will become important to the discussion below.

A working definition of relative factor-bias links this framework to the literature:

Technological change is *relatively skill biased* if $\dfrac{\beta_S^{'}(t)}{\beta_S(t)} > \dfrac{\beta_L^{'}(t)}{\beta_L(t)}$. (10)

Relative skill bias implies that the output elasticity of skilled labor increases at a faster rate than that of unskilled labor. To justify this usage, consider the implications of relative skill-bias. Assuming perfectly competitive labor markets and holding relative wages constant, the relative demand for skilled workers is:

$$\frac{w_S}{w_L} = \frac{MP_S}{MP_L} = \frac{\beta_S(t)(Y/S)}{\beta_L(t)(Y/L)} = \frac{\beta_S(t)}{\beta_L(t)}\frac{L}{S}$$

$$\Leftrightarrow \frac{S}{L} = \frac{\beta_S(t)}{\beta_L(t)}\frac{w_L}{w_S}$$

Eq. (10) implies that the relative demand for skilled workers is increasing over time:

[49] If inputs are to be forever useful in production and subject to diminishing marginal returns, then the standard restriction $0 < \beta_f(t) < 1$ must also be imposed for all factors *f* and time *t*.

$$
\begin{aligned}
\frac{\partial}{\partial t}\left(\frac{S}{L}\right) &= \frac{w_L}{w_S}\cdot\frac{\partial}{\partial t}\left[\frac{\beta_S(t)}{\beta_L(t)}\right] \\
&= \frac{w_L}{w_S}\left[\frac{\dot{\beta}_S}{\beta_L} - \frac{\beta_S}{\beta_L}\frac{\dot{\beta}_L}{\beta_L}\right] \\
&> \frac{w_L}{w_S}\left[\frac{\dot{\beta}_S}{\beta_L} - \frac{\beta_S}{\beta_L}\frac{\dot{\beta}_S}{\beta_S}\right] = 0.
\end{aligned}
$$

Conversely, holding the ratio of inputs fixed, relative skill-bias implies that the relative wage of skilled workers increases. For a Cobb-Douglas production function, it also implies that the wagebill share of skilled workers increases. These three implications have been treated as alternative symptoms of skill-biased technological change in the literature (Bound & Johnson, 1992; Katz & Murphy, 1992).

This framework also allows estimation of the absolute, as opposed to relative, bias of technological change. Define technological change as:

absolutely f-biased if $\beta_f^{'} > 0$,

absolutely f-saving if $\beta_f^{'} < 0$.

That is, technological change is absolutely *f*-biased if the marginal product of factor *f* increases over time (beyond the neutral increase ρ), holding inputs constant.

In the two factor Cobb-Douglas model without capital, unchanging returns to scale imply that $\beta_S^{'} = -\beta_L^{'}$, so absolute and relative skill bias are equivalent, and skill-biased technological change is equivalent to labor-saving technological change. The three factor model, even with unchanging returns to scale, is more flexible. For instance, technology could be absolutely biased against both *s* and *l*, but relatively biased toward *s*. Assuming unchanging returns, absolute skill-bias and absolute capital-bias imply absolute labor-saving since $\beta_L^{'} = -\beta_S^{'} - \beta_K^{'}$.

Factor Bias and Productivity Growth

To study the effect of factor bias on productivity change, note that the *β'* terms also reflect the effect of factor quantities on changes in total factor productivity since:

$$\frac{dTFP}{dt} = \left.\frac{\partial y}{\partial t}\right|_{l,s,k} = \rho + \beta_L^{'}(t)\cdot l + \beta_S^{'}(t)\cdot s + \beta_K^{'}(t)\cdot k\,. \qquad (11)$$

Here the partial derivative of *y* with respect to time is a change in total factor productivity since inputs are held constant. The factor bias term is the cross-partial of output with respect to time and input *f*.[50] For example, if technological change is absolutely skill-biased, then TFP growth must be faster the greater the level of skilled labor in production. Eq. (11) suggests that one way of estimating factor bias terms is by regressing the TFP growth rate on the levels of inputs.

Figure 2.5 illustrates a relatively skill-biased technological change as the shift of an (unit) isoquant in *S*-*L* space, holding *K* constant. For a country or an industry at point B, the *S/L* ratio is given by the slope of the vector OB, and the productivity gain is given by the length of the segment BC -- the decrease in inputs required to produce one unit of output. This technological change is relatively skill-biased since at the original relative wage ratio, illustrated by the slope of the line tangent to the isoquant at B, the new isoquant requires a higher *S/L* ratio (at point D). In contrast, a country or

[50] This property is not specific to the Cobb-Douglas. It follows from the symmetry of cross-partial derivatives.

industry with a lower *S/L* ratio, say at point A, experiences no productivity gain. The size of the differential productivity gain between A and B is given by the factor bias coefficients $\beta'(t)$, which are estimated below.

Assuming unchanging returns to scale, β'_L has the following convenient interpretation: if one industry has twice the *K/L* ratio and twice the *S/L* ratio as another, the TFP growth rate of the former will be $-\beta'_L$ faster. Anticipating our results, β'_L will be negative, so the former will grow faster. The URS assumption is not strictly necessary for what follows but it is convenient and it allows for more precise estimates. In some cases, the data will insist that the factor bias terms sum to a negative number, implying that returns to scale decline over time. Yet the a replication argument implies that returns to scale should remain unchanging: if declining returns were true at the firm level, large firms could split into smaller pieces to increase productivity. If declining returns were true at the industry level, large industries could send production abroad to increase productivity. Either way, in equilibrium we should not observe declining returns to scale. The plausibility of URS will come up again in interpreting estimates, though the thrust of the evidence for factor bias will not require this assumption.

Estimating Equations and Measurement Issues

Estimation requires a functional form assumption for $\beta(t)$. Imposing a linear parameterization,[51]

$$\beta_f(t) = \beta_f + \gamma_f \cdot t \text{ so } \beta_f^{'}(t) = \gamma_f ,$$

onto Eq. (9) yields one way of estimating factor bias terms. Eq. (11) provides the second method. Under this simplification, unchanging returns to scale hold if the sum of factor-bias coefficients adds to zero: $\sum \gamma_f = 0$. Constant returns to scale hold if $\sum \beta_f = 1$.[52]

The data we use to estimate production function parameters are three-dimensional panels of manufacturing industries across a range of countries. Each decade has a separate dataset in which each industry-country observation is observed twice. The set of available countries and industries differs between the datasets so they cannot be linked together into a single panel.

Measurement issues complicate estimation for two reasons. First, inputs are measured inconsistently. The definitions of skilled and unskilled labor sometimes differ conceptually across countries. For instance, middle income countries are more likely to undersample small firms, which tend to have lower proportions of skilled

[51] This restriction should be thought of as a short term approximation. The linear functional form implies that if γ_f is nonzero, factor f will eventually have an output elasticity outside the [0,1] interval.
[52] This condition is not technically correct since time shows up in the sum of elasticities: $\Sigma\beta_f(t) = \Sigma\beta_f + t\Sigma\gamma_f$. However, it is sufficient for CRS in the first (t =0) period, which is how "constant returns" will be used in the remainder of this paper. Alternatively, this restriction also implies constant returns to scale if the factor-bias coefficients sum to zero (i.e., under URS).

workers, leading them to overestimate the proportion of skilled workers.[53] The quality of inputs may also differ across industries. More generally, Griliches and Jorgenson (1967) demonstrate that mismeasurement of input quality can lead to substantial mistakes in TFP accounting. Assume that capital, skilled labor, and unskilled labor are measured with a country-specific error of proportionality. In logarithms, the observed quantity is then the sum of the true quantity and a country-factor specific error:

$$f_{ict}^{m} = f_{ict} + u_c^f \text{ for } f \in \{L, S, K\}.$$

Besides inconsistent measurement of factor qualities, a second source of potential measurement error is in price comparisons across countries and industries. National price indexes from the Penn World Tables are not completely corrected for quality, which is likely to differ disproportionately across industries because of market power, particularly for nontraded goods. These fixed industry-country specific price differences are absorbed by an industry-country specific level effect, α_{ic}, which also absorbs fixed productivity differences, measurement error in output, and any industry-country specific measurement error in quantities. These measurement errors may be substantial considering that the data are collected from disparate sources without the intention of making them comparable. Including a country-period specific productivity level δ_{ct} and an industry-specific productivity trend in output growth ρ_i in (9) yields:

$$\begin{aligned} y_{ict} = {} & \alpha_{ic} + \delta_{ct} + \beta_L(l_{ict} + u_c^L) + \beta_S(s_{ict} + u_c^S) + \beta_K(k_{ict} + u_c^K) \\ & + \rho_i t + \left[\gamma_L(l_{ict} + u_c^L) + \gamma_S(s_{ict} + u_c^S) + \gamma_K(k_{ict} + u_c^K)\right] \cdot t + \varepsilon_{ict}. \end{aligned} \tag{12}$$

[53] The measured proportion of skilled workers in Japanese manufacturing jumped from 46 to 53 percent between the 1975 and 1978 surveys when the firm size threshold for the "long form" with the skilled worker question changed.

Differencing (12) over time removes the time-invariant measurement error from β coefficients but not from γ coefficients. Labeling the periods $t = 0$ and $t = 1$:

$$\Delta y_{ic} = \beta_L \Delta l_{ic} + \beta_S \Delta s_{ic} + \beta_K \Delta k_{ic} + \gamma_L l_{ic} + \gamma_S s_{ic} + \gamma_K k_{ic} \\ + \rho_i + [\Delta\delta_c + \gamma_L u_c^L + \gamma_S u_c^S + \gamma_K u_c^K] + \Delta\varepsilon_{ic}. \tag{13}$$

Under these assumptions the elasticity coefficients β and the factor-bias coefficients γ are identified despite the measurement error. The country effect includes all the bracketed terms: the country-specific change in productivity $\Delta\delta_c$ and the three terms involving country-specific measurement error in factors. There is a symmetric argument for industry-factor specific measurement error, u_i^f, which can be accommodated in the same way, compromising identification of industry specific changes in productivity, ρ_i, but not affecting identification of the elasticities and factor bias terms.

One final measurement issue is that physical units of value added are not observed. We can measure PY (sales net of intermediate inputs), or $p + y$ in logarithms. This is a familiar problem in production function estimation whenever the price deflator is suspect: unobserved price changes may be correlated with changes in input use, generating an omitted variable bias. The ability to estimate industry effects adds a novel element to the solution. Consider the reduced form regression of Δp on Δy (which cannot be estimated for lack of data):

$$\Delta p_{ic} = a_i + b_c + m\Delta y_{ic} + \nu_{ic}. \tag{14}$$

Here a_i and b_c are industry and country fixed effects in price changes. The coefficient m cannot be signed as it is an average of the (inverse) demand and supply elasticities of industry output, weighted by the variances of demand and supply shifts. Since these variances are conditional on common industry effects across countries, they can be interpreted as local supply and demand shifts. For instance, *m* will be positive if the variance of local demand shifts exceeds that of local supply shifts and the price elasticity of demand exceeds that of supply (in magnitude). The scalar *m* would be small if trade makes product demand quite elastic.

Adding Δy to both sides of (14) and then substituting (13) on the right hand side yields:

$$\begin{aligned}\Delta p_{ic} + \Delta y_{ic} &= a_i + b_c + (1+m)\Delta y_{ic} + v_{ic} \\ &= (1+m)\begin{bmatrix}\beta_L \Delta l_{ic} + \beta_S \Delta s_{ic} + \beta_K \Delta k_{ic} \\ + \gamma_L l_{ic} + \gamma_S s_{ic} + \gamma_K k_{ic} \\ + \rho_i + \delta_c + \Delta\varepsilon_{ic}\end{bmatrix} \\ &\quad + a_i + b_c + v_{ic}.\end{aligned} \tag{15}$$

Thus, unmeasured price changes introduce an ambiguity. The coefficients of (13) are identified only up to a proportion (1+*m*). If we assume constant returns, then the sum of the estimated β coefficients is $\sum_f (1+m)\beta_f = (1+m)$, which provides an estimate of *m*.

An alternative approach to estimating the factor-bias terms is to use the relationship in Eq. (11) by regressing TFP growth on the level of inputs.[54] Assuming constant returns to scale and competitive markets, the rate of TFP growth is:

$$\Delta TFP_{ic} = \Delta y_{ic} - \psi_{ic}^{L} \Delta l_{ic} - \psi_{ic}^{S} \Delta s_{ic} - \psi_{ic}^{K} \Delta k_{ic},$$

[54] This approach is similar to that taken by Kahn and Lim (1998) in their study of skill-augmenting technological change in the US. In their estimating equation the shares appear as covariates and they are forced to impose an adding up constraint on these, similar to URS.

where the weights ψ^f are the value-added shares of each factor. This calculated TFP change becomes the dependent variable in the specification implied by (11):

$$\Delta TFP_{ic} = \gamma_L l_{ic} + \gamma_S s_{ic} + \gamma_K k_{ic} + \rho_i + \delta_c + \Delta \nu_{ic}. \tag{16}$$

The same fixed effects and time trends as the output regression in (13) are also included.

Both regression specifications have analogous two-factor variants. Let E be total employment. The output specification in (13) becomes:

$$\Delta y_{ic} = \beta_E \Delta e_{ic} + \beta_K \Delta k_{ic} + \gamma_E e_{ic} + \gamma_K k_{ic} + \rho_i + \delta_c + \Delta \varepsilon_{ic}, \tag{17}$$

while the TFP version is:

$$\Delta TFP_{ic} = \gamma_E e_{ic} + \gamma_K k_{ic} + \rho_i + \delta_c + \Delta \nu_{ic}. \tag{18}$$

The 1980s data allow us estimate both two- and three-factor specifications and then bound the omitted variable bias (see Appendix). The comparison implies that the labor-bias coefficient from a two factor regression, γ_E, is an attenuated version of the three factor one (γ_L).

2.4 Data

This project uses three versions of the United Nations' industrial statistics database. All three provide total employment and wagebill, value added, and gross fixed capital formation for manufacturing industries across multiple nations but differ in temporal and geographic coverage. Table 2.2 summarizes these differences. The first dataset is the General Industrial Statistics (GIS) Database, which runs from 1967-1993 and covers 28 manufacturing industries (3-digit ISIC rev.2) across multiple countries (United Nations, 1992). The GIS's main advantage over the other two datasets is that it provides employment and wage data for both skilled and unskilled workers (operatives and non-operatives in the UN's terminology). This distinction allows us to estimate rates of both absolute and relative labor-saving technological change using the three-factor specifications in Eqs. (13) and (16).

After 1993, the collection of industrial statistics passed from the United Nations Statistics Division to the UN Industrial Development Organization (UNIDO). The second dataset, INDSTAT3 2006, is the direct descendant of the GIS, down to the same 3-digit industry codes (UNIDO, 2006). This database extends into the late 1990s and provides a wider range of countries, including many newly industrializing, low income economies. The most notable difference between this version and the older GIS is that INDSTAT3 does not provide separate statistics for skilled and unskilled workers. Because only total employment and wages are observed, we can only estimate labor-saving technological change using the two-factor estimating equations (17) and (18).

The third dataset (INDSTAT4 2008) is intended to address some unreported methodological shifts in INDSTAT3 (see the Appendix for details). This newest version of UNIDO's database contains data from the 1990s onward with an emphasis on within-country consistency in reporting (UNIDO, 2008). Because INDSTAT4 reports data at the more detailed 4-digit industry level, fewer countries are available. The industrial classifications in INDSTAT4 were also updated to ISIC rev.3 and unfortunately cannot be directly matched to those in INDSTAT3. Though the two databases cannot be directly merged together, it is still possible to estimate the two-factor regression specifications and gauge the extent of labor-saving technological change using the cleaner, more recent, data.

The 19 nations in the GIS sample are arranged into two income groups: a high income group with GDP per capita exceeding $10,000 (1985 US$) in 1980 and a middle income group with GDP per capita between $2,000 and $10,000 in 1980.[55] The two INDSTAT editions use slightly modified income cutoffs; middle income countries are those with per capita income between $8,000 and $18,000 (2000 US$) while high income countries have GDPs above that.[56] These income categories are slightly higher in real terms than their 1980s analogs but contain a relatively similar set of countries. There is also the addition of a new, low income group, comprised of those with per capita GDP between $1,000 and $8,000. Twenty-seven countries, grouped into these three income tiers, are used from INDSTAT3. Ten of these countries fall into the low-income category; another ten are high income, and the remaining seven are middle income. INDSTAT4 contains useable data for 17 countries, of which five are middle income and the remainder are high-income.

Skilled vs. Unskilled

The measure of skill in the GIS data is classification into non-production and production workers. The term "production worker" usually refers to employees directly engaged in production or related activities of the establishment. It includes

[55] GDP data and deflators are from the Penn World Tables 5.6, where the base is the 1985 US dollar.
[56] The two INDSTAT databases are deflated using PWT 6.2, where the base is the 2000 US dollar.

clerks or working supervisors whose function is to record or expedite the production process. Employees of a similar type engaged in activities ancillary to the main activity of the establishment and those engaged in truck driving, repair and maintenance and so on, are also classified as production workers.

This is far from the ideal measure of "skill," which would include elements of education and training. In addition, the educational level of these worker categories is likely to differ across countries. However, two pieces of evidence indicate that non-production workers do indeed have higher educational attainment than production workers. First, past examinations of matched worker and employer surveys have revealed a fairly tight relationship between years of schooling, occupation and non-production categories, at least in the 1990s (Berman, Bound, & Machin, 1997; Machin, Ryan, & Van Reenen, 1996; Harris, 1999). In addition, non-production workers tend to be uniformly better paid. Quality indices based on a comparison of CPS and ASM data in the US suggest that about ½ of skill upgrading in US manufacturing took place within non-production and production categories over the 1980s (Berman, Bound & Griliches, 1994). While the aggregation problems are worse than usual for these categories, within country comparisons are probably reasonable measures over periods as long as a decade, as is done here. Between country comparisons, especially across income ranges, should be viewed with caution.

Summary Statistics

Table 2.4 reports descriptive statistics for the nineteen GIS countries used. The ten middle income countries are from Asia, Europe and South America. This group includes several countries with large manufacturing sectors: (the former) Czechoslovakia, Korea, and Spain. The high income group includes nine countries ranging in income from Japan to the US. The choice of 1985 exchange rates favors the US, but note that US value added per worker is twice as high in 1980 as that of West Germany, the second-ranked country in this group. The US is also the largest manufacturing employer, with 19 million workers, followed by Japan with 10.5 million, the UK with 6.5 million and West Germany with 6.3 million.

Total factor productivity growth in this sample is only slightly higher among the developed countries than among middle income countries. The standard deviation is almost three times as high among middle-income countries, reproducing Figure 2.3's pattern of selective convergence. Note also that manufacturing industries in high income countries have a much faster absolute decline in production worker employment.

Table 2.5 provides summary statistics for the regression variables in the 1990-1997 INDSTAT3 data. To reiterate, these newer data differ from the 1980s panel in two important aspects. First, they include a new tier of low-income countries which allows us to examine newly-industrializing countries.[57] Second, they do not include disaggregation by skill groups.

[57] Berman and Machin (2000) did not find evidence of SBTC in low-income countries in the 1980s.

Within this sample, value added increased by about 3.1% annually in the 1990s, matching the 3.1% increase among the nineteen countries in the GIS sample. Output growth appears to be negatively correlated with income; manufacturing value added in the lowest tier of countries in the sample grew by over 6% annually, even faster than in the middle income countries in the previous decade. Manufacturing growth in developed countries was more muted than in the 1980s; the middle-income countries in INDSTAT3 experienced output growth of about 2.3% annually, compared to 4.2% in the GIS data, while output growth in the richest countries also declined by about half, from 2.1% to 1.2%.

In the 1990s, employment in manufacturing grew by 0.62% annually, compared to declines of 0.34% in the 1980s. This increase is mostly driven by the low income developing economies where employment increased by a whopping 4% annually. Employment trends in the other 17 countries are similar to their 1980s pattern; the number of workers slightly increased in middle income countries and sharply decreased in the high income ones.

One additional observation from Table 2.5 bears mentioning: some of the standard deviations are disturbingly large compared to their 1980s values, the most glaring of which are those for value added and employment in high income countries. In the GIS data, the nine wealthiest countries only had a standard deviation of 2.4 and 2.7 for production employment and output, respectively. In the 2006 INDSTAT, this had increased to 6.9 and 7.5. This variation likely reflects data irregularities in

INDSTAT3 rather than true heterogeneity in growth rates since the summary statistics for the more recent INDSTAT4 show smaller standard deviations.[58]

Rapid capital accumulation continued in the 1990s; countries sampled accumulated capital by 4.9% annually, compared to 2.6% in the previous decade.[59] While low income countries experienced the most manufacturing investment, the 4.9% yearly increase in capital stock overall is not driven solely by the low income tier. Both middle and high income countries demonstrated accelerated capital accumulation. Capital stock in middle income countries grew by 5.3% in the early 1990s, compared to 3.7% previously. Among the wealthiest countries investment also grew faster in the 1990s, rising to 2.7% annually from 1.6% in the 1980s.

Surprisingly, TFP growth actually appears to be negative in Table 2.5. This is due to a combination of two factors: large fluctuations in national price levels and the aforementioned reporting errors among some of the high and middle income countries in this dataset. The first issue is a potential concern since value added, and hence TFP, is measured in dollars, making them susceptible to fluctuations in price indices.[60] However, when productivity growth is calculated without adjusting value added for price increases, the means are all significantly positive.[61] In addition, regression specifications will include country fixed effects, absorbing measurement error in

---

[58] The Appendix documents further evidence that these jumps are due to noise, not true manufacturing growth.
[59] See Appendix for details about constructing capital accumulation from reported "gross fixed capital formation."
[60] Capital is also measured in dollars, but price fluctuations tend to be dampened by the length of the time series.
[61] The unadjusted average annual TFP growth rates are: low-income 136% (116), middle-income 49% (41), high-income 22% (45).

suspect national price deflators. While the second issue cannot be corrected in our data, the analysis can be repeated on a secondary data source for these countries. While we cannot validate the results for low income developing countries using INDSTAT4, they are also the least afflicted by these reporting changes.

Table 2.6 reports summary statistics for the data from INDSTAT4. Manufacturing value added grew by 1.5% annually overall, most of which was in middle income countries (3.2%). This rate is about a percentage point higher than in the 1990s (2.3%), which is somewhat surprising since output growth slowed down between the 80s and early 90s data. The average growth rate of manufacturing value added for high income countries has continued to decline, from 1.16 to 0.73% annually. While they are still noticeably larger than the original GIS counterparts, the standard deviation for labor in both middle and high income countries are more reasonable than those from INDSTAT3 even though there are far fewer observations in INDSTAT4. This suggests that the discontinuities in total employment in the 2006 INDSTAT are less noticeable in the newest data but not completely eliminated.

Potential Pitfalls in Estimation

While the TFP specification in Eq. (16) is less restrictive than the Cobb-Douglas production function version in (13), it also suffers from additional estimation issues. First, *measurement error* is likely in the levels of factors, which is both transitory and industry-country specific, so industry and country effects will not absorb it. This could be anything from fluctuations in unmeasured quality, to price changes in capital to coding error. One implication of transitory measurement error is that it appears on both sides of Eq. (16), creating the potential for spurious correlation between factor levels and ΔTFP. To illustrate, let f_t be a vector of measured factors in period t. Then $f_t = f_t^* + u_t$ where f_t^* is the true level and u_t is classical measurement error, uncorrelated with f or y. The change in TFP would then be calculated as:

$$\begin{aligned} \Delta TFP &= \Delta y - \psi^{'}\Delta f \\ &= \Delta y - \psi^{'}\Delta f^* - \psi^{'}\Delta u \\ &= \Delta TFP^* - \psi^{'}\Delta u. \end{aligned}$$

As Δu appears in the dependent variable and as part of the regressors, this measurement error creates a spurious negative correlation with f_t and a spurious positive correlation with f_{t-1}.

A convenient solution is to use the average level of factors over time as regressors in (16). Let $f = (f_t + f_{t-1})/2 = f^* + (u_t + u_{t-1})/2$ and Σ_t denote the variance of u_t. The spurious covariance is $\psi^{'}\operatorname{cov}(\Delta u, \Delta u)\psi/2 = \psi^{'}(\Sigma_t - \Sigma_{t-1})\psi/2$, which will be zero if the variance of the measurement error is unchanged over time.

A related problem arises with the factor shares $\psi^f = w^f F/Y$ (where $w^f F$ is the wagebill of factor f, and capital's share is calculated as a residual). These include the level of factor f on the left-hand side of (16); transitory measurement error appears in levels on the left-hand side and in logarithm on the right-hand side, inducing a spurious correlation. This correlation can be prevented by predicting ψ_{ic} from a regression of shares on industry and country indicators and using the predicted values to calculate TFP.

A second, more standard, implication of measurement error in factors of production is that bias due to measurement error is exacerbated by differencing, due to the reduction in the signal-to-noise ratio (the ratio of the true variance to the variance of the measurement error). This implies that the estimated elasticities β in Eq. (12) are biased downward. This is a common problem in estimating production functions in differences; the estimated capital coefficient in firm data is often near zero (Griliches & Mairesse, 1995).

The potentially biased β estimates are for the most part incidental, but they could transmit bias to the estimated γ terms through the covariance of estimated coefficients. To see this, consider the least squares estimation of the vectors β and γ, where $X_1 = \Delta f$, and $X_2 = f$ and Δf is correlated with the error term, but f is not. Then the least squares estimator is:

$$\begin{bmatrix} b \\ g \end{bmatrix} = (X'X)^{-1} X'\left[X_1\beta + X_2\gamma + \varepsilon\right]$$

$$\text{so} \begin{bmatrix} b-\beta \\ g-\gamma \end{bmatrix} = (X'X)^{-1} \begin{bmatrix} X_1'\varepsilon \\ X_2'\varepsilon \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} X_1'\varepsilon \\ X_2'\varepsilon \end{bmatrix}$$

Assume $E(X_1'\varepsilon \mid X) = a$ and $E(X_2'\varepsilon \mid X) = 0$.

$$\begin{aligned} \text{Then } & E(b-\beta \mid X) = Aa \\ \text{and } & E(g-\gamma \mid X) = Ca = CA^{-1}E(b-\beta \mid X) \\ & \qquad = \operatorname{cov}(g, b')V(b)^{-1}E(b-\beta \mid X) \end{aligned}$$

under homoskedastic errors.

Aggregation to the industry level helps in this respect, as measurement error between firms tends to cancel, raising the ratio of signal to noise. Defining f as an average over time also helps. It reduces the spurious negative covariance of Δf with f due to measurement error, thus reducing the spurious covariance between estimated β and γ coefficients. A third treatment consists of using prior beliefs about the values of β to bound the possible bias on γ.

A third potential source of bias in the estimation of Eq. (13) is the endogenous response of factor use to an industry-country specific change in productivity or prices. This induces a positive covariance with the error term, $\text{cov}(\Delta f, \Delta \varepsilon) > 0$, and a generally upward bias in the estimated β. Experimenting with restrictions on the estimated β can help gauge how much of this bias is transmitted to the estimated γ.

As input levels are replaced with their time averages in every regression specification, a related concern is that endogenous response will induce a positive correlation between the average *level* of f and the error term, since f_t appears in f. This problem can be treated in the production function specification by using lagged inputs f_{t-1} as instruments since they are determined before $\Delta\varepsilon_t$ is observed.[62] The γ coefficients in Eq. (13) are then identified by cross-industry variation in lagged levels of inputs, which could arise from variation in historical industry-country specific demand or supply conditions in labor, capital, or product markets.

Endogeneity bias is more problematic in the TFP specification since lagged values of inputs are not valid instruments. In the presence of transitory measurement

---

[62] Strictly speaking, that instrument will be invalid in the production function specification, since $\text{Cov}(u_{t-1}, [u_t + u_{t-1}]/2) > 0$, where u_t is the measurement error in measuring the factors f_t. Nevertheless, the induced bias is probably no worse than the standard least squares attenuation bias (which involves the covariance of $(u_t + u_{t-1})/2$ with itself, but also a larger denominator) and would likely tend only to bias estimates towards zero.

error, the error term in the TFP regression will include $-\psi^{'}\Delta u$. Instrumenting with the variable f_{t-1}, which includes the lagged measurement error u_{t-1}, will tend to bias the estimated γ away from zero (in addition to any bias due to classical measurement error, probably toward zero).

In summary, identifying the factor-bias terms in the production function specification appears to be feasible, as the major sources of potential bias can be controlled. However, in the TFP specification a potential endogeneity bias remains untreatable. In practice, comparing the results of the two approaches will turn out to be informative.

2.5 Results for 1980s

This section reports the results for the 1980s panel, in which three factor TFP can be calculated and separate output elasticities can be estimated for both skilled and unskilled labor. Table 2.7 reports the result of estimating the translog specification in Eq. (13). The first three rows report the factor bias coefficients (γ) on log levels of inputs, while the next three report the elasticities (β) on changes in logarithms. Looking first at the β coefficients in the leftmost row, note that they are large, with an estimated β_K of 77.4 and returns to scale of 139. This is not an unusual result in cross-country regressions with developing countries. It may be due to endogenous adjustment of inputs, especially capital, to price and productivity shocks. It may also be due to a positive correlation of prices and quantities of product, reflected in a positive *m* coefficient. These excessive returns to scale recede when we include country effects. The estimated β_K declines to a more reasonable 44.8. This change indicates that the high coefficient on Δk in the leftmost column may have been due to country-specific, cyclical increases in measured productivity. The β's sum to 109; if constant returns holds, the bias due to unmeasured prices is rather small – the estimated coefficients are about 9% too high in absolute value. The reasonable size of the estimated β's from the "country effects" column also provides some reassurance about bias in the estimated β that may be transmitted to the estimated γ coefficients.

The third column adds industry effects in productivity growth, as specified in Eq. (13). This does not much change the estimated β's. Under constant returns, m is estimated at 8%. The addition of industry effects corrects a positive omitted variable bias on the estimate of γ_L in the previous column, changing it from -1.24% to -2.15%. Conditional on country effects, industries with high production worker employment tended to have high measured TFP growth, implying a sector bias (Haskell & Slaughter, 1998) *toward* unskilled workers (or at least industry-specific time-invariant measurement error in inputs). Subtracting 0.16 due to m, the estimated γ_L of -2.15% implies that annual productivity growth is almost 2% slower in industries with twice as many production workers. The estimated standard error is 0.51, indicating strong evidence of absolute labor saving technological change.

The estimated coefficient on skilled labor, γ_S, is positive, at 0.69, but not statistically significant, providing weak evidence of absolute skill bias. Evidence for relative skill-bias is strong, as the estimated value of $\gamma_S - \gamma_L$ is 2.41% (s.e.=1.05%) (not shown in the Table).[63] The estimated coefficient on capital, γ_K is 0.87% (s.e=0.41%), providing strong evidence of absolute capital bias in technological change.

The second to last row reports the change in returns to scale $\gamma_S + \gamma_L + \gamma_K$, which should be zero under unchanging returns to scale (URS). The estimated sum is -0.59%, indicating that increased productivity of skilled labor and capital does not fully compensate for declining productivity of unskilled labor. (This does not imply a productivity decline, since the equation allows Hicks-neutral productivity change.) Changing returns to scale are an uncomfortable finding. They conflict with the replication argument offered in the previous section, since they imply that industries of different sizes have systematically different TFP growth rates. (In this case smaller industries have higher growth rates.) Those objections, and the clear interpretation that URS allows, argue for exploring what happens if URS is imposed.

Restricting the sum of factor bias coefficients raises the estimated skill and capital bias coefficients, yielding an implied γ_L estimate of -1.80%, which is less negative than the unrestricted estimate, or -1.67% corrected for *m*. In other words, conditional on industry and country effects, an industry with twice the capital/unskilled labor ratio and twice the skilled/unskilled labor ratio has an annual TFP growth advantage of 1.67%!

Are these results driven by some outlier, rogue industry or misbehaving country? Figure 2.6 illustrates a leverage plot of the estimated γ_L. It graphs the growth

[63] Relative skill bias as defined in (10) requires that $\gamma_S/\beta_S > \gamma_L/\beta_L$, but we lack the precision to estimate these ratios (which involve four coefficients). The tested hypothesis, $\gamma_S > \gamma_L$, implies (10) under the relatively innocuous assumptions that both β's are positive (since γ_L is negative and γ_S is positive assuming unchanging returns).

rate of value added against the log of production employment, once both have been conditioned on all the other covariates (in the linear regression sense). The upper left panel is a simple scatterplot. The upper right panel is drawn with circles proportional to the weights used in the regression (value-added shares within country). The lower two panels are labeled by country and industry. Combined, the four panels make it clear that estimated labor-saving technological change is not driven by outliers. As a separate robustness check the regression was run dropping a single country each time, which had no substantial effect on the factor bias coefficients.

Table 2.8 presents some specification checks to address possible endogeneity issues. One potential source of bias is the endogenous reaction of factors (l, s, k) to industry-country specific productivity or price changes, which would appear in the residual, $\Delta\varepsilon$. Since factors are measured at their average level between the beginning and end of the period, this may bias estimated coefficients, probably towards one. Using lagged levels (l_{t-1}, s_{t-1}, k_{t-1}) as instrumental variables can treat this problem since these are determined before a productivity or price shock. The column labeled "lagged levels as instruments" reports these instrumental variable estimates. These are essentially identical to the least squares estimates in the previous table. A Hausman test reveals that we cannot reject the hypothesis of identical coefficients: endogenous reaction of factors to productivity or price shocks is not a source of discernible bias.

Another potential source of bias discussed above was bias transmitted from the β coefficients to the γ coefficients. Since the β coefficients are estimated without an instrument in all specifications, they are vulnerable to bias due to endogenous response to productivity or price shocks. Regardless of the source of potential bias, the most suspicious estimated β coefficient is that on the change in non-production workers. At 48.6, it is much higher than the non-production wagebill share in value added. One way to approach the potential transmitted bias is to force this coefficient to take a lower value and observe the change in γ estimates. A possible restriction would be constant returns to scale, imposed in the next column. This exercise has little effect on the β's, so it is not surprising that the γ's are not much changed. A more drastic step is to force the estimated β_S coefficient to be zero, in order to provide an upper bound on the possible transmitted bias. This reduces the estimated γ_S coefficient from 0.69 to 0.47 but has little effect on the other factor bias coefficients. The URS-restricted γ_L estimate rises from -1.80 to -1.64, which can be thought of as an upper bound on the rate of labor saving technological change.

The main conclusions of Table 2.7 are robust to corrections for endogeneity and measurement error biases: very strong evidence that technological change in the 1980s had an absolute labor saving bias, weaker but statistically significant evidence of an absolute capital-bias, and evidence of absolute skill-bias on the borderline of statistical significance. Evidence of relative skill-bias is quite strong, which is consistent with the labor economics literature.

TFP Function Estimates

The total factor productivity specification is more flexible in many ways than the production function as it requires no functional form assumptions except on the factor bias terms. In particular, it does not impose unitary elasticity of substitution between factors. It does require the (standard) assumptions of constant returns to scale and competitive markets to define TFP. (Note that constant returns in the initial period were not rejected in the specifications estimated in Table 2.7, except in the first, which did not include country effects).

Table 2.9 reports estimated factor bias terms as specified in Eq. (16). Despite the difference in specification, the γ coefficients are quite similar to those obtained from the production function, though smaller in absolute value. Our preferred specification, in the third column, includes country and industry effects. The estimated coefficient on production workers is large and negative at -1.66% (s.e.=0.45%), indicating absolute labor saving technological change. The coefficients on non-production workers and capital are positive at 0.44% and 0.68% but neither is significantly different from zero, providing weak evidence of absolute skill-bias and absolute capital bias. The sum of factor bias terms is -0.55% (s.e.=0.33%), providing weak evidence of a decline in returns to scale. If we assume unchanging returns (column 4), the implied estimate of γ_L from the restricted regression is -1.35% (s.e.=0.40%). That estimate is only slightly smaller in absolute value than the restricted estimate of γ_L (-1.80%) from the production function specification. Like the production function estimates, these estimates imply substantially faster TFP growth for skill- and capital-intensive industries.

Omitting industry effects changes the estimated γ_L coefficient to -0.91. This change indicates TFP growth is disproportionately concentrated in industries with high levels of production employment, conditional on country (as in the production function specification). Omitting country effects as well tends to lower the estimated γ_L and γ_S coefficients in absolute value, while raising the coefficients on capital, i.e. countries with high levels of capital and low levels of employment tended to have faster calculated TFP growth.

As discussed earlier, lagged levels of inputs are not valid instruments in the TFP specification, so we cannot correct for endogeneity using IV. However, endogeneity bias was not a discernible problem in the production function estimates, as shown by the similarity of instrumental variable and least squares estimates. If the major form of factor adjustment is through unskilled labor, which has the lowest adjustment costs, then this bias could explain why the TFP estimates have a less negative γ_L estimate. Without an instrument, a conservative approach would be to borrow the estimated m (8%) from the production function specification and deflate the URS-restricted estimate of γ_L from -1.35 to -1.25.

Both approaches show the same pattern: statistically significant evidence of absolute labor-saving technological change, weaker evidence of absolute skill-biased technological change and evidence of capital-biased technological change that is statistically insignificant in the TFP specification but significant in the production function specification. The restricted γ_L estimate summarizes the results neatly (though the sum of factor bias terms rejects that restriction, the unrestricted estimates would make the following a slight *understatement*): conditional on industry and country effects, and allowing for fixed country and industry specific measurement error, a manufacturing industry in the 1980s with double the *K/L* ratio and double the *S/L* ratio is predicted to have an annual TFP growth rate 1.4 to 1.8 percent higher. This is a remarkable level of labor saving technological change, compared with the sample average TFP growth rate of 1.65%.

Middle Income Countries and the Technology Transfer Hypothesis

Tables 2.7-2.9 all report extremely high rates of labor-saving technological change. Are these estimates too large to be believed? In the Cobb-Douglas specification the γ_L coefficient represents the shift in the value-added share of production workers. Using this approach, the shifts reported in Table 2.1 suggest values of γ_L between -0.2% and -0.5%, which are only a fraction of the estimates in Tables 7-9 (-1.4% to -1.8%).[64] Estimated labor saving technological change is also high in another sense. If β_L is about 0.3 and γ_L is about -.015, then by 2010 production workers will be quite useless in production!

A possible explanation for such strong evidence of factor-bias comes from the hypothesis of skill-biased technology transfer. Previous research suggested that during the 1980s, middle income countries absorbed several vintages of technology from high income ones (Berman & Machin, 2000; Conte & Vivarelli, 2011). Perhaps this accelerated technological catch-up induced factor bias in the 1980s for middle income countries at a rate much faster than that experienced at the technological frontier. For

[64] Part of the difference may be due to reallocation of production between industries. Table 2.10 will suggest that these reallocations favor production workers in middle-income countries but disfavor them in high income countries. Yet reallocation between industries is too small to provide most of the answer. A more likely culprit is overly restrictive assumptions about supply and demand in labor market, which underlie that calculation. In particular, Cobb-Douglas implies a unitary elasticity of factor demand. If manufacturing demand for unskilled labor is elastic, then a decline in demand for less skilled workers could result in a very small decline in their wagebill share.

example, if technological convergence is four times as fast in middle income countries as the rate of advance at the frontier, then the labor-saving rate would be $4\gamma_L$ in middle-income countries.

Accelerated factor-biased technology transfer in middle income countries implies that evidence of factor-bias should be stronger in middle income countries than in the high income in the 1980s. Table 2.10 provides a test of that implication, reporting separate regression estimates for the nine high income countries and the ten middle income countries. Dividing the sample reduces precision. For simplicity, only the URS-restricted results are reported.

The high income countries provide a surprise. While the estimates without industry effects are similar to those reported for the sample as a whole, the preferred specification (with country and industry effects) reports labor-biased technological change which is capital-saving. These coefficients are statistically insignificant, so they should not be interpreted as overturning the large body of evidence in the literature suggesting skill-bias in the US and other high income countries. It is more likely that at the level of resolution these data allow, we cannot find skill-bias in these countries.

More interesting is the contrast between the estimated factor-bias coefficients in middle and high income countries. Unlike the high income countries, the 10 middle income countries show strong evidence of capital-bias and labor saving in technological change (in the preferred specification, including country and industry effects). The coefficient indicating skill-bias is positive but imprecisely estimated. The implied γ_L estimate is -2.71% (s.e. = 0.84%), indicating very strong evidence of substantial labor-saving technological change in middle income countries. These results reinforce the view that middle income countries absorbed several vintages of factor-biased manufacturing technology in the 1980s, so that a γ_L estimate of -1.5% (or even -2.5%) overestimates the trend rate of labor-saving technological change at the frontier.

The contrast between estimates with and without industry effects in high and middle income countries sheds light on the sector-bias hypothesis of Haskell and Slaughter (1998). Apparently, industry-specific measured productivity growth disfavored production workers in high income countries.[65] In the middle income countries, the contrast between the results with and without industry effects indicates that industry effects in measured productivity favored *production* workers. Overall the pattern in both subsamples of countries is consistent with the prediction of Heckscher-Ohlin trade theory in a period of declining trade restrictions: price changes favored capital and skill intensive industries in countries with high skill and high capital intensity, while price changes favored industries intensive in unskilled labor in countries with low skill and low capital intensity.[66] Once industry effects in productivity growth are accounted for, the full extent of labor-saving technological change in middle income countries is evident.

---

[65] These results suggest that the ambiguity expressed by Kahn and Lim (1998) about the interpretation of their estimates as evidence of skill augmenting technological change was well founded. They could not include industry effects in the same way as they had only one country to work with.

[66] The pattern of these price effects is inconsistent with the argument that demand for skills increased in middle income countries because of foreign outsourcing to low income countries (Feenstra & Hanson, 1996), as that would predict industry effects in the opposite direction.

Bias Due to Aggregate Employment

Before moving on to results from more recent data, the 1980s GIS data also allow us to analyze omitted variable bias due to having only total, rather than skilled and unskilled, labor data. As most employment in manufacturing is unskilled, the coefficient on employment in the two-factor regressions should be close to that for unskilled labor. Aggregating the two types of labor in the GIS allows us to gauge the attenuation in estimating labor-saving technological change when estimation is limited to just two factors of production (see Appendix). The results here suggest that the regression specifications reported below, which will use only total employment rather than skilled and unskilled workers, underestimate labor-saving technological change by about 75%. Estimates of capital-bias remain relatively unchanged.

Table 2.11 provides estimates of two- and three-factor regressions from both output and TFP specifications using the GIS data. The first four rows report the factor bias coefficients γ while the next four rows report elasticities β. The last row imposes the URS restriction and estimates the factor-bias coefficient on unskilled labor (or total employment in the case of the two-factor model). For convenience, Column (1) reproduces the estimates from the preferred specification in the third column of Table 2.7, along with its γ_L under URS. The rows labeled "Employment" and "Δ Employment" provide estimates for the output elasticity of labor in the two-factor equations in (17) and (18). The coefficient on γ_E in Column (2) is only 74% as large as the γ_E in column (1), -1.59 compared to -2.15. The sum of elasticities is also slightly larger at 111 instead of 108; after correcting for *m*, the two-factor regression in the second column implies that labor-intensive industries grow -1.4% slower than those with half as many workers.

The other results are similar to the three-factor estimates. The capital-bias coefficient is estimated to be 0.89, compared to 0.87 previously, while its β coefficient was estimated to be 39.75 instead of 37.9. Both regressions reject URS, but the sum of factor-bias coefficients is slightly larger in magnitude in the second column.

The imposition of URS is particularly interesting because it is a very restrictive assumption in the two-factor model. In the three-factor case, $\gamma_L = -\gamma_S - \gamma_K$; capital-bias ($\gamma_K > 0$) and skill-bias ($\gamma_S > 0$) allows us to estimate the degree of labor-saving. The two-factor model, which cannot separately identify labor-saving and skill-augmenting technology, will result in an estimate of γ_E closer to zero. This is borne out in the last row of Column (2); γ_E is estimated to be -0.99 (0.44) compared to -1.59 without the URS restriction. Whereas γ_L^{URS} was almost as large as the unrestricted estimate in the original regression, it is only 62% of the unrestricted γ_L for the two-factor regression. Note that even with the limitations of only two factors of production and URS, there is still evidence of labor-saving technological change; an industry with twice the capital-employment ratio is predicted to grow about 1% faster per year. The coefficients between the two and three factor regressions have the same signs and magnitudes within both high and middle income countries, implying that evidence of technology transfer is still apparent in these specifications with total employment.

The last two columns of Table 2.11 repeat this exercise for the TFP specification. Recall that the factor-bias coefficients were similar between the three-factor output and TFP regressions, though the estimates were smaller in magnitude in the latter. Similarly, the two-factor TFP estimates are comparable to those from the (two-factor) output specification and also closer to zero. Even in the attenuated TFP variant, the results still indicate significant labor-saving technological change in the 1980s; γ_E is -1.25 (0.52). Industries with double the employment experience 1.25% slower TFP growth. Compared to the output regression in Column (2), this estimate of γ_E is only three-quarters as large. As before, we fail to reject the assumption of URS, and imposing URS brings the estimate of γ_E even closer to zero, so much so that it is not statistically different from zero (γ_E^{URS} is -0.76 with a standard error of 0.45).

In summary, the estimation results in Tables 7-11 suggest a number of observations about technological progress in the 1980s. First, there is significant evidence of labor-saving technological change among the 19 middle and high income countries in this sample. Because we are able to distinguish between skill and unskilled labor in the GIS, this productivity growth specifically favors capital and, to an extent, skilled workers over unskilled ones. Productivity growth appears to be 1.8-2.2% higher annually for industries with twice the levels of capital and skill. Endogeneity bias is negligible. Second, factor-biased productivity growth is even evident in the more flexible TFP specification, though the extent of estimated labor-saving technological change is lower (1.4-1.7%). Since the output regressions all failed to reject CRS, the implicit assumption of constant returns in calculating TFP is relatively harmless. Third, labor-saving technological change is more prominent in middle income countries, which can import production technology from the highest-income, technology-innovating nations. Middle income countries with twice the capital- and skill- intensity grow 2.7% faster annually while high income countries did not demonstrate significant labor-saving technological change. Fourth, even the more restrictive two-factor (labor and capital) specifications provide evidence of significant labor-saving (and capital-biased) technological change. The extent of labor-saving, however, is noticeably muted since two opposing effects (skill-bias and labor-saving) are both loaded onto the single labor coefficient; γ_E is only 74% as large as γ_L.

Do these results hold for the 1990s? Do we see labor-saving technological change in that decade, both among middle- and low-income countries? If so, is it greater among the newly industrializing, low-income countries, as the technology transfer hypothesis would suggest?

2.6 Results for 1990s

The top panel of Table 2.12 presents estimates of the output specification in (17) using INDSTAT3. The first column includes all 27 countries, reporting weak evidence of labor-saving technological change: γ_E is estimated to be -1.16 (s.e. = 0.60). The estimate has low precision, and likely exhibits attenuation due to having only one labor measure. In contrast to the 1980s results, m is actually negative, implying that the output estimates here are slightly underestimated. Correcting for m implies a labor-bias coefficient of -1.20. The data provide strong evidence that technological change in the 1990s was significantly capital-biased (γ_K = 1.03, s.e. = 0.49), though the point estimate is smaller than that of labor.

As with the results for the 1980s (Table 2.7), we cannot reject constant returns for the full sample of countries. However, unlike the 1980s results, we also fail to reject URS. Imposing URS provides a smaller, borderline significant estimate of the factor-bias coefficient on labor (γ_E^{URS} = -1.01, s.e. = 0.49). Due to the limitation of having only two factors of production, only one factor-bias coefficient is estimated in the URS-constrained results of Table 2.12; evidence of labor-saving technology in this row mathematically implies capital-bias as well. Manufacturing industries with twice the capital per worker have output growth that is about 1% higher annually, a slightly larger rate than in the 1980s. The sum of the factor-bias coefficients has the same sign as in the 1980s, but is now much smaller in size, -0.13 compared to –0.59. If smaller industries did experience faster growth in the 1980s, this growth differential has become somewhat muted in the 1990s, potentially due to the inclusion of the newly industrializing low-income nations.

Are these results robust to outliers? Figure 2.8 contains added variable plots for the coefficient γ_E corresponding to the full sample estimates, labeled by both country and industry. There are no obvious outliers at either the country or industry level. Omitting the single country-industry point in the upper-left corner does not qualitatively change our finding of labor-saving and capital-biased technological change in the 1990s.[67]

The bottom panel of Table 2.12 lists the results from using the TFP specification in Eq. (18). The first column reports weak evidence of labor-saving and capital-biased technological change, with larger magnitudes than in the output specification, but with very large standard errors. The sum of factor bias coefficients is again negative (-0.65), though of comparable size to the GIS results (-0.58). Imposing unchanging returns yields an estimate of γ_E of -1.47, implying that in the 1990s, industries with twice the capital intensity experience annual TFP growth that is approximately 1.5% than those with lower capital per worker. This coefficient is larger than that from the production function specification, though it is not statistically significant.

Low-Income Countries and Technology Transfer

Is the technology transfer that occurred in middle income countries in the 1980s also at work in the 1990s for the low income, newly industrializing countries? The last three columns of Table 2.12 explore this question by repeating the output and TFP regressions by income group. For the richest countries, both the output and TFP specifications indicate labor-saving and capital-biased technological change, though only the output version has coefficients statistically different from zero. Labor-

[67] That point corresponds to Macao's "other minerals" industry. Omitting Macao entirely does not change the primary finding of significant labor-saving and capital-biased technological change.

intensive industries in high income countries demonstrate an annual growth disadvantage of about 1.3%, after correcting for *m*.

Middle income countries provide a surprise in that the estimated production function coefficient on capital has a negative sign (though it is mercifully insignificant). While the standard errors in the previous column were uniformly larger than those from the pooled regression in the first column, the ones here dwarf the size of the estimates themselves; none of the coefficients are even marginally significant.

The final column demonstrates that the finding of labor-saving technological change in the first column is solely driven by data from the low income countries; the labor-bias coefficient for just low income countries is large, negative, and significant. Industries in newly industrializing countries with half the labor are expected to grow 2.3% faster annually compared to those with higher employment. The difference in the degree of capital-bias is even starker; high-capital industries in low-income countries have a growth advantage of 2.6% annually, compared to 1.7% for high-income nations.

Even more striking is the rejection of constant returns; the sum of elasticities implies that *m* is negative; the γ_E here actually *underestimates* the degree of labor-saving. Alternatively, it could be that the sum of elasticities is actually too small; β_E is 68 and close to the (pooled) 1980s estimate of 72 but the elasticity of capital is 13 and quite far from the old estimate of 40. It might be that the CRS rejection seen in the 1990s is actually due to β_K being far too small. Because these low-income countries are accumulating capital extraordinarily fast (7% annually in Table 2.5), these investments may not have had enough time to become fully productive, resulting in the smaller-than-expected β_K.

Low income countries also demonstrate strong capital-biased technological change as γ_K is positive and significant (and underestimated if we believe the rejection of CRS). Industries with twice the capital levels grow 2.6% more each year compared to those with less capital. This coefficient does not change much even if we impose URS; industries with twice the capital intensity experience 2.6% higher growth.

The fact that the degrees of labor-saving and capital-bias among developing countries are both higher than those for high-income ones is evidence for the technology transfer hypothesis. This is not too surprising since the replication vs. innovation argument applies here, as it did for the middle income countries in the 1980s. The effect might be even stronger for poorer countries since they can import vintages of technology from both the middle and high income countries, resulting in high levels of labor-saving and capital-biased growth. While the labor-saving coefficients do not appear too different between the high and low income countries, the capital-bias coefficient is nearly a full percentage point higher in the last column, potentially because it is easier/quicker/cheaper to import or upgrade new mechanical production improvements than it is to attract and retain high-skill workers.

The income-specific TFP estimates hint at an even more drastic level of labor-saving and capital-bias in low income countries. Assuming constant returns to scale, TFP growth in low-employment industries in the lowest income tier is 4.9% higher than those that use more labor. Capital-bias is also evident in the positive and significant estimate of $\gamma_K = 4.61$ (s.e. = 1.71), which is also much larger than even the high-income nations. Imposing unchanging returns reduces these coefficient slightly ($\gamma_E^{URS} = -4.64\%$) but it is still significantly different from zero (s.e. = 1.73). Capital-intensive manufacturing industries in low income countries appear to have TFP growth that is 4-5% higher annually, though this comes with the caveat that constant returns was rejected in the production specification but we assumed it anyway to calculate TFP.

It appears that low-income countries in the 1990s do exhibit similar, if not higher, levels of labor-saving and capital-biased technological change than the middle income countries of the 1980s. Even within the TFP estimates, there is evidence that newly developing economies may be importing production technology from those on the innovation frontier.

INDSTAT4 and the Middle Income Countries

While there is significant evidence of labor-saving, capital-bias and technology transfer among low income countries in the 1990s, concerns about data quality impeded similar analysis for middle and high income countries. INDSTAT4 allows us to replicate these regressions over a slightly different time window for high and middle income countries. INDSTAT4 coverage ends in 2004, but we begin observing about half of the countries in our data in the mid-1990s (1995-1997).[68] The panels for

---

[68] The countries who first reported to INDSTAT4 in the mid-1990s are: Austria, Belgium, Denmark, Ireland, Israel, Italy, Macao, Portugal, and the US. Half of the 266 country-industry observations in our analysis are from these countries.

the other eight countries are slightly longer since they entered the data in the early-1990s.

The first column of Table 2.13 estimates Eqs. (17) and (18) (the two factor model) using more recent data from INDSTAT4 2008. As before, the estimates imply significant labor-saving (γ_E = -2.79, s.e. = 1.03) and capital-biased (γ_K = 2.01, s.e. = 0.70) technological change. Both of these coefficients are far larger than the 1980s results (Column (2) of Table 2.11). Even after correcting for the new *m*, the labor bias coefficient has grown from -1.6 to -2.2 while γ_K has gone from 0.9 to 1.60. The sum of factor bias coefficients is again negative though only marginally significant. Imposing URS implies a precisely estimated γ_E^{URS} of -2.01, implying greater labor savings than the GIS estimates. Whereas industries with twice the K/L ratio experienced a growth advantage of 1% in the 1980s, these results imply twice that growth advantage in the 1990s, an acceleration of labor saving technological change.

Unlike the 1980s results, these estimates are not driven solely by the middle income countries. Even after splitting the sample, both income tiers show significant labor-saving and capital-biased technological progress. Both factor-bias coefficients are larger in magnitude for middle income nations than they were in the 1980s. Industries with double the labor in middle income countries grow about 2.7% (after correcting for *m*) slower than those with lower employment. This same industry would have had a smaller growth disadvantage of 1.7% in a high income country.

Pooled together, the TFP results in the lower panel indicate labor-saving technological change, though to a lesser degree than in the output version. However, both coefficients are somewhat poorly estimated in the TFP regression; the restricted URS estimate is better estimated though smaller in magnitude. Industries in the late 1990s with twice the capital per worker have 1.5% faster annual productivity growth.

In contrast to the output results in the upper panel, the estimates from the TFP regression appear to be driven mostly by the middle income countries. The high income ones display weak evidence of labor-saving and capital-bias, while both coefficients for the middle income countries are strongly significant, even assuming URS. High employment industries in middle income countries have productivity growth about 3.6-3.8% slower than those with half the labor.

The bottom panel of Figure 2.8 provides the added-variable plot for the γ_E coefficient in the first column of Table 2.13. Macao has two outliers, but they are not quite as obvious as the one in Figure 2.8. Norway also has some industries of concern, but omitting Macao and Norway both individually and together do not significantly alter the regression results.

Taken together, the two INDSTAT versions indicate that technological change in the 1990s is significantly labor-saving and capital-biased. In the 1990s, manufacturing sectors with double the capital per worker experience productivity growth that is 1.7 – 2.4% higher per annum than lower intensity ones. In addition, there is very strong evidence that these occur in low income countries in the 1990s at a similar, if not higher, rate than in the middle income countries of the 1980s. Estimates of the labor saving coefficient for low income countries range from 2.3% to 4.9%. Technology transfer is one plausible explanation for the high rates of labor-saving and capital-bias in low income developing countries as their coefficients are much higher than those for high income ones Even in the late 1990s, both high and middle income developed countries still demonstrate significant labor saving.

2.7 Implications

The 1980s estimates from the GIS (assuming unchanging returns) lend themselves to a straightforward interpretation. The US has about twice the measured *K/L* and *S/L* ratios as Cyprus and Portugal. The estimated rates of labor-saving bias, between 1.4% and 1.8% annually, imply TFP growth rates 1.4 to 1.8 percent higher in US manufacturing than in the manufacturing sectors of those countries. Thus, all other things equal, manufacturing value added per worker will diverge quite quickly, with the labor productivity gap doubling every 39-50 years. So why don't we observe divergence? Capital intensity in middle income countries is about half that of high income countries, and skill intensity is about 2/3 (though correcting for measurement error would lower that figure). For lower income countries the factor intensity gap is even larger.

One possible explanation for lack of TFP divergence was suggested at the outset: replication is faster than invention, and this technological catch up compensates for factor bias. Another possibility is that URS does not hold in the 1980s, despite the replication argument offered: smaller industries truly had higher TFP growth rates, a force which favored convergence and partially compensated for the factor bias effect. This is the pattern suggested by the data as the sum of estimated

γ coefficients was consistently negative. Note that these estimates cannot be interpreted as evidence for technological catch up across countries (or industries), as they are present in specifications that already include country effects.

The extent of factor bias compensation (through these or some other mechanisms) can be roughly examined by seeing how much of the cross-country variance in TFP growth rates is explained by country effects in a (URS restricted) regression which allows factor-bias.[69] Figure 2.7 reports the result of this exercise in a plot of TFP growth rates against GDP per capita. Points labeled are the country effects in the industry and country effects specification for the pooled sample, reported in the rightmost column of Table 2.7.[70] Squares represent TFP growth rates for these countries, as in the right panel of Figure 2.3. Estimated country effects exceed TFP growth in all the middle income countries and are lower than the TFP growth rate in the high-income ones. Thus, country effects and income per capita are negatively correlated (illustrated by the regression line), indicating that once we account for factor-bias, there is evidence of TFP convergence. This negative correlation should not be overemphasized, as t = -0.9 for this regression. On the other hand, if middle-income countries did not tend to overstate measured skill intensity, the slope would be even more negative. Similarly, if we used the middle-income factor-bias coefficients from

[69] This calculation is not completely accurate: estimated country effects include not only the true country effect in TFP growth but also an estimation bias due to measurement error in factor levels. For instance, if a country miscodes less-skilled labor as skilled, and $\gamma_S + \gamma_L$ is negative, the estimated country effect will be biased downwards.

[70] A constant has been added to estimated country effects so that their mean is the same as that of the TFP growth rate. Otherwise they would reflect the conditional mean TFP growth rate with S/L and K/L set equal to unity, which would be an unusual country indeed.

Table 2.10, the slope would also be more negative. For these two reasons, TFP convergence conditional on factor-bias is stronger than indicated by the figure.

A final implication of labor saving technological change is this: if ratios of capital and skilled labor to unskilled labor are increasing (as would be efficient), TFP must *accelerate* under the simplifying assumption that $\beta_f^{'} = \gamma_f$. Eq. (11), together with this restriction, implies that:

$$\frac{d^2 TFP}{dt^2} = \sum_f \gamma_f f^{'}(t) .$$

This condition is difficult to test as TFP fluctuates considerably over time. Nevertheless, two things are worth noting: first, in the very long run, measured labor productivity has accelerated (Kremer, 1993), and second, this TFP acceleration is a fairly direct implication of the considerable evidence of skill-bias in the labor economics literature.

2.8 Conclusion

Factor-biased technological change, a familiar finding for developed countries in the labor economics research, also provides a plausible explanation for the lack of cross-country convergence in total factor productivity. In the 1980s, most of the cross-national variation in growth rates of manufacturing value added per worker is TFP growth. Thus a factor-bias explanation for lack of convergence in TFP growth rates provides most of the explanation for lack of convergence in value added per worker in manufacturing. These, in turn, are highly correlated with (non-convergent) growth rates in GDP per capita.

The empirical literature generally attributes slow international convergence in income levels to country-specific institutional and geographic factors and market failures within individual countries.[71] *Within*-country variance from the manufacturing industries of various countries, both developed and developing, provides a fresh source of information, orthogonal to the finding of "conditional" convergence.

The data yield strong evidence that technological change is absolutely labor-saving, absolutely capital-biased and relatively skill-biased. Estimates are large, suggesting that a country or industry with twice the capital and skill intensity will have a total factor productivity growth rate 1.4% - 1.8% higher annually. The data are unusually rich, allowing us to estimate factor-bias coefficients that allow for country and industry effects in TFP growth. Estimated factor bias coefficients are driven for the most part by the ten middle-income countries, suggesting that accelerated technology transfer to these countries in the 1980s caused unusually rapid, factor-biased technological change.

The findings of absolute labor-saving and capital-bias also extend in the 1990s, across both middle and high income countries. The estimated growth advantage for capital-intensive industries here is about 1.5-2% annually. In addition, low income countries, who were excluded from the 1980s results for lack of data, demonstrate even stronger factor-biased technological change. Capital intensive industries in

[71] For a survey see Barro and Sala-I-Martin (1995), Ray (1998) or Weil (2000).

developing countries exhibited annual growth rates that were 2.5% higher. As middle income countries in the 1980s were apparently able to utilize technology transfer to achieve high manufacturing output growth, so too were low income developing countries in the 1990s.

More generally, these results are based on manufacturing data from individual decades, so extrapolation to entire economies over longer periods should be done with caution. These data show considerable similarity to the Baumol-Barro-style 1960-2000 non-convergence diagram (the triangles and correlations of Section 2.2), but suggest that a country accumulating skill and capital intensity experiences a twofold benefit: both an immediate increase in labor productivity and a repositioning which increases the benefit from future (absolute) skill and (absolute) capital bias in technological change. In this second sense current savings increase future growth. Yet, Solow convergence through factor accumulation is quite slow (in these data or in Mankiw, Romer and Weil (1992), for example).[72] Whatever economic mechanisms slowed factor accumulation in poorer countries over recent decades positioned them badly for factor-biased TFP growth.

Does factor-bias forever stifle convergence? Theory suggests not. Return to the two-factor illustration in Figure 2.5 and imagine a (closed economy) Solow or Ramsey growth model augmented with labor-saving technological change. Designate B as the Ramsey steady state in which the marginal product of skill (human capital) is equal to the rate of time preference. Cross country convergence would be the motion from A to

[72] That is the prediction of a model with constant returns to skill and capital combined (Barro, 1991). Interestingly, these data cannot reject that possibility, especially for the middle income countries for which the point estimates indicate slightly increasing combined returns for skill and capital.

B, as skill-scarce (but otherwise identical) countries increase skill-intensity (*S/L*) and thus decrease the disparity in income per capita. The relative wages of skilled workers fall along this path till they reach their Ramsey steady state level. Now consider the comparative statics of a (surprise) skill-biased technological change that shifts the isoquant for all countries from F_{t-1} to F_t. The new Ramsey steady state will be at a point like D, where the marginal product of skill is again equal to the rate of time preference. The shift in isoquants implies faster TFP growth for countries with higher skill intensity and causes divergence in income per capita.

The Ramsey model augmented with factor-biased technological change admits both periods of divergence and periods of convergence. This interpretation of the cross-country data is inherently hopeful about convergence. Despite factor-bias, Solow's decreasing returns mechanism eventually induces all countries to arrive at point D, with equal income per capita. This argument, combined with empirical estimates of large and pervasive labor-saving technological change in manufacturing, underscores the importance of establishing the relative importance of factor-bias, market failures in accumulation, failures in technology transfer or absorption, and other factors in explaining slow convergence not just within manufacturing, but in other economic sectors and across a range of countries. The recent evidence of strong skill-bias in services, combined with the growing importance of the service sector in both developed and developing countries (Jorgenson & Tiller, 2011; Hendricks, 2010), suggests that the gains to skill accumulation for today's economies could be even larger than documented here.

2.9 Acknowledgement

Chapter 2 is currently being prepared for submission for publication of the material. Berman, Eli; Tiffany Chou. The dissertation author was a primary investigator and author of this material. We acknowledge the helpful comments of Robert Barro, Francesco Caselli, Simon Gilchrist, Samuel Kortum, Kevin Lang, Dilip Mookherjee and Andrew Weiss; seminar participants at Boston University, Case Western Reserve University, University of Pittsburgh, Rice University, UC San Diego, NBER Labor Studies and Productivity sessions, NEUDC and SOLE conferences. This work continues research with Stephen Machin during which a version of these data were prepared.

Table 2.1: Factor Shares in Value Added

	Middle Income			High Income		
	~1980	~1990	**change**	~1980	~1990	**change**
Production	25.5	23.1	**-2.4**	30.3	25.7	**-4.6**
Non-production	10.1	10.2	**+0.1**	18.4	18.0	**-0.4**
Capital	64.4	66.7	**+2.3**	51.3	56.3	**+5.0**

Notes: The wagebill shares of production and non-production workers are wagebill/value added. The capital share is the complement so that the three shares sum to one. These figures are calculated from the UN GIS database, using the same sample as regression results in the other Tables, which is restricted to countries for which capital can be calculated both near the beginning and near the end of the 1980s. Middle Income countries have GDP per capita between $2,000 and $10,000 US in 1980. They are: Turkey, Columbia, Czechoslovakia, Malta, Portugal, Chile, Cyprus, South Korea, Ireland and Spain. High Income countries (those with GDP per capital above $10,000 US) are: Japan, UK, Austria, Finland, Denmark, West Germany, Sweden, Australia, US.

Table 2.2: Summary of Data Sources

Data Source	GIS	INDSTAT3	INDSTAT4
Years covered	1967 to 1993	1963 to 2002	1991 to 2005
- Years used for first diff.	~1980 to ~1990	~1990 to 1997	~1995 to 2004
Number of countries used	19	27	17
- Low income	-	10	-
- Middle income	10	7	5
- High income	9	10	12
Number of industries identified	28	28	61
- Used in regression	28	23	17
Income Categories (in thousands)			
- Low income	-	1-8	1-8
- Middle income	2-10	8-18	8-18
- High income	>10	>18	>18
- In year…	1980	1990	1990
Currency measured in:	1985 USD	2000 USD	2000 USD
- Using PWT version…	5.4	6.2	6.2
Disaggregated employment?	Yes	No	No

Notes: INDSTAT4 identifies 61 industries at the 3-digit level, but these are aggregated up to the 2-digit level to provide similar industry classifications as GIS, of which 17 are used in regression. Five industries were dropped between the GIS and INDSTAT3 (tobacco, industrial chemicals, other chemicals, petroleum refining, and other petroleum productions).

Table 2.3: Sample Descriptive Statistics - 1980

Country	GDP/capita (1985 $)	Manuf. Value Added per Worker ($)	Manuf. Value Added as % of GDP	Manuf. Employment (1000s)	Annual Production Wage ($)	Annual Non-Production Wage ($)	Fraction Non-production	Years of Education per adult
A. Middle-income Countries								
Turkey	2872	5780	14	795	3290	4312	0.22	2.6
Colombia	2948	4662	23	508	2660	5139	0.27	4.2
S. Korea	3093	6764	28	2015	3346	4772	0.21	6.8
Czechoslovakia	3731	5651	-	2472	2780	3064	0.27	
Chile	3898	7472	21	206	4711	14496	0.27	5.9
Malta	4488	7790	-	25	5826	11584	0.15	6.3
Portugal	4982	2390	-	663	4157	6766	0.14	3.2
Cyprus	5289	6990	-	36	4884	7252	0.16	7.2
Ireland	6828	11894	-	212	12929	18383	0.19	7.6
Spain	7391	8835	-	1159	11842	16478	0.23	5.2
B. High-income Countries								
Japan	10068	18467	29	10500	10506	11908	0.46	8.2
UK	10161	13988	27	6462	14559	19045	0.30	8.3
Austria	10499	15657	25	679	11602	19309	0.30	6.2
Finland	10843	16256	28	531	13645	20597	0.24	9.6
Denmark	11333	15664	20	381	22356	29948	0.28	10.1
W. Germany	11916	20262	-	6302	20810	31450	0.28	8.5
Norway	12141	14360	15	354	18619	25869	0.26	10.3
Sweden	12447	17813	23	853	17520	27207	0.29	9.5
Australia	12518	15702	19	1138	16380	19517	0.26	10.1
US	15311	40078	22	19200	18357	28145	0.28	11.9

Notes: All manufacturing figures are from the United Nations General Industrial Statistics Database. These apply to 1980 except where otherwise noted. GDP per capita is from the Penn World Table. Percent of GDP in manufacturing is from World Development Indicators, 1999. Years of education/adult (aged 25 or more) are from the Barro-Lee data. All pecuniary figures reported in 1985$ deflated by the implicit Laspeyres GDP deflator in the Penn World Table. Employment reflects the sample rather than the entire population as surveys typically include only plants with ten or more employees.

Table 2.4: Descriptive Statistics, GIS

	All 19 countries		10 middle income countries		9 high income countries	
Growth rate (x100):						
Total factor productivity	1.57	(3.40)	1.51	(4.52)	1.64	(1.76)
Value added	3.11	(5.30)	4.16	(6.90)	2.10	(2.71)
Employment	-0.34	(3.08)	0.34	(3.57)	-1.00	(2.34)
Production	-0.73	(3.16)	-0.02	(3.65)	-1.42	(2.42)
Non-production	0.62	(3.36)	1.44	(3.91)	-0.18	(2.49)
Capital	2.64	(3.47)	3.70	(4.24)	1.62	(2.07)
Log level of:						
Employment	10.79	(1.88)	9.94	(1.73)	11.62	(1.64)
Production	10.44	(1.83)	9.64	(1.73)	11.20	(1.59)
Non-production	9.46	(2.05)	8.44	(1.82)	10.45	(1.76)
Capital	21.53	(2.09)	20.40	(1.90)	22.62	(1.63)
Observations	422		197		225	

Notes: Observations are at the country-industry level. Of the 532 potential observations (28 industries x 19 countries) 422 are available (79%). Appendix 1 provides details of coverage by industry. Standard deviations in parentheses. Observations are weighted by their within-country value-added share. Total factor productivity is calculated using wagebill shares of value added as weights. These weights are predicted by regression using a full set of country and industry indicators. Production worker weights are predicted with an R^2 of 0.84 and non-production worker weights are predicted with an R^2 of 0.77. Capital weights are calculated as the complement so that the weights sum to one. The middle income countries are Turkey, Colombia, South Korea, Czechoslovakia, Chile, Malta, Portugal, Cyprus, Ireland, and Spain. The high income countries are Japan, the UK, Austria, Finland, Denmark, West Germany, Sweden, Australia, and the US.

Table 2.5: Descriptive Statistics, INDSTAT3

	All 27 countries		10 low income countries		7 middle income countries		10 high income countries	
Growth rate (x100):								
Total factor productivity	-0.07	(4.39)	-0.22	(5.66)	-0.90	(3.24)	0.59	(3.66)
Value added	3.14	(7.43)	6.13	(7.35)	2.32	(6.11)	1.16	(7.50)
Employment	0.62	(6.85)	4.00	(5.88)	0.14	(6.10)	-1.91	(6.91)
Capital	4.89	(5.34)	7.16	(6.43)	5.29	(4.70)	2.74	(3.61)
Log level of:								
Employment	10.71	(1.81)	10.73	(1.63)	9.82	(1.80)	11.27	(1.74)
Capital	21.79	(2.14)	21.51	(1.60)	20.92	(2.49)	22.56	(2.04)
Observations	524		190		133		201	

Notes: Observations are at the country-industry level. Of the 621 of potential observations (27 countries x 23 industries), 524 are used (84%). Standard deviations are in parentheses. Differences are calculated from 1990 to 1997 and weighted by within-country value-added share. Total factor productivity is calculated using wagebill's share of value added as the weights. These weights are predicted by regression using a full set of country and industry indicators. Capital weights are calculated as the complement so the weights sum to one. The low income countries are India, Indonesia, Iran, Jordan, Malaysia, Mexico, Panama, the Philippines, Sri Lanka, and Turkey. The middle income countries are Cyprus, Greece, Ireland, Israel, Malta, South Korea, and Spain. The high income countries are Austria, Finland, Hong Kong, Italy, Japan, Macao, Norway, Singapore, the UK, and the US.

Table 2.6: Descriptive Statistics, INDSTAT4

	All 17 countries		5 middle income countries		12 high income countries	
Growth rate (x100):						
total factor productivity	-0.29	(3.42)	-0.29	(4.77)	-0.29	(2.66)
value added	1.47	(4.94)	3.18	(5.89)	0.73	(4.28)
employment	-0.45	(4.20)	0.74	(4.75)	-0.96	(3.85)
capital	3.65	(3.62)	5.85	(3.34)	2.71	(3.31)
Log level of:						
employment	11.19	(1.54)	10.90	(1.24)	11.32	(1.64)
capital	22.97	(1.83)	22.73	(1.33)	23.07	(2.00)
Observations	266		78		188	

Notes: Observations are at the country-industry level. Of the 289 of potential observations (17 countries x 17 industries), 266 are used (92%). Standard deviations are in parentheses. Differences are calculated from the mid-90s to early 00s, depending on data availability, and weighted by within-country value-added share. Total factor productivity is calculated using wagebill's share of value added as the weights. These weights are predicted by regression using a full set of country and industry indicators. Capital weights are calculated as the complement so the weights sum to one. The middle income countries are Ireland, Israel, South Korea, Portugal, and Spain. The high income countries are Austria, Belgium, Denmark, Finland, France, Italy, Japan, Macao, Norway, Singapore, the UK, and the US.

Table 2.7: Factor Bias Estimates from Production Function

		country effects	.. & industry effects	.. & imposing unchanging returns
Production	-1.46	**-1.24**	**-2.15**	-
	(0.72)	(0.44)	(0.51)	
Non-production	-0.17	**0.77**	0.69	**0.89**
	(0.69)	(0.29)	(0.43)	(0.44)
Capital	**1.51**	0.58	**0.87**	**0.91**
	(0.61)	(0.45)	(0.41)	(0.42)
Δ Production	34.9	22.5	21.5	19.3
	(12.9)	(14.5)	(13.8)	(14.2)
Δ Non-production	27.1	41.7	48.6	49.9
	(9.9)	(9.4)	(9.2)	(9.5)
Δ Capital	77.4	44.8	37.9	37.3
	(8.3)	(8.5)	(6.6)	(7.0)
19 country effects		X	X	X
28 industry effects			X	X
R^2	0.65	0.84	0.87	0.87
Sum of elasticities	139	109	108	107
(β's)	(10)	(09)	(11)	(11)
Sum of factor bias	-0.11	0.10	-0.59	0
term (γ's)	(0.20)	(0.13)	(0.24)	-
γ_L (assuming URS)				**-1.80**
				(0.51)

Notes: All specifications include 422 observations of industries within countries. Standard errors (in parentheses) are heteroskedasticity-consistent, allowing a country specific grouped error term. Factor bias coefficients in bold are significant at the 5% level. The dependent variable is the annualized change in log value added (x100). Observations are weighted by value added share within each country. The sum of factor bias coefficients sums estimated coefficients of production workers, non-production workers and capital. The coefficient γ_L assuming unchanged returns to scale is the estimated coefficient on production workers, using the same specification but restricting the three factor bias coefficients to sum to zero. For descriptive statistics see Table 2.4. Estimating equation is (13) in text.

Table 2.8: Factor Bias Estimates from Production Function - Specification Checks

	lagged level as instruments	.. & un-changing returns	impose constant returns	..& un-changing returns	impose $\beta_s = 0$	..& un-changing returns
Production	**-2.10**	-	**-2.07**	-	**-2.07**	-
	(0.55)		(0.53)		(0.58)	
Non-production	0.63	0.82	0.70	0.89	0.47	0.71
	(0.42)	(0.44)	(0.42)	(0.43)	(0.52)	(0.52)
Capital	**0.89**	**0.93**	**0.82**	**0.87**	**0.87**	**0.93**
	(0.42)	(0.43)	(0.38)	(0.41)	(0.43)	(0.45)
Δ Production	21.5	19.4	17.1	15.8	58.2	56.7
	(13.9)	(14.2)	-	-	(10.4)	(10.9)
Δ Non-production	48.6	49.8	48.4	49.6	0	0
	(9.2)	(9.5)	(9.1)	(9.5)	-	-
Δ Capital	37.9	37.3	34.5	34.6	48.5	48.1
	(6.6)	(7.0)	(5.4)	(5.5)	(7.4)	(7.9)
R^2	0.87	0.87	0.78	0.77	0.84	0.84
Sum of elasticities	108	107	-	-	107	105
(β's)	(11)	(11)			(12)	(12)
Sum of factor bias terms	-0.57	0	-0.55	0	-0.72	0
(γ's)	(0.24)	-	(0.26)	-	(0.30)	-
γ_l (assuming URS)	**-1.76**	**-1.76**	**-1.76**	**-1.76**	**-1.64**	**-1.64**
	(0.57)	(0.57)	(0.52)	(0.52)	(0.55)	(0.54)

Notes: All specifications include 422 observations of industries within countries. Standard errors (in parentheses) are heteroskedasticity-consistent, allowing a country specific grouped error term. Factor bias coefficients in bold are significant at the 5% level. The dependent variable is the annualized change in log value added (x100). Observations are weighted by value added share within each country. The sum of factor bias coefficients sums estimated coefficients of production workers, non-production workers and capital. The coefficient γ_L assuming unchanged returns to scale is the estimated coefficient on production workers, using the same specification but restricting the three factor bias coefficients to sum to zero. The "constant returns" specification imposes constant returns to scale. The dependent variable in that case is Δlog(value added) - Δlog(production). For descriptive statistics see Table 2.4. Estimating equation is (13) in text.

Table 2.9: Factor Bias Estimates from TFP Specification

		country effects	..& industry effects	..& unchanging returns
Production	-1.20	-0.91	**-1.66**	-
	(0.73)	(0.45)	(0.45)	
Non-production	0.04	0.73	0.44	0.62
	(0.73)	(0.37)	(0.40)	(0.38)
Capital	1.08	0.42	0.68	0.73
	(0.58)	(0.47)	(0.44)	(0.45)
country effects		X	X	X
industry effects			X	X
R^2	0.09	0.57	0.63	0.63
Sum of factor bias terms (γ's)	-0.09	0.24	-0.55	0
	(0.21)	(0.14)	(0.33)	-
γ_L (assuming URS)	-1.06	**-0.97**	**-1.35**	**-1.35**
	(0.54)	(0.45)	(0.40)	(0.40)

Notes: All specifications include 422 observations of industries within countries. Standard errors (reported in parentheses) are heteroskedasticity-consistent and allow a country specific grouped error term. The dependent variable is the annualized change in TFP (x100). Observations are weighted by their value added share within each country. Total factor productivity is calculated using wagebill shares in value added as weights. The sum of factor bias coefficients sums estimated coefficients of production workers, non-production workers and capital. The coefficient γ_L assuming unchanged returns to scale is the estimated coefficient on unskilled labor, calculated using the same specification but restricting the three factor bias coefficients to sum to zero. For descriptive statistics see Table 2.4. Estimating equation is (16) in text.

Table 2.10: Factor Bias Estimates – High vs. Middle Income Countries (assuming URS)

	High Income			Middle Income		
		country effects	..& industry effects		country effects	..& industry effects
Non-production	**1.47**	0.67	0.23	-1.25	1.05	1.28
	(0.52)	(0.45)	(0.33)	(1.33)	(0.67)	(1.28)
Capital	0.12	0.27	-0.70	**2.48**	0.62	**1.42**
	(0.35)	(0.25)	(0.63)	(0.89)	(0.89)	(0.62)
Δ Production	61.7	49.8	55.6	21.8	12.4	6.1
	(14.8)	(6.6)	(7.4)	(16.4)	(20.8)	(19.0)
Δ Non-production	-0.1	21.8	30.9	43.6	51.7	60.6
	(10.4)	(6.7)	(5.2)	(13.0)	(14.1)	(11.1)
Δ Capital	23.0	19.2	11.2	84.2	50.5	41.7
	(11.8)	(4.4)	(5.4)	(10.0)	(9.9)	(10.4)
country effects		X	X		X	X
28 industry effects			X			X
R^2	0.65	0.76	0.83	0.70	0.85	0.89
Sum of elasticities	85	91	98	149	114	108
(β’s)	(10)	(6)	(8)	(11)	(13)	(16)
γ_L (assuming URS)	**-1.60**	**-0.94**	+ 0.46	-1.23	**-1.67**	**-2.71**
	(0.20)	(0.23)	(0.41)	(1.00)	(0.45)	(0.84)
Observations		225			197	

Notes: Standard errors (in parentheses) are heteroskedasticity-consistent, allowing a country specific grouped error term. Factor bias coefficients in bold are significant at the 5% level. The dependent variable is the change in log value added (x100). Observations are weighted by value added share within each country. The sum of factor bias coefficients sums estimated coefficients of production workers, non-production workers and capital. The coefficient γ_L assuming unchanged returns to scale is the estimated coefficient on production workers, using the same specification but restricting the three factor bias coefficients to sum to zero. For descriptive statistics see Table 2.4. Estimating equation is (13) in text.

Table 2.11: Total Employment vs. Skilled-Unskilled

Dependent variable:	Δlog(value added) x100					
	(1)	(2)	Mid (1)	Mid (2)	High (1)	High (2)
Production	**-2.15**	-	**-3.18**	-	0.38	-
	(0.51)		(0.82)		(0.38)	
Non-production	0.69	-	0.94	-	0.20	-
	(0.43)		(1.21)		(0.30)	
Employment	-	**-1.59**	-	**-2.77**	-	0.50
		(0.49)		(0.71)		(0.56)
Capital	**0.87**	0.89	1.37	**1.77**	-0.69	-0.65
	(0.41)	(0.43)	(0.67)	(0.57)	(0.62)	(0.58)
Δ Production	21.5	-	8.96	-	55.6	-
	(13.8)		(18.0)		(7.28)	
Δ Non-production	48.6	-	43.2	-	30.7	-
	(9.2)		(9.99)		(5.21)	
Δ Employment	-	71.50	-	71.8	-	85.9
		(8.55)		(11.7)		(8.32)
Δ Capital	37.9	39.75	59.7	49.1	11.9	12.0
	(6.6)	(7.12)	(10.8)	(8.2)	(5.51)	(4.6)
R^2	0.869	0.860	0.891	0.879	0.834	0.840
Sum of elasticities	108	111.24	111.83	120.9	98.3	97.9
(β's)	(11)	(10.01)	(15.54)	(11.8)	(7.31)	(7.21)
Sum of factor bias	-0.59	-0.70	-0.87	-1.00	-0.12	-0.15
terms (γ's)	(0.24)	(0.26)	(0.36)	(0.38)	(0.44)	(0.46)
γ_L(assuming URS)	**-1.80**	**-0.99**	**-2.71**	**-1.91**	0.46	0.65
	(0.51)	(0.44)	(0.84)	(0.51)	(0.41)	(0.59)
Number of Countries	19		10		9	
Observations	422		197		225	

Notes: Data come from the GIS. All regressions include country and industry fixed effects, and observations are weighted by within-country industry size. Robust standard errors, clustered by country, are in parentheses. Factor bias coefficients in bold are significant at the 5% level. The coefficient γ_L assuming unchanged returns to scale is the estimated coefficient on unskilled labor, calculated using the same specification but restricting the factor bias coefficients to sum to zero. In the specifications with total employment, this amounts to constraining the two factor bias coefficients to have opposite signs.

Table 2.11: Total Employment vs. Skilled-Unskilled, Continued

Dependent variable:	ΔTFP x100	
	(1)	(2)
Production	**-1.67**	-
	(0.45)	
Non-production	0.44	-
	(0.40)	
Employment	-	**-1.25**
		(0.52)
Capital	0.68	0.67
	(0.44)	(0.43)
Δ Production	-	-
Δ Non-production	-	-
Δ Employment	-	-
Δ Capital	-	-
R^2	0.632	0.629
Sum of elasticities (β's)	-	-
Sum of factor bias terms (γ's)	-0.55	-0.58
	(0.33)	(0.32)
γ_L(assuming URS)	**-1.35**	-0.76
	(0.40)	(0.45)
Number of Countries	19	
Observations	422	

Table 2.12: Factor Bias Estimates, INDSTAT3

Dependent variable: Δlog(value added) x100

	All	High only	Mid only	Low only
Employment	-1.16	-2.02	0.74	**-2.31**
	(0.60)	(0.93)	(1.23)	(0.92)
Capital	**1.03**	1.65	-0.57	**2.56**
	(0.49)	(0.82)	(0.68)	(0.83)
Δ Employment	88.53	107.85	90.56	67.34
	(8.21)	(8.36)	(9.62)	(8.65)
Δ Capital	8.69	11.94	-15.84	12.97
	(6.39)	(16.37)	(21.04)	(7.17)
R^2	0.87	0.94	0.89	0.83
Sum of elasticities	97.22	118.79	74.73	80.31
(β's)	(7.38)	(11.28)	(17.15)	(7.49)
Sum of factor bias	-0.13	-0.36	0.16	0.25
terms (γ's)	(0.32)	(0.43)	(0.98)	(0.64)
γ_L(assuming URS)	**-1.01**	-1.52	-0.55	**-2.55**
	(0.49)	(0.81)	(0.68)	(0.79)

Dependent variable: ΔTFP x100

	All	High only	Mid only	Low only
Employment	-2.19	-0.96	0.24	**-4.86**
	(1.19)	(1.41)	(0.97)	(1.74)
Capital	1.53	0.18	-0.98	**4.61**
	(0.96)	(0.92)	(0.65)	(1.71)
R^2	0.43	0.43	0.43	0.55
Sum of factor bias	-0.65	-0.78	-0.74	-0.25
terms (γ's)	(0.40)	(0.56)	(0.98)	(0.87)
γ_L(assuming URS)	-1.47	-0.09	-1.12	**-4.64**
	(0.98)	(0.77)	(0.69)	(1.73)
Number of countries	27	201	133	190
Observations	524	10	7	10

Notes: Differences are calculated between 1990 and 1997. All regressions include country and industry fixed effects, and observations are weighted by industry size. Robust standard errors, clustered by country, are listed. Factor bias coefficients in bold are significant at the 5% level. Regression specifications are (17) for the top panel and (18) for the bottom. The coefficient γ_L assuming unchanged returns to scale is the estimated coefficient on employment, calculated using the same specification but restricting the two factor bias coefficients to have opposite signs

Table 2.13: Factor Bias Estimates, INDSTAT4

Dependent variable: Δlog(value added) x100

	All	High only	Mid only
Employment	**-2.79**	**-2.12**	**-3.61**
	(1.03)	(0.79)	(0.67)
Capital	**2.01**	**1.46**	**2.58**
	(0.70)	(0.64)	(0.40)
Δ Employment	98.26	89.17	111.35
	(8.61)	(6.43)	(7.10)
Δ Capital	27.34	37.81	20.96
	(7.44)	(7.82)	(13.95)
R^2	0.88	0.88	0.92
Sum of elasticities (β's)	125.60	126.98	132.31
	(12.43)	(11.38)	(20.93)
Sum of factor bias terms (γ's)	-0.78	-0.66	-1.03
	(0.48)	(0.31)	(0.69)
γ_L(assuming URS)	**-2.01**	**-1.47**	**-2.67**
	(0.70)	(0.59)	(0.42)

Dependent variable: ΔTFP x100

	All	High only	Mid only
Employment	-1.64	-0.98	**-3.83**
	(0.84)	(0.68)	(0.61)
Capital	**1.47**	0.73	**3.60**
	(0.69)	(0.60)	(0.83)
R^2	0.35	0.48	0.44
Sum of factor bias terms (γ's)	-0.16	-0.24	-0.23
	(0.31)	(0.23)	(1.22)
γ_L(assuming URS)	**-1.49**	-0.76	**-3.59**
	(0.69)	(0.61)	(0.89)
Number of countries	17	12	5
Observations	266	188	78

Notes: Differences are calculated from mid-90s to early 00s, depending on data availability. All regressions include country and industry fixed effects, and observations are weighted by industry size. Robust standard errors, clustered by country, are listed. Factor bias coefficients in bold are significant at the 5% level. Regression specifications are (17) for the top panel and (18) for the bottom. The coefficient γ_L assuming unchanged returns to scale is the estimated coefficient on employment, calculated using the same specification but restricting the two factor bias coefficients to have opposite signs.

Table 2.14: Industry Coverage

Industry	Observations: 10 Middle Income Countries	9 High Income Countries	Total	Value Added Share (%)
Food	8	9	17	11.8
Beverages	8	8	16	4.3
Tobacco	7	7	14	1.9
Textiles	10	9	19	5.7
Apparel	9	9	18	4.0
Leather Products	10	8	18	0.5
Footwear	9	8	17	1.1
Food Products	9	8	17	2.5
Furniture	8	8	16	1.7
Paper Products	10	9	19	4.7
Printing & Publishing	10	9	19	5.5
Industrial Chemicals	6	8	14	3.6
Other Chemicals	7	7	14	4.4
Petroleum Refineries	6	6	12	1.5
Petroleum and Coal	3	5	8	0.2
Rubber Products	8	8	16	1.2
Plastic Products	5	9	14	2.3
Pottery & China	4	7	11	0.3
Glass Products	6	8	14	1.0
Non-metallic minerals n.e.c.	7	8	15	3.2
Iron and Steel	4	8	12	2.5
Nonferrous metals	5	7	12	2.2
Metal Products	6	9	15	5.6
Machinery	8	9	17	9.8
Electrical Machinery	8	9	17	9.2
Transportation Equipment	8	8	16	6.4
Professional Goods	4	9	13	1.3
Other Goods	4	8	12	0.8
Total	197	225	422	100

Notes: Observations record the number of countries reporting for each industry at both the beginning and end of the 1980s so that a useful observation existed. There are 28 2.5 digit ISIC industries and 19 countries so the potential number of industry-country observations is 532, of which 422 useful observations are available. The value added share reports the average share of manufacturing value added in that industry for countries reporting at the end of the 1980s.

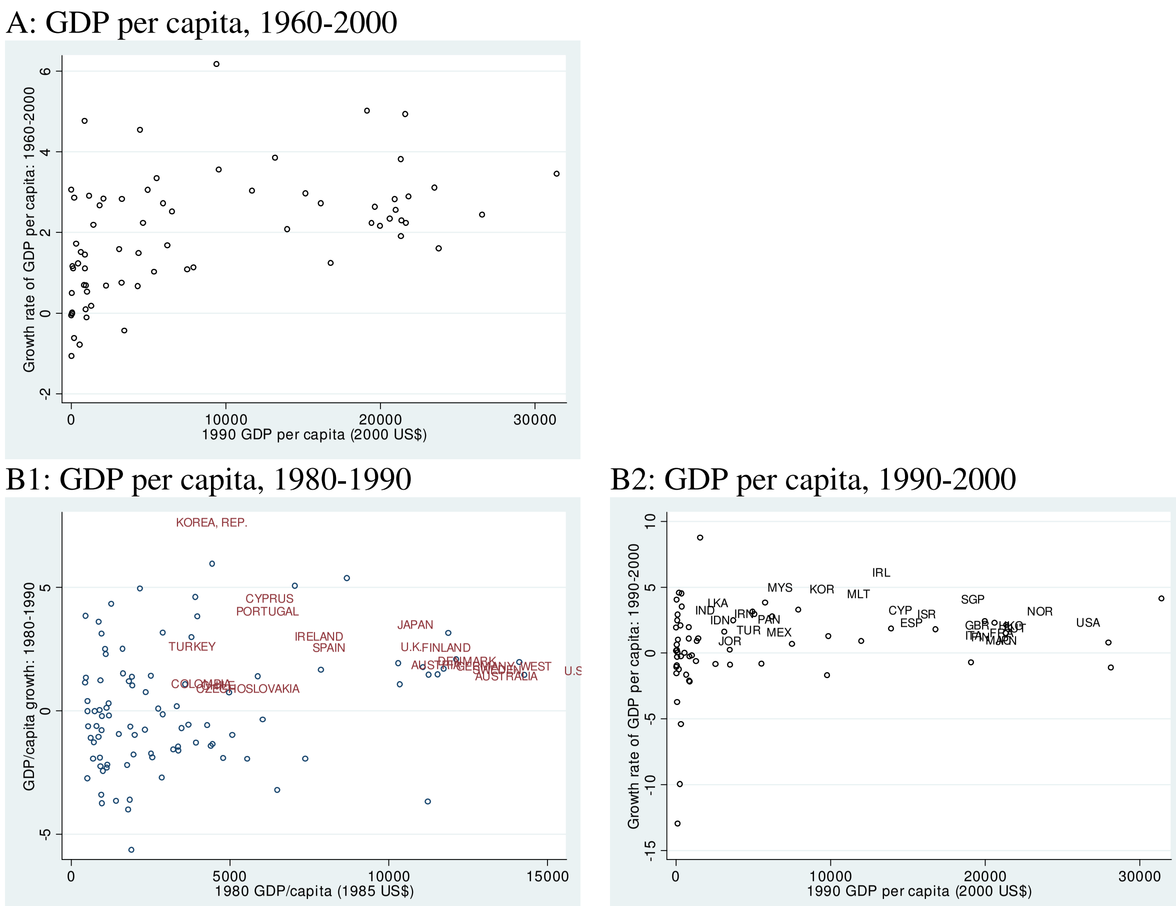

Figure 2.1: Non-convergence of Growth Rates

Notes: GDP data from PWT 6.2. Labeled countries in panel B are those included in the manufacturing sample. Growth rates are annualized and in percentages.

A: 1980s

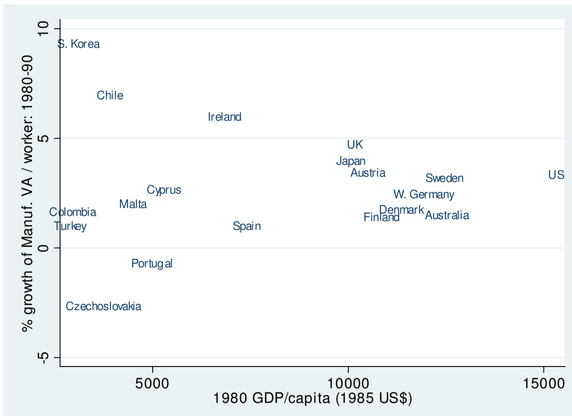

B: 1990s

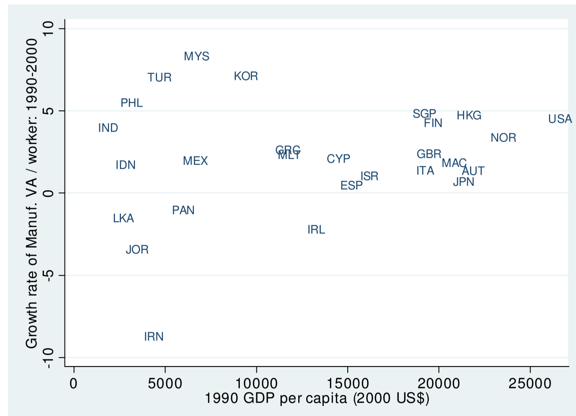

Figure 2.2: Non-Convergence of Manufacturing Value Added (per worker)

Notes: Manufacturing data in the 1980s (left) are from the GIS and deflated using PWT 5.4. Manufacturing data in the 1990s (right) are from INDSTAT3, deflated by PWT 6.2.

A: Factor Accumulation B: TFP Growth

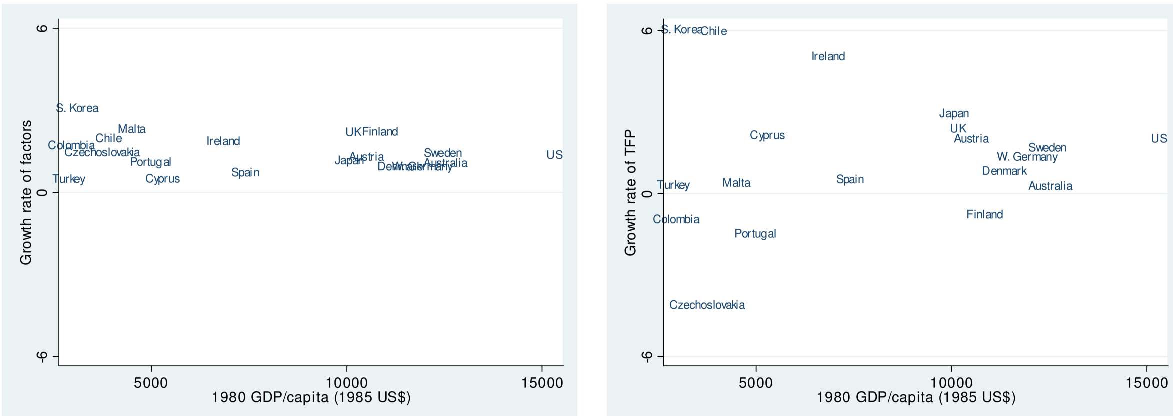

Figure 2.3: Decomposition of Increases in Manufacturing Value Added (per worker), 1980s

Notes: Growth in manufacturing value added per worker decomposed into factor accumulation (left panel) and TFP growth (right panel). GDP statistics from Penn World Table 5.6 while manufacturing data from GIS. Factor of production are skilled labor, unskilled labor, and capital.

A: Factor Accumulation B: TFP Growth

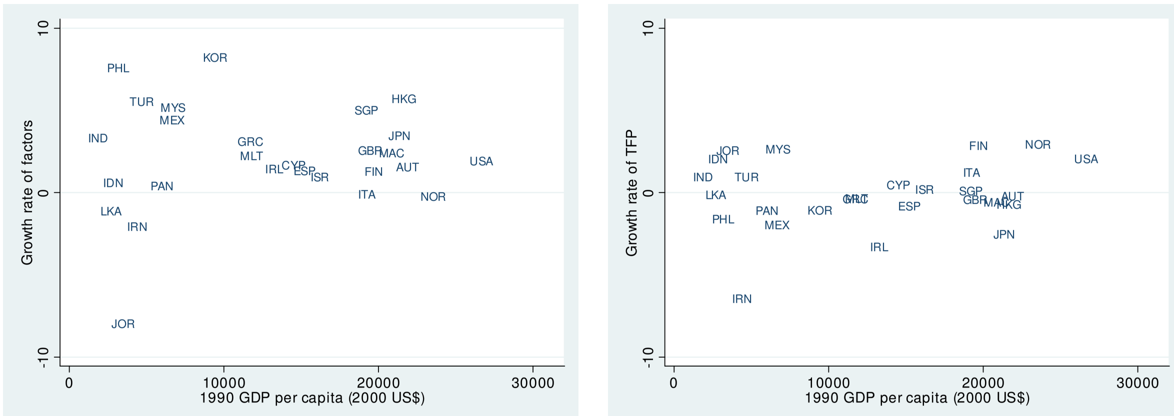

Figure 2.4: Decomposition of Increases in Manufacturing Value Added (per worker), 1990s

Notes: Manufacturing value added per worker decomposed into factor accumulation (left panel) and TFP growth (right panel). GDP statistics from PWT 6.2 while manufacturing data are from INDSTAT3. Factors of production are labor and capital, so factor accumulation is just changes in capital intensity (K/L). There is no measure of skill accumulation here.

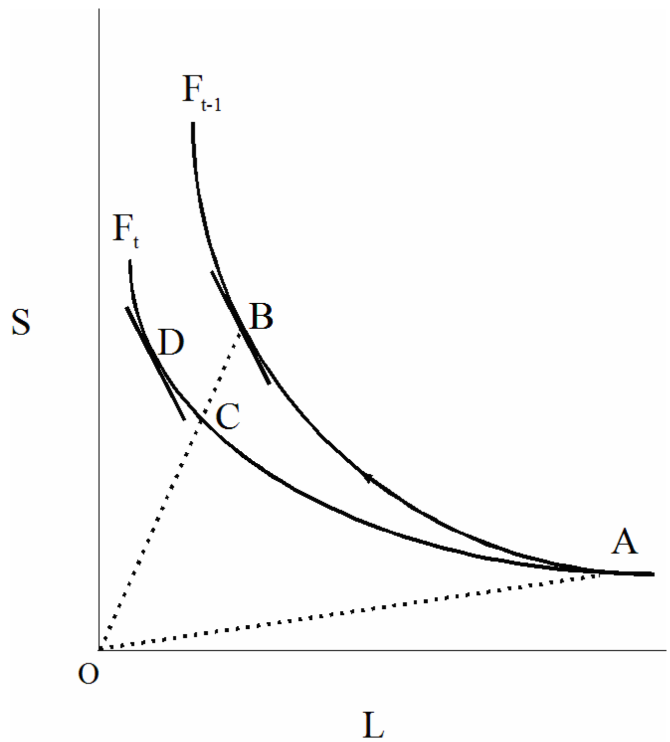

Figure 2.5: Technological Change with Relative Skill Bias

Notes: F_{t-1} and F_t are unit isoquants, holding capital constant. The shift from F_{t-1} to F_t is relatively skill-biased because the S/L ratio is higher at D than B even though both points have the same relative wages. Equivalently, the wage ratio w_L/w_S at C is lower than at B, implying an increase in skilled wages (relative to unskilled), even though both points have the same S/L ratio. Country B exhibits faster TFP growth (the length of the segment BC) than country A because B had a higher S/L ratio in the previous period.

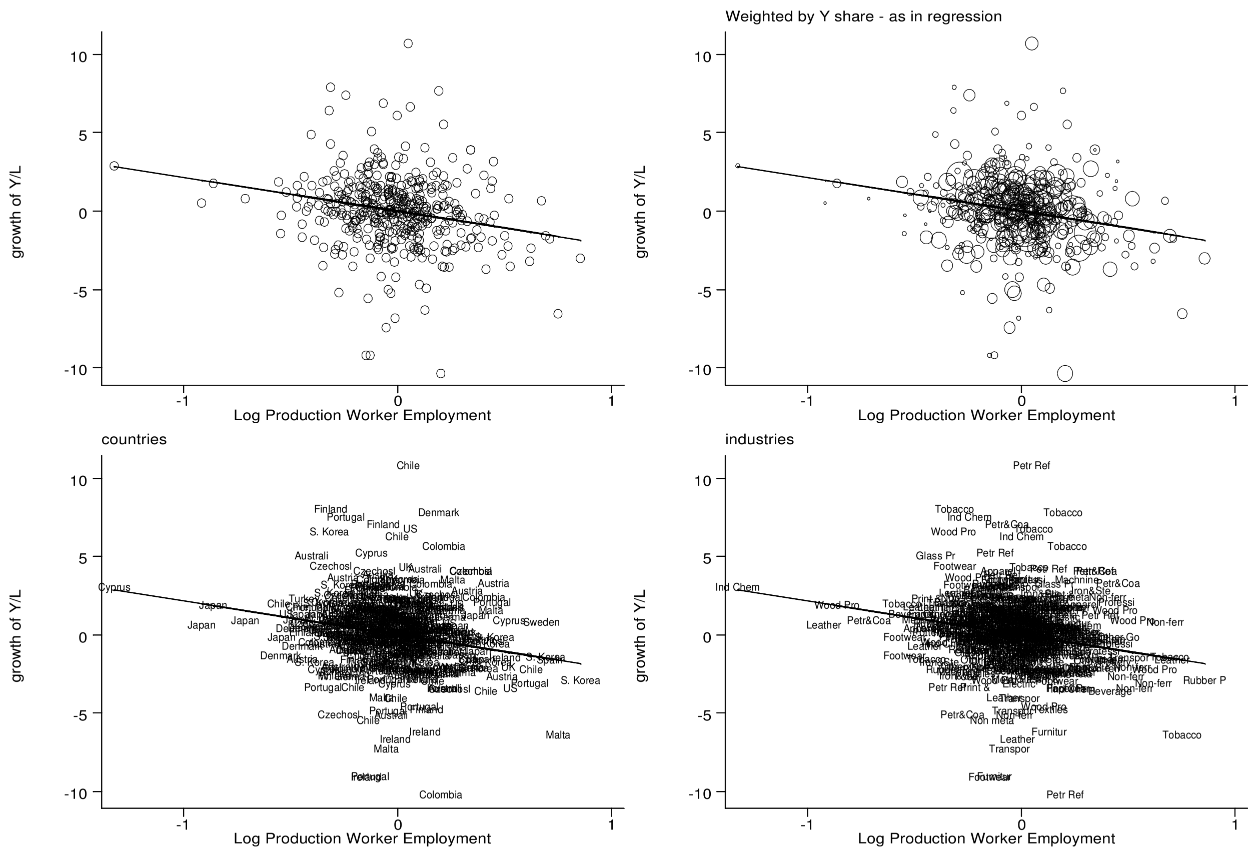

Figure 2.6: Leverage Plots, GIS

Notes: Added variable plots for the factor bias coefficient on production employment (third column of Table 2.7). Top left graph is unweighted; top-right plot has marker sizes relative to weight in regression. Bottom plots label points by country (left) or industry (right).

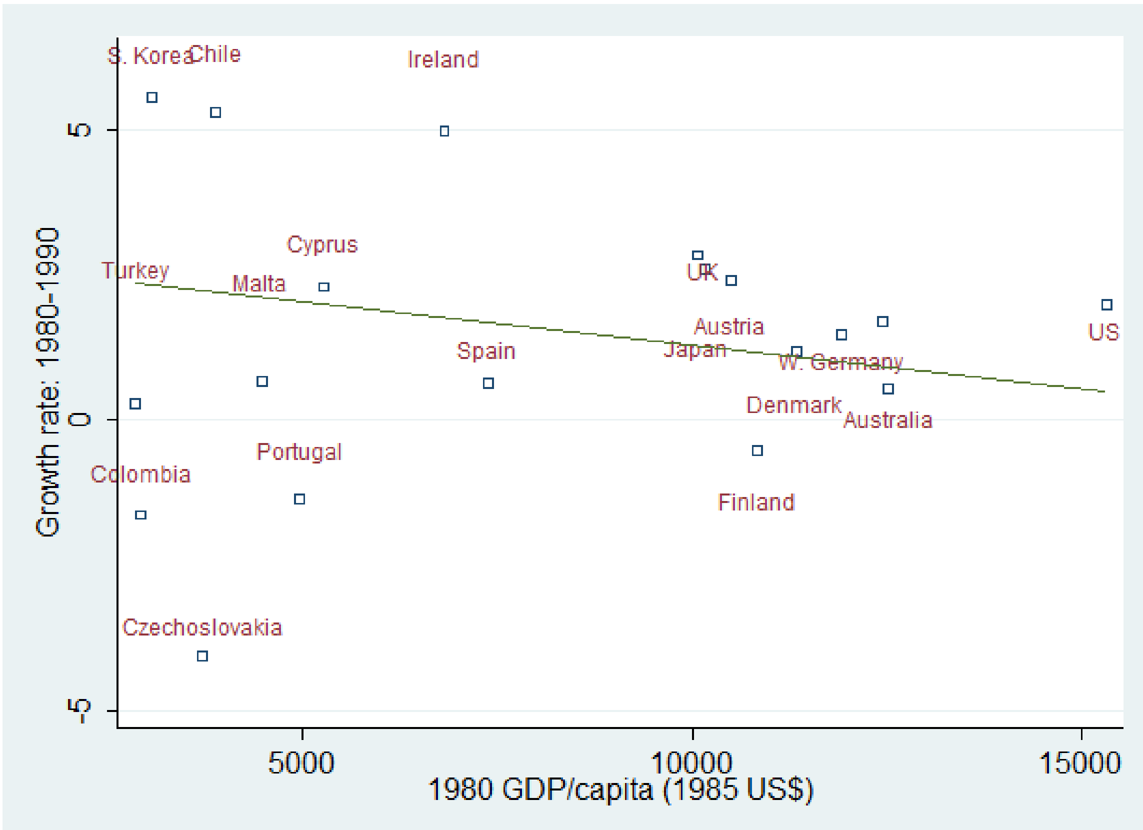

Figure 2.7: Conditional Convergence or Factor Bias?

Notes: Labels are the country effects in the industry and country effects specification for the regression reported in the rightmost column of Table 2.7. Squares represent TFP growth rates for these countries, as in the right panel of Figure 2.3. All country effects are shifted up by a constant so they have the same mean as TFP growth rate. Otherwise they would reflect the conditional mean TFP growth rate with S/L and K/L set equal to unity.

A: INDSTAT3

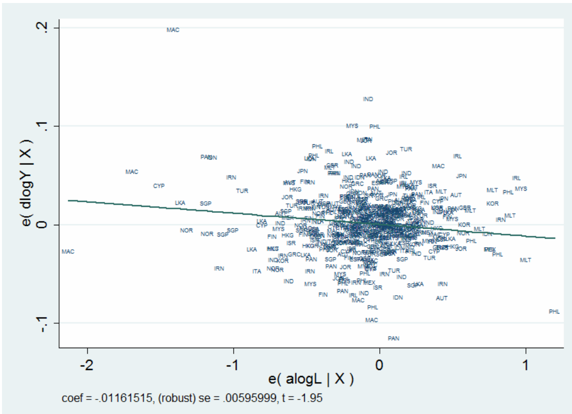

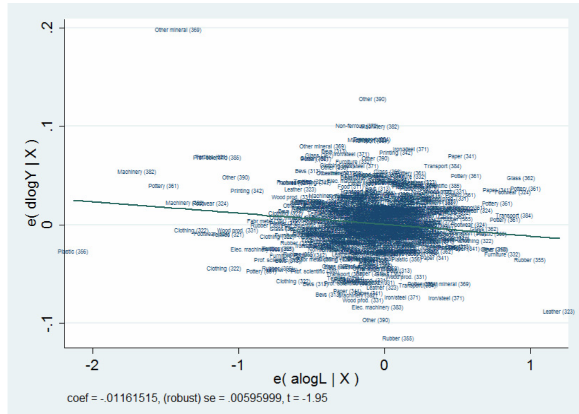

B: INDSTAT4

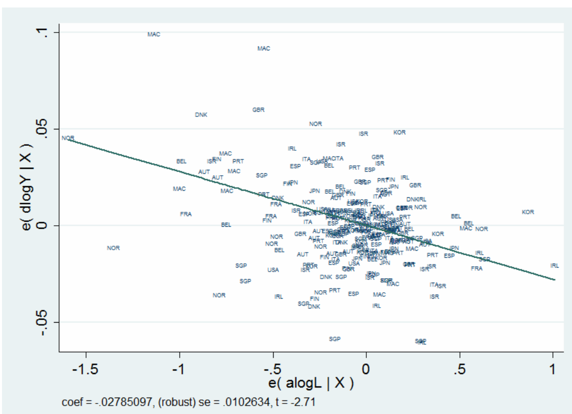

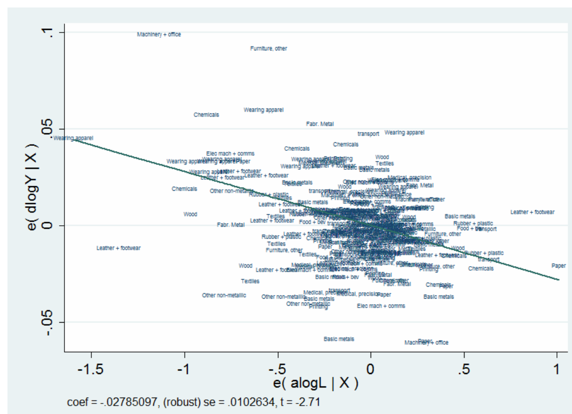

Figure 2.8: Leverage Plots for Output Specification, INDSTAT

Notes: Regression equation is Eq. (13) in the text. The left panel labels points by country, while the right one labels by industry.

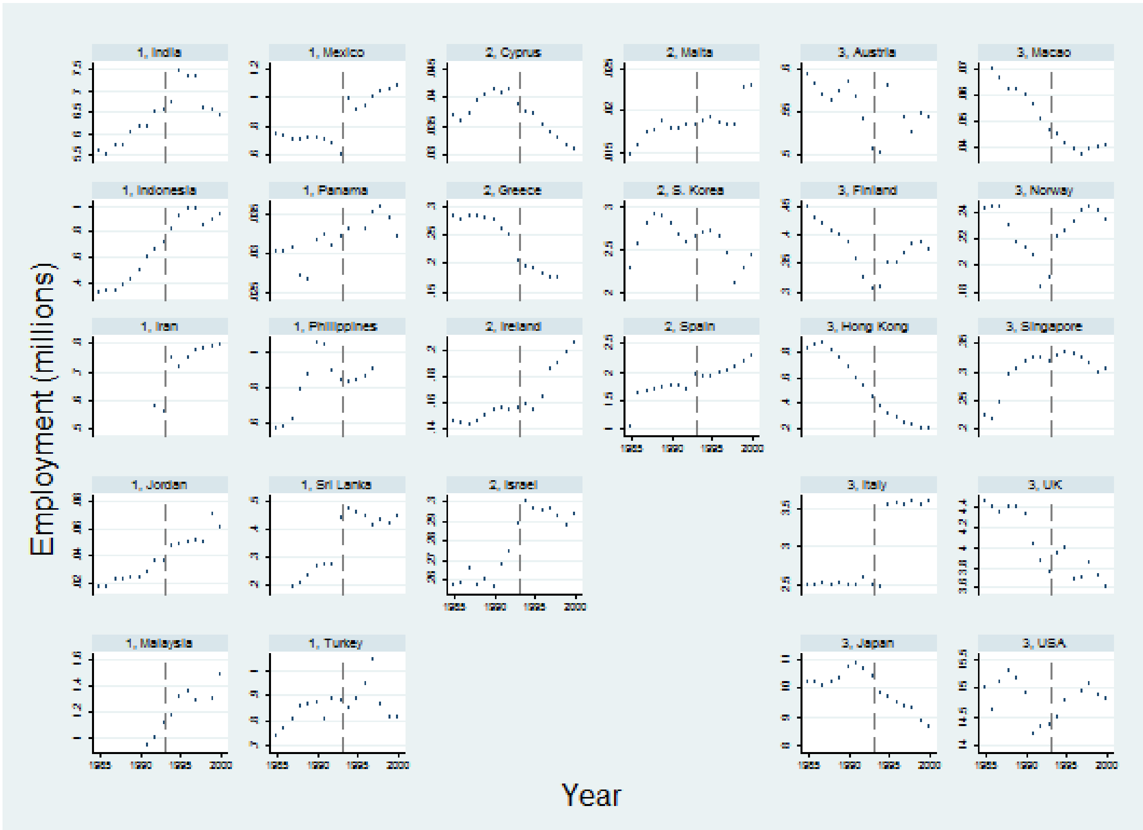

Figure 2.9: Total Manufacturing Employment, By Country

Notes: Data from INDSTAT3.

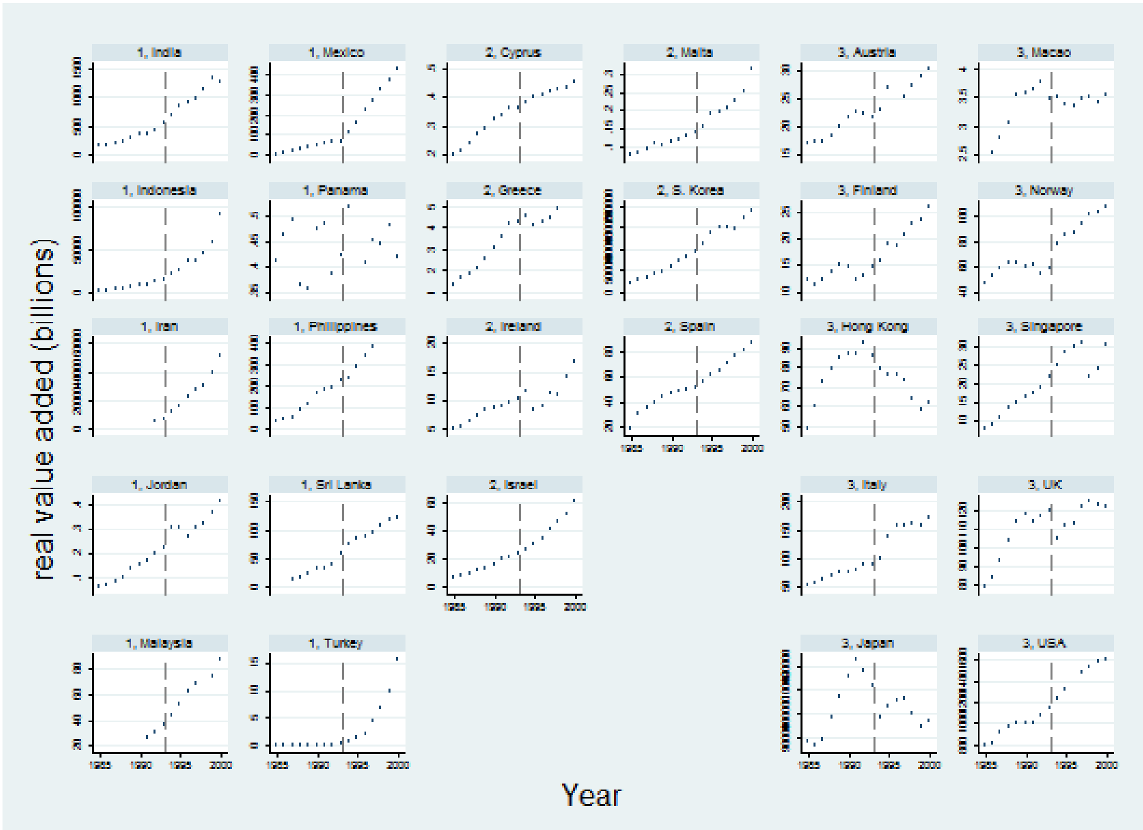

Figure 2.10: Total Manufacturing Value Added, By Country

Notes: Data from INDSTAT3. Value added deflated to 2000 US$ using Penn World Table 6.2.

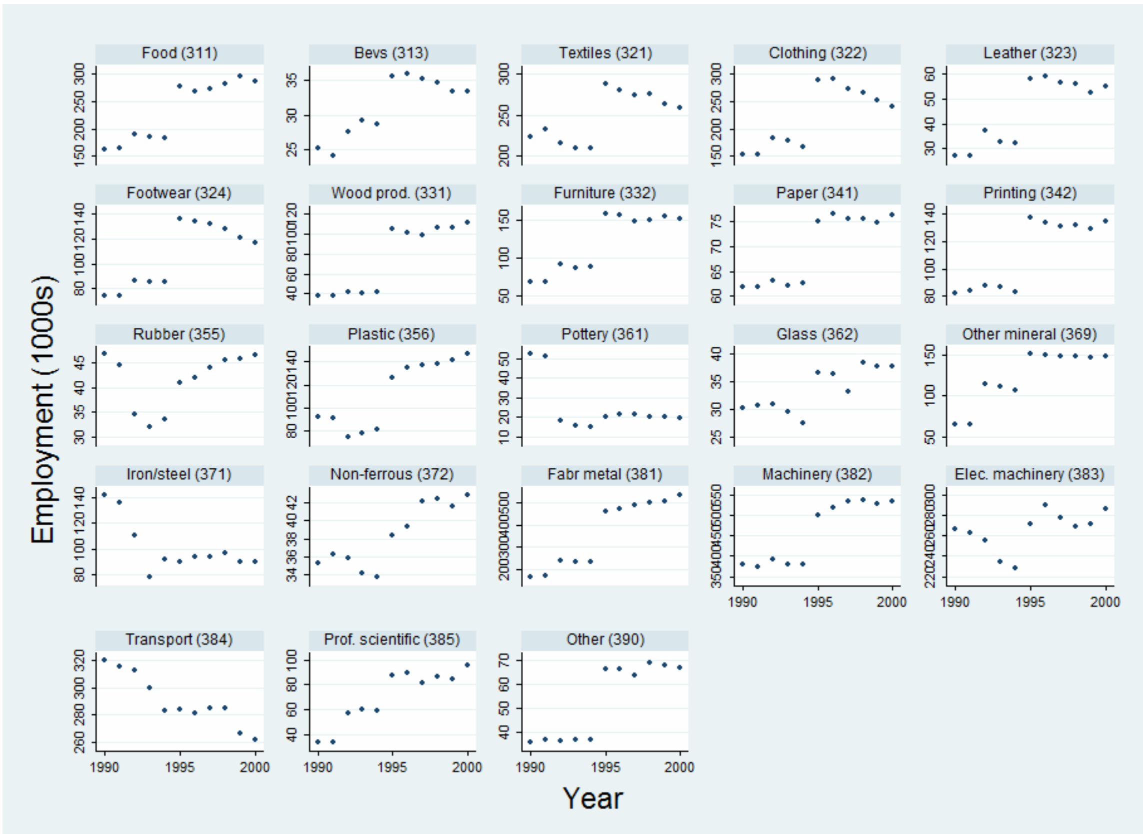

Figure 2.11: Italian Employment, By Industry

Notes: Data from INDSTAT3.

2.10 Appendix

Skilled-unskilled labor to total employment

The GIS distinguishes between skilled and unskilled workers, allowing us to estimate the bias inherent in approximating a three-factor production function using only two factors.

Ignore the log-linearity of the production function and suppose that $e = s + l$. The TFP specification from (16), with fixed effects suppressed for simplicity, is:

$$\begin{aligned}\Delta TFP_{ic} &= \gamma_L l_{ic} + \gamma_S s_{ic} + \gamma_K k_{ic} + \Delta u_{ic} \\ &= \gamma_L l_{ic} + \gamma_S (e_{ic} - l_{ic}) + \gamma_K k_{ic} + \Delta u_{ic} \\ &= (\gamma_L - \gamma_S) l_{ic} + \gamma_S e_{ic} + \gamma_K k_{ic} + \Delta u_{ic}.\end{aligned}$$

The two-factor regression in (18), which only uses total employment and capital, omits the variable l_{ic} from the above equation. A textbook omitted variable bias calculation yields:[73]

$$\begin{aligned}\gamma_E &= \gamma_S + (\gamma_L - \gamma_S)\frac{\text{cov}(l_{ic}, e_{ic})}{\text{var}(e_{ic})} \\ &= \gamma_s \left[1 - \frac{\text{cov}(l_{ic}, e_{ic})}{\text{var}(e_{ic})}\right] + \gamma_L \frac{\text{cov}(l_{ic}, e_{ic})}{\text{var}(e_{ic})}.\end{aligned}$$

This last equation makes it clear that under the two-factor TFP specification, the coefficient on employment is a weighted average of the two labor coefficients from the original three-factor regression.

The results from the GIS (γ_l = -2.15, γ_s = 0.69, γ_E = -1.59) imply that the weight on unskilled labor, L, is 0.803. If the factors that determine employment

---

[73] All covariance and variance terms are understood to have fixed effects and capital appropriately "partialled out."

variance were the same between the 1980s and 1990s, the γ_E estimated using the INDSTAT data is actually attenuated. The two-factor estimate of labor-saving underestimates the true rate of factor-bias.

Data discrepancies in INDSTAT3

Several of the countries in INDSTAT3 have large discontinuities in employment and, to a lesser extent, value added, in the mid-1990s (see Figure 2.9 and Figure 2.10 respectively). While the discontinuities in value added are hardly noticeable, the same cannot be said for employment. The magnitudes of some of jumps in Figure 2.9 are too large to be true variation in the number of manufacturing workers. For example, Italian manufacturing employment increased by about 1.5 million workers between 1994 and 1995, a 60% increase in a single year. Furthermore, these employment discontinuities are not industry-specific, as Figure 2.11 illustrates. All Italian manufacturing sectors exhibit simultaneous increases in the number of workers, which is more plausibly explained by a change in statistical or reporting practices than an unexpected growth in Italian manufacturing across all sectors.

We take this as evidence of a methodological or reporting change that occurred during this period but was undocumented in INDSTAT3. These discontinuities are especially troublesome since our regressions are estimated in first-differences, and the coefficients are identified by the magnitude of those jumps. Because they vary both across countries and industries, these inconsistencies will not be absorbed by any of the included fixed effects, probably attenuating our estimated production function coefficients (β's).

Calculation of Capital Stock

Estimation also requires a measure of capital stock, constructed from gross fixed capital formation[74] as a sum of discounted lagged investments (Berman & Machin, 2000). The capital stock for an industry at year t with data on T lagged investments available is:

$$\hat{K}_t^T = b^T \sum_{\tau=1}^{T} (1-\delta) I_{t-\tau} + c^T I_{t-T} \ ,$$

where b^T and c^T employ superscripts rather than exponents. The coefficients b^T and c^T for each T, along with the depreciation factor δ, were estimated from the US Annual Survey of Manufactures (Gray & Bartelsman, 1995). The minimum lag length used was 8; the maximum was 23.[75] Nominal capital flows are discounted using the investment price indices from the appropriate Penn World Tables.

INDSTAT4 does not include data before the 1990s, so the capital stock calculation for this panel requires merging the historical capital formation values from INDSTAT3 in order to include a sufficient number of lagged investments.

[74] UN documentation defines "gross fixed capital formation" as "the value of purchases and own-account construction of fixed assets during the reference year less the value of corresponding sales."
[75] All lags larger than 23 used the coefficients corresponding to T = 23, namely b^{23} and c^{23}.

2.11 References

Acemoglu, Daron & Fabrizio Zilibotti. (2001). Productivity Differences. *Quarterly Journal of Economics,* 116(2), 563-606.

Autor, David, Lawrence F. Katz, & Alan Krueger. (1998). Computing Inequality: Have Computers Changed the Labor Market? *Quarterly Journal of Economics*, 113(4), 1169-1214.

Bartelsman, Eric & Wayne Gray. (1994). National Bureau of Economic Research Manufacturing Productivity Database. National Bureau of Economic Research, mimeo.

Barro, Robert. (1991). Economic Growth in a Cross-Section of Countries. *Quarterly Journal of Economics,* 106(2), 407-444.

Barro, Robert & Jong-Wha Lee. (1997). International Measures of Schooling Years and Schooling Quality. ftp://www.nber.org/pub/barro.lee.

Barro, Robert & Xavier Sala-I-Martin. (1995). *Economic Growth*, New York, McGraw Hill.

Basu, Susanto & David N. Weil. (1998). Appropriate Technology and Growth. *Quarterly Journal of Economics*, 113(4), 1025-1054.

Baumol, William J. (1986). Productivity Growth, Convergence and Welfare: What the Long-run Data Show. *American Economic Review*, 76(5), 1072-1085.

Berman, Eli, John Bound & Zvi Griliches. (1994). Changes in the demand for skilled labor within U.S. manufacturing industries: Evidence from the Annual Survey of Manufacturing. *Quarterly Journal of Economics*, 109(2), 367-397.

Berman, Eli, John Bound & Stephen Machin. (1998). Implications of Skill-Biased Technological Change: International Evidence. *Quarterly Journal of Economics*, 113(4), 1245-1280.

Berman, Eli & Stephen Machin. (2000). Skill-biased technology transfer around the world. *Oxford Review of Economic Policy*, 16(3), 12-22.

Bound, John & George Johnson. (1992). Changes in the structure of Wages During the 1980s: An Evaluation of Alternative Explanations. *American Economic Review*, 82(3), 371-392.

Bresnahan, Timothy F. & Manuel Trajtenberg. (1997). General Purpose Technologies: ‘Engines of Growth’? *Journal of Econometrics,* 65(1), 1091-1126.

Caselli, Francesco & Wilbur John Coleman II. (2006). The World Technology Frontier. *American Economic Review*, 96(3), 499-522.

Christensen, Laurits R., Dale W. Jorgenson and Lawrence J. Lau. (1973). Transcendental Logarithmic Production Frontiers. *Review of Economics and Statistics*, 55(1), 28-45.

Ciccone, Antonio & Elias Papaioannou. (2009). Human Capital, the Structure of Production, and Growth. *Review of Economics and Statistics*, 91(1), 66-82.

Conte, Andrea & Macro Vivarelli. (2011). Imported Skill-Biased Technological Change in Developing Countries. *The Developing Economies*, 49(1), 36-65.

Doms, Mark, Timothy Dunne, & Kenneth R. Troske. (1997). Workers, Wages and Technology. *Quarterly Journal of Economics*, 112(1), 253-290.

Eaton, Jonathan and Samuel Kortum. (1996). Trade in Ideas: Patenting and Productivity in the OECD. *Journal of International Economics,* 40(3-4), 251-278.

Eaton, Jonathan and Samuel Kortum. (1999). International Technology Diffusion: Theory and Measurement. *International Economic Review*, 40(3), 537-570.

Feliciano, Zadia M. (2001). Workers and trade liberalization: The impact of trade reforms in Mexico on wages and employment. *Industrial and Labor Relations Review,* 55(1), 95-115.

Feenstra, Robert C. & Gordon Hanson. (1996). Globalization, Outsourcing, and Wage Inequality. *American Economic Review Paper and Proceedings*, 86(2), 240-245.

Freeman, Richard & Lawrence Katz. (1994). Rising wage inequality: The United States vs. other advanced countries. In Freeman (eds.) *Working Under Different Rules*. New York: Russell Sage Foundation.

Goldin, Claudia & Lawrence F. Katz. (1996). Technology, Skill and the Wage Structure: Insights from the Past. *American Economic Review,* 86(2), 252-257.

Goldin, Claudia & Lawrence F. Katz. (1998). The Origins of Technology-Skill Complementarity. *Quarterly Journal of Economics,* 113(3), 693-732.

Gottschalk, Peter & Mary Joyce. (1998). Cross-National Differences in the Rise in Earnings Inequality: Market and Institutional Factors. *Review of Economics and Statistics*, 80(4), 489-502.

Griliches, Zvi. (1969). Capital-Skill Complementarity. *Review of Economics and Statistics*, 51(4), 465-468.

Griliches, Zvi & Dale W. Jorgenson. (1967). The Explanation of Productivity Change. *Review of Economic Studies,* 34(3), 249-283.

Griliches, Zvi & Jacques Mairesse. (1995). Productivity Functions: The Search for Identification. National Bureau of Economic Research Working Paper 5067, March.

Griffith, Rachel, Stephen Redding, & John Van Reenen. (2004). Mapping the Two Faces of R&D: Productivity Growth in a Panel of OECD Industries. *Review of Economics and Statistics*, 86(4), 883-895.

Hanson, Gordon H. & Ann Harrison. (1995). Trade, Technology and Wage Inequality. National Bureau of Economic Research Working Paper 5110.

Haskel, Jonathan E. & Matthew J. Slaughter. (1998). Does the Sector Bias of Skill-Biased Technological Change Explain Changing Wage Inequality? National Bureau of Economic Research Working Paper 6565.

Helpman, Elhanan. (1998). Introduction. In Helpman (ed.) *General Purpose Technologies and Economic Growth.* Cambridge, MA: MIT Press.

Helpman, Elhanan & Gene M. Grossman. (1991). *Innovation and Growth in the Global Economy*. Cambridge, MA: MIT Press.

Hendricks, Lutz. (2010). Cross-country variation in education attainment: structural change or within-industry skill upgrading? *Journal of Economic Growth*, 15, 205-233.

Heston, Alan, Robert Summers, & Bettina Aten. (2006). Penn World Table Version 6.2, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania, September.

Jorgenson, Dale & Marcel P. Timmer. (2011). Structural Change in Advanced Nations: A New Set of Stylised Facts. *Scandinavian Journal of Economics*, 113(1), 1-29.

Kahn, James A. & Jong-Soo Lim. (1998). Skilled-Labor Augmenting Technical Progress in U.S. Manufacturing. *Quarterly Journal of Economics*, 113(4), 1281-1308.

Katz, Lawrence F., Gary W. Loveman & David G. Blanchflower. (1995). A comparison of Changes in the Structure of Wages in Four OECD Countries. In

Katz & Freeman (eds). *Differences and Changes in Wage Structures.* Chicago: University of Chicago Press.

Katz, Lawrence F. & Kevin M. Murphy. (1992) Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *Quarterly Journal of Economics*, 107(1), 35-78.

Klenow, Peter. (1993). Ideas vs. Rival Human Capital: Industry Evidence on Growth Models. *Journal of Monetary Economics,* 42(1), 3-24.

Kremer, Michael. (1993). Population Growth and Technological Change: One Million B.C. to 1990. *Quarterly Journal of Economics,* 108(3), 681-716.

Lawrence, Robert & Matthew Slaughter. (1993). International Trade and U.S. wages in the 1980s: Giant sucking sound or small hiccup. *Brookings Papers on Economic Activity*, vol. 2, 161-227.

Machin, Stephen, & John Van Reenen. (1998). Technology and changes in skill structure: Evidence from Seven OECD Countries. *Quarterly Journal of Economics*, 113(4), 1215-1244.

Mankiw, N. Gregory, David Romer, & David N. Weil. (1992). A Contribution to the Empirics of Economic Growth. *Quarterly Journal of Economics*, 107(2), 407-437.

Murphy, Kevin M. & Finis Welch. (1992). The Structure of Wages. *Quarterly Journal of Economics*, 107(1), 285-326.

Nelson, Richard & Edmund Phelps. (1966). Investment in Humans, Technological Diffusion and Economic Growth. *American Economic Review Papers and Proceedings,* 56(1/2), 69-75.

Psacharolpoulos, George. (1994). Returns to Investment in Education: A Global Update. *World Development,* 22(9), 1325-1343.

Psacharolpoulos, George. (1999). The Opportunity Cost of Child Labor: A Review of the Benefits of Education. University of Athens mimeo.

Ray, Debraj. (1998). *Development Economics.* Princeton, NJ: Princeton University Press.

Robbins, Donald J. (1995). Trade, Trade Liberalization and Inequality in Latin America and East Asia- Synthesis of Seven Country Studies. Harvard University mimeo.

Solow, Robert. (1956). A Contribution to the Theory of Economic Growth. *Quarterly Journal of Economics,* 70(1), 65-94.

Solow, Robert. (1957). Technical Change and the Aggregate Production Function. *Review of Economics and Statistics*, 39(3), 312-320.

Schumacher E.F. (1973). *Small is Beautiful: Economics as if People Mattered.* New York: Harper and Row.

Summers, Robert & Alan Heston. (1991). Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950-1988. *Quarterly Journal of Economics*, 106(2), 327-368.

United Nations, Dept. of Economic and Social Affairs, Statistical Office, *Industrial Statistics Yearbook, Volume I: General Industrial Statistics*, (New York, NY: 1992).

United Nations Industrial Development Organization. (2006). INDSTAT3 2006 ISIC Rev.2.

United Nations Industrial Development Organization. (2008). INDSTAT4 2008 ISIC Rev.3.

Weil, David N. (2000). *Economic Growth*, Brown University manuscript.

Young, Alwyn. (1995). The Tyranny of Numbers: Confronting the Statistical Realities of the East Asian Growth Experience. *Quarterly Journal of Economics,* 110(3), 641-680.

Chapter 3: Does Development Assistance Reduce Violence? Evidence from Afghanistan

Abstract

Current military doctrine emphasizes the importance of development spending in reducing insurgent violence. We ask whether development aid in Afghanistan is violence-reducing. We use data from three distinct development programs, the Afghan National Solidarity Program, USAID's Local Governance and Community Development Program and the US military's Commander's Emergency Response Program (CERP), combined with military records of insurgent-initiated events. Overall spending has no clear effect on rebel attacks. Moreover, the types of development program most effective at reducing violence in Iraq – small CERP projects—does not appear to do so in Afghanistan. We speculate as to why, considering troop strength, conditionality of aid, effectiveness of aid in producing benign outcomes, and measurement issues. Policymakers might re-evaluate development spending in Afghanistan.

3.1 Introduction

Current counterinsurgency (COIN) doctrine emphasizes the role of benign development assistance as a key component in any campaign to enhance security in conflicted and post-conflict regions.[76] As a consequence, significant resources have been spent on rebuilding Afghanistan's institutions and livelihoods with the intention that such projects achieve both conventional development goals[77] and donors' security objectives. While counterinsurgency is almost as old as war itself, there has been relatively little empirical research into whether these reconstruction efforts have generated security improvements as intended. Using unique data on insurgent attacks and three reconstruction programs in Afghanistan, this paper examines whether this development spending decreases insurgent violence.

The questions of when, where, and how development assistance builds stability are especially relevant to policy-makers as the military intervention in Afghanistan enters its tenth year and international donors begin to shift their attention to other conflicted areas such as the Middle East and Africa. The "hearts and minds" theory underlying current counterinsurgency doctrine is quite intuitive: in a conflict between the government and rebel forces, the local population has actionable information on insurgent activities, which it can either share with the government and other allied

---

[76] The COIN Field Manual explicitly states that "Durable policy success requires balancing the measured use of force with an emphasis on nonmilitary programs… COIN programs for political, social, and economic well-being are essential to developing the local capacity that commands popular support when accurately perceived." (US Army, 2006, Section 2-5).

[77] Having experienced thirty years of continuous conflict, Afghanistan ranks as one of the poorest countries in the world. In 2009, GDP per capita was estimated to be $486. For comparison, neighboring Pakistan's GDP was nearly twice as high ($955). Even in the absence of national security interests, Afghanistan would be a prime candidate for conventional development aid.

forces or not (Berman, Shapiro, & Felter, 2011). In the Afghan context, this could be villagers witnessing insurgents planting a roadside bomb, knowledge of which they can either share with the local military commander or keep to themselves. Naturally, the insurgents would prefer that the community keep quiet since non-cooperation results in a successful attack that damages government or allied forces. The key insight from the theoretical model is that the government can induce information sharing by providing basic necessities or other goods and services.[78] Other interpretations of "hearts and minds" theory argue that noncombatants are influenced not by improved governance but by grievances allayed, jobs provided, or because their leaders are co-opted, and that the consequential act of noncombatants is not information sharing but active resistance to rebel activity, taxation or recruitment. Nevertheless, all these models share the implication that development spending reduces violence.

Though the current strategy of combining military operations with civilian development has been somewhat successful in Iraq, the results here suggest that development efforts in Afghanistan have ambiguous effects on conflict. Using detailed project-level data from three separate reconstruction programs (NSP, LGCD, and CERP)[79] and incident-level military reports on insurgent attacks, we find that

---

[78] As we argue in Section 3.3, the government does not need to be a social welfare optimizer for this implication to hold. In the model, service provision is purely instrumental and it is still in the government's interest to provide some services to incentivize cooperation by the population.

[79] NSP is the Afghan government's National Solidarity Program. USAID operates the Local Governance and Community Development (LGCD) program while CERP is the US military's Commander's Emergency Response Program. Section 3.4 discusses these in more detail.

development spending does not appear to reduce (or increase) the level of rebel violence once district fixed effects are controlled for.[80]

The difference in results between CERP and the two other programs suggests that aid "conditionality" is an essential, but currently under-emphasized, prerequisite for stability-enhancing development. The model predicts stark differences in effectiveness between aid that is contingent on community cooperation ("conditional aid") and aid that isn't: only conditional aid reduces insurgent violence. Development projects provided independent of information sharing have no effect on violence because they do not make the community more likely to share information on the margin. Out of the three programs examined here, only CERP practices conditionality and hence is the only one predicted to have violence-reducing potential. Our empirical results are consistent with this conjecture as CERP is the only one to have consistently negative, if poorly estimated, effects on rebel violence.

While overall spending does not appear to be "winning hearts and minds," there is some heterogeneity across different types of spending. In particular, we find preliminary evidence that small-scale CERP projects might be more effective at reducing violence than larger ones. This finding is consistent with the theoretical prediction that projects or places where the government is more effective at providing services should exhibit stronger violence reduction. However, these estimates are imprecise and only small-spending delivered through the US military demonstrates

[80] Districts are the next level of administrative division after provinces, analogous to a county in the US. As of 2005, Afghanistan has 398 districts spread across 34 provinces.

this effect; small-scale development through USAID's LGCD program does not appear to have much effect on rebel activity.

The findings discussed here have important implications for both academics and aid practitioners. From a practical standpoint, it is not obvious that the resources currently being devoted to reconstruction in Afghanistan are having any stability-enhancing effect. Future research on reconstruction and stability should closely examine (1) aid conditionality, (2) the effectiveness of development programs in providing services –including the importance of the government's institutional capacity to adeptly provide services. Future efforts to rebuild contested and post-conflict areas should not necessarily focus on spending more money, but rather on using it more effectively.

The next section summarizes current theories on the relationship between service provision, governance, and insurgent violence. Section 3.3 outlines a simple model of counterinsurgency with an emphasis on two empirically testable hypotheses regarding the relationship between development spending and insurgent violence. Section 3.4 discusses both the military records on violence and the institutional details of the three different development programs used here. Section 3.5 presents the main empirical findings and Section 3.6 concludes.

3.2 Insurgency and Development Aid as COIN

What distinguishes an insurgency from a traditional inter-state military conflict? Both are contests between armed parties in pursuit of political power, but unlike conventional armed conflict, insurgencies emphasize the pivotal role of noncombatants (US Army, 2006, Section 1-3). Civilians, who are generally unarmed and may not even share political ideologies with the rebels, are crucial to the success of insurgent campaigns because they can provide actionable information that makes military operations more effective (Kalyvas, 2006). Rather than being merely passive observers of a conflict, the populace is an active player in insurgencies, one that responds swiftly to both state and rebel actions (Galula, 1964; Popkin, 1979).

While counterinsurgency researchers and practitioners appear to agree on the importance of popular support in determining the outcome of insurgent conflicts, the question of *how* to gain it is still actively debated. “Hearts and minds” proponents argue that the government can win civilian support by addressing grievances, thus reducing the “demand” for rebellion (Gurr, 1970; Horowitz, 1985). Others argue that rebels, like secular criminals (Becker, 1968), might be more sensitive to the opportunity costs and potential payoffs of rebellion (Grossman, 1991; Fearon, 2008). This would be especially true in weakly governed places where the state cannot successfully “buy off” potential rebels, either through legitimate work opportunities or other income transfers, nor can they effectively utilize a cooperative populace’s information.

The empirical evidence on the relative importance of grievances (“demand”) compared to employment/income-generation (“supply”) as motivations for insurgent violence has been somewhat mixed. Fearon and Laitin (2003) find that civil war is predicted by low income per capita and difficult terrain, both of which are indicative of poor state capacity and low opportunity cost of rebellion. However, Berman, Felter, Shapiro, and Callen (2011) show that employment rates are actually positively correlated with insurgent violence in Iraq, the Philippines, or Afghanistan. On the other hand, rebel attacks seem to increase after coalition-induced civilian casualties, suggesting that the "supply" of insurgent activity in Afghanistan is somewhat responsive to government actions (Condra, Felter, Iyengar, & Shapiro, 2010).

Theoretically, reconstruction and service provision by the state signals competent and committed governance, and should be effective at inducing information-sharing and improving security. This appears to be the case in Iraq: development projects channeled through the US military's CERP were effective at reducing rebel violence (Berman, Shapiro, & Felter, 2011). That is also the conclusion of Beath et al. (2010) for the experimental subsample of the NSP, a paper we discuss further below. However, corruption and poor governance are common complaints among Afghans, and these weaknesses can dampen or even reverse the effect of reconstruction on stability should they provide more rents for rebels to capture (Wilder, 2009; Crost, Felter & Johnston, 2010) or signal incompetent or ambivalent governance (Rashid, 2008).

3.3 Theoretical Background

Brief Description of the Model

The model developed in this section follows Berman, Shapiro, and Felter (2011). The three active players in the game, Government, Rebels, and the

Community, play a four-stage game. Initially, Nature determines a preference parameter for the Community ("norms") which is revealed only to the Community. In Stage 2, Government and Rebels simultaneously chose their actions (detailed below). Having observed both *G* and *R*'s actions, the Community chooses its level of information-sharing. Finally, Nature resolves the uncertainty about whether *G* or *R* has "control" at the end of the game and payoffs for all three players are realized.

The remainder of this section formally derives the two main testable implications examined in the empirical estimation. The first hypothesis is that a regression of violence on reconstruction spending will yield a negative coefficient when controlling for rebel strength, community norms, and other local characteristics (Hypothesis H_1). That is to say, development aid is violence-reducing. Second, the violence-reducing impact of reconstruction spending will be greater when government forces have better knowledge of local community needs (Hypothesis H_2).

Players, Actions, and Payoffs

There are three players in the game, denoted *G*, *R*, and *C*. The key state of the game that determines payoffs is whether *G* or *R* has control at the end, denoted by the binary variable *a* where $a = 1$ if *G* is in control, $a = 0$ if *R* is in control. The Community has political norms regarding rebel control, *n*, which are conceptualized as a utility penalty if *G* is in control at the end of the game.

The Community's sole action is to choose a level of information-sharing $i \in [0,1]$. The Rebels also only have one action, to choose a level of violence $v \geq 0$, which targets the Government but still negatively affects the Community. The Government has two actions: it can combine benign social services, $g \geq 0$, with active operations to mitigate violence, $m \geq 0$.

The Community's payoffs are as follows: if $a = 1$ and the Government has control, it receives secular consumption *c* and government-provided services *g* but also experiences a penalty for having shared information *n*; if $a = 0$ and the Rebels have control, it still receives consumption *c* but also suffers from violence *v*. The payoff function for *C* is:

$$U_C(c, g, v, n, a) = a \cdot u(c + g - n) + [1 - a] \cdot u(c - v),$$

where $u(\cdot)$ is a well-behaved[81] utility function. Note that a key assumption is that *g* is "conditional": the Government can and will only provide services if it is in control at the end of the game. This is a rather unconventional assumption and will be discussed later in this section.

The Rebels' goal is to impose costs on government, presumably to extort concessions. Violent actions benefit Rebels according to the function $A(v)$ but only if they are in control at the end of the game. Violence costs Rebels $B(v)$ regardless of the ending state. The payoff function for *R* is:

$$U_R(v, a) = [1 - a] \cdot A(v) - B(v),$$

where $A(\cdot)$ and $B(\cdot)$ are both C^2 and increasing. $A(\cdot)$ is concave while $B(\cdot)$ is convex. Assume that no violence results in no damage: $A(0) = 0$.

[81] $u(\cdot)$ is twice continuously differentiable (C^2) and monotonically increasing.

Both the Community and Rebels are expected utility maximizers. The Government seeks to minimize a combination of violence and costs. If *R* has control at the end of the game, *G* suffers damage $A(v)$, otherwise it is unharmed by rebel violence. Both violence mitigation *m* and service provision *g* incur costs, defined by $D(m)$ and $H(g)$ respectively, regardless of which player is in control at the end. The Government's total cost function is:

$$C_G(v,m,g,a) = [1-a]\cdot A(v) + D(m) + H(g)\,.$$

Cost functions $D(\cdot)$ and $H(\cdot)$ are C^2, increasing and convex, and scaled such that $D(0) = H(0) = 0$. To rule out the case where mitigation is never effective, assume that $A(n_U) > D'(0)$. Intuitively, this condition says that even in the "worst case scenario" (i.e. areas with the highest proclivity toward violence), it costs less to provide a tiny amount of counterinsurgency effort than it does to suffer full damage from Rebel violence. Hence, it is always in the Government's interest to provide nonzero counterinsurgency effort.

The final component of the model is how *G* converts mitigation *m* and information *i* into control. Let *p* denote the probability that $a = 1$. *G* can combine mitigation and information to increase its probability of winning control according to:

$$p = \Pr(a = 1) = h(m)\cdot E(i)\,,$$

where $h(m): \mathrm{R}^+ \to [0,1]$ is a "contest success function" (Skaperdas, 1996). Higher COIN effort *m* increases the probability that *G* is in control, but this mitigation also faces decreasing returns; $h(m)$ is increasing but concave. $h(0) = 0$ and $h(m) \to 1$ as

$m \to \infty$. Note that information sharing is necessary but not sufficient for control: if $i = 0$, then $p = 0$, but $i = 1$ does not guarantee that $p = 1$.

Description of the Game

The game has four stages but strategic interaction only occurs in Stage 2 and Stage 3. In Stage 1, Nature draws norms $n \sim U[n_L, n_U]$, and this parameter is revealed only to *C*. The support of n is assumed to be wide enough that neither *G* nor *R* can fully determine the outcome of the game through his actions alone.[82] In Stage 2, *G* and *R* simultaneously move. In Stage 3, *C* observes the actions of the previous stage $\{v, m, g\}$ and chooses its level of information sharing. Finally, Nature draws the final state $a \sim bernoulli(p(m,i))$ and payoffs to *G*, *R*, and *C* are determined.

Equilibrium

Solve for the subgame perfect Nash equilibrium via backward induction. The Community's objective is to choose *i* to maximize:

$$\begin{aligned} EU_C(c, g, v, n, a) &= E(a) \cdot u(c + g - n) + [1 - E(a)] \cdot u(c - v) \\ &= p \cdot u(c + g - n) + [1 - p] \cdot u(c - v) \\ &= h(m) \cdot i \cdot u(c + g - n) + [1 - h(m) \cdot i] \cdot u(c - v). \end{aligned}$$

[82] More specifically, $n_L \leq v + g \leq n_U$.

Since this function is linear in *i*, the only solutions are on the boundaries.[83] *C* will choose to share information if $u(c+g-n) > u(c-v)$; otherwise, it will not share at all. Since $u(\cdot)$ is monotonically increasing, this implies that *C*'s best response is:

$$i^* = \begin{cases} 1 \text{ if } g-n > -v \\ 0 \text{ if } g-n < -v. \end{cases}$$

Given the distributional assumption about *n*, this implies that:

$$\begin{aligned}\Pr(i^* = 1) &= \Pr(n < g+v) \\ &= \frac{g+v-n_L}{n_U - n_L} \\ &= f \cdot (g+v-n_L) \text{ where } f = 1/(n_U - n_L).\end{aligned}$$

Plugging this back into the definition of *p* results in:

$$\begin{aligned}p^*(m,g,v) &= h(m)\cdot i^* \\ &= h(m)\cdot f \cdot (g+v-n_L).\end{aligned} \tag{19}$$

Turning to the previous stage of the game, *G* and *R* will simultaneously optimize, knowing that *C*'s actions will result in the final state $a=1$ with probability p^* defined by Eq. (19). *R*'s problem is simply to choose violence to maximize:

$$\begin{aligned}EU_R(v,a) &= [1-E(a)]\cdot A(v) - B(v) \\ &= [1-p^*(m,g,v)]\cdot A(v) - B(v).\end{aligned}$$

The first-order condition for *v* is:

$$\begin{aligned}\frac{\partial}{\partial v}EU_R &= [1-p^*]A'(v) - A(v)\frac{\partial}{\partial v}p^* - B'(v) \\ &= [1-p^*]A'(v) - A(v)h(m)f - B'(v) = 0,\end{aligned}$$

[83] Trivial solutions occur in the case where $h(m) = 0$ or $g - n = -v$. In either case, any value of *i* is optimal. Since $m = 0$ is never optimal and $h(m)$ is increasing, there are no other values of *m* that might yield Case 1.

which results in a best-response function $v^*(m, g)$. Differentiating implies that v^* is decreasing in both its arguments (see Appendix). Holding *m* constant, Rebels respond to increased service provision with lower violence. Similarly, Rebels respond to higher COIN effort by lowering violence, holding *g* constant.

G's problem is to choose both *g* and *m* to minimize:

$$\begin{aligned} EC_G(v,m,g,a) &= [1-a]\cdot A(v) + D(m) + H(g) \\ &= [1-p^*(m,g,v)]\cdot A(v) + D(m) + H(g). \end{aligned}$$

The first-order condition with respect to *m* is:

$$\frac{\partial}{\partial m} EU_G = -A(v)\frac{\partial}{\partial m} p^* + D'(m) = 0.$$

The first-order condition with respect to *g* is:

$$\frac{\partial}{\partial g} EU_G = -A(v)\frac{\partial}{\partial g} p^* + H'(g) = 0.$$

Solving the first-order conditions provide best response functions $m^*(g,v)$ and $g^*(m,v)$. Differentiating implies that both COIN effort and service provision are increasing in *v* (see Appendix). Furthermore, for a given level of Rebel activity, mitigation and services are complements:

$$\begin{aligned} \left.\frac{\partial m^*}{\partial g}\right|_v &= -\left[\frac{\partial^2 EU_G}{\partial m^2}\right]^{-1}\left[\frac{\partial^2 EU_G}{\partial g \partial m}\right] \\ &= -[(+)]^{-1}\left[-A(v)\frac{\partial^2 p^*}{\partial g \partial m}\right] \\ &= [(+)]^{-1}\left[A(v)h'(m)f\right] \\ &> 0. \end{aligned}$$

The subgame perfect Nash equilibrium (SPNE) is defined by the best response functions $m^*(g,v)$, $g^*(m,v)$, $v^*(m,g)$, and i^* derived above.

Testable Implications

The first testable hypothesis (H_1) is that $\partial v^* / \partial g \,|_m$ is negative: holding local characteristics and Government counterinsurgency effort constant, an increase in government spending reduces violence.

The second hypothesis concerns the relative effectiveness of particular types of service provision. Note that the Community's utility function implicitly assumes that the marginal utility of services is unity. To derive the second testable implication, we relax this assumption and allow it to have its own coefficient β_g so the Community's utility is now:

$$U_C(c,g,v,n,a) = a \cdot u(c + \beta_g g - n) + [1-a] \cdot u(c - v).$$

Then *C*'s best response becomes:

$$i^* = \begin{cases} 1 \text{ if } n < \beta_g g + v \\ 0 \text{ if } n > \beta_g g + v \end{cases},$$

implying that $p^* = h(m) \cdot f \cdot (\beta_g g + v - n_L)$. Then *R*'s optimal response to a change in government services is:

$$\left.\frac{\partial v^*}{\partial g}\right|_m = \frac{\beta_g A'(v) fh(m)}{A''(v)[1-p^*] - 2A'(v) fh(m) - B''(v)} < 0$$

When β_g is high, p^* is high and $1-p^*$ is small. Since $1-p^*$ appears in the denominator, this derivative gets more negative (i.e. larger in magnitude) as β_g gets larger. Hence, services that provide higher marginal utility to the Community have stronger violence-reducing effects (H_2). We interpret this empirically as saying that small projects, which are quicker to implement and more adaptable to community needs, should exhibit stronger violence-reducing effects than large ones. In other words, the coefficient on small spending should be more negative.

Necessary Condition: Conditionality of Aid

Recall that *C* only benefited from the Government's service provision if *G* was in control at the end of the game; if the Rebels are in control, *g* does not appear in *C*'s payoff. Since information is necessary but not sufficient for *G* to have control, service provision is actually "conditional" on information-sharing by the community. At first glance, this seems to be a rather extreme assumption since it cannot be true of certain projects (e.g. infrastructure). However, aid conditionality is a necessary condition for *g* to be violence-reducing in the model. Intuitively, unconditional service provision does not affect the Community's behavior on the margin since *C* benefits from it in both states of the world and it cancel out in the optimality condition for i^*.

More formally, suppose that overall service provision *g* is actually divided up into conditional services, g_c, and unconditional services, g_u. Then *C*'s expected utility becomes:

$$EU_C(c,g,v,n,a) = p \cdot u(c + g_c + g_u - n) + [1 - p] \cdot u(c + g_u - v),$$

and optimal behavior is still determined by the expected tradeoff between utility in the two states of the world. As before, C will share information if and only if the payoff from doing so is greater than not. By monotonicity of $u(\cdot)$:

$$\begin{aligned} u(c + g_c + g_u - n) &> u(c + g_u - v) \\ c + g_c + g_u - n &> c + g_u - v \\ c + g_c - n &> c - v \end{aligned} \quad .$$

Hence unconditional service provision has no effect on information sharing. Since spending by traditional development agencies is not conditional on cooperation, the model predicts stark differences in the effectiveness between the military's CERP and the other, unconditionally provided, aid programs. In particular, NSP and LGCD will have no effect on violence while CERP, whose guidelines emphasize conditionality, should be violence-reducing.

While the model's prediction about unconditional spending is quite clear, we consider it a positive rather than normative statement. The Government can still provide g_u to increase welfare, if not to induce information sharing. In practice, some reconstruction projects, like paving roads or building power plants, provide logistical benefits to the government in addition to improving service provision to locals.

3.4 Data

To test the empirical hypotheses derived in Section 3.3, we use data on insurgent violence, in addition to project records for three distinct development programs: NSP, LGCD, and CERP. Since more populous districts are likely to have more insurgent attacks and receive more development assistance, both violence and spending will be scaled by district population.[84]

Insurgent Violence: CIDNE

To form a measure of insurgent activity, we use declassified incident records from the US military's Combined Information Data Network Exchange (CIDNE) database. Our records from CIDNE consist of 60,075 events of "significant activity" (SIGACT) from April 2002 through January 2010. Each event record comes with date, time, attack type, and geographic coordinates. The fields provided allow us to precisely geo-locate each incident and create a detailed district-month panel of insurgent activity.

A few limitations of our violence data are worth discussing. First, to qualify as a SIGACT, an event must be insurgent-initiated; events initiated by coalition or Afghan forces are not included. To the extent that rebels attack civilians or conduct criminal activity, our violence measure will undercount true violence. As the theoretical model is framed as rebels attacking the government, we consider SIGACTs an appropriate measure of insurgent activity to test the model's predictions. Second, SIGACTs can vary in scale and complexity, ranging from direct fire incidents to improvised explosive devices (IEDs), and we do not have information about the damage caused or units involved in such attacks.[85] In addition, individual military

[84] Cross-sectional district population estimates are extrapolated from Landscan population densities and generously shared by Nils B. Weidmann.
[85] In Afghanistan, the vast majority of events are either "direct fire" incidents or IEDs. Since 2005, 44% of the total SIGACTs in our data are direct fire while IEDs constitute another 34% of the observations.

units are likely to differ in their reporting thresholds of what constitutes a "significant" event. Since insurgent violence appears on the left-hand side of our regressions, classical measurement error in SIGACTs should only result in larger standard errors but should not bias estimates.

A more salient concern with SIGACTs is the interpretation of zero-violence districts. Since our district-month panel is constructed using recorded events, a district that appears to have no events could either have no attacks at all (a "true" zero) or have no military personnel around to report those events. Since we do not have data on the allocation or placement of US forces, we cannot directly control for this omitted variable. One way to address this issue is to condition on a proxy for the location of soldiers. As will be argued later, large-scale projects are likely to require more protection than small ones, so spending or presence of large projects could serve as a proxy for unobserved counterinsurgency effort.

An additional measurement issue somewhat unique to our situation bears mentioning. While CERP and LGCD have the stated function of enhancing "stability" –which is generally understood to mean the security of noncombatants-- our theory and regressions were developed using violence directed against combatants as the outcome. Implicitly, we have assumed that SIGACTs appropriately proxy for district stability or government control, which might not be the case. For example, there could be a nonlinear relationship between our observed outcome, SIGACTs, and unobserved rebel control simply because there are no military targets to attack in insurgent zones

---

Since the model does not provide strong implications about tactical choice by insurgents, we just pool the event types.

of control, nor would there be anyone around to record the incident. As government or coalition forces start to enter these insurgent strongholds, the number of SIGACTs could increase as the rebels are presented with more potential targets.[86] Since the government (and ISAF) sometimes expand into regions where it previously had little control, this could be viewed as a stability improvement even though reported violence is actually increasing.

Figure 3.2 plots observed SIGACTs per capita on the vertical axis against a composite index of district-level security perceptions (x-axis).[87] Along the left side of the curve, we see that stability improvements are correlated with decreasing violence. However, the interpretation is reversed for places on the right-hand side: decreasing violence moves along with *decreasing* security perceptions. This problem is not limited to just a few outlying districts; almost half the districts plotted in the figure are on the right-hand side of the curve. This inverted U-shape implies that the non-monotonic relationship between SIGACTs and unobserved stability might be a realistic problem in evaluating reconstruction, and potentially other policy interventions, in Afghanistan. As an additional check on our results, we split the estimation sample based on the constructed stability index to see if development aid has differential effects based on unobserved district security.

[86] This is a more pernicious problem than that of using crime reports to infer crime rates. Additional police officers might improve reporting, but they are unlikely that they actually *attract* more crime. However, additional military units both improve reporting but also draw the attention of rebels.
[87] Details about the construction of the stability index are provided in the Appendix.

Reconstruction Programs: NSP, LGCD, and CERP

We have detailed project-level data for three different reconstruction programs in Afghanistan. All three programs fund a variety of projects types though project selection is likely to differ based on the incentives of the different stakeholders and involved parties.

The first development program for which we have detailed data is the National Solidarity Program (NSP). Started in 2003, NSP is intended to help individual communities build and manage their own development projects (MRRD, 2007). Logistically, NSP allocates block grants to individual rural areas and aids a Community Development Council (CDC) in identifying and developing projects to use those funds. The election of a CDC is a precondition for receipt of a grant. These block grants are calculated based on the number of households in the community ($200 per family). Grants are capped at $60,000 though this does not appear to be a binding constraint as the average size of grants is well under the maximum ($33,000 per CDC). Our NSP data cover almost $680 million in project expenditures spread across 316 districts.

NSP differs from the other two development programs in a few dimensions. First, it will help villages establish a CDC if one does not currently exist. CDCs were originally intended to aid project implementation, but some evidence suggests that they also provide auxiliary benefits in the form of local governance and dispute resolution (Beath, Christia, Enikolopov, & Kabuli, 2010). Since CDC formation is a prerequisite for project implementation, we cannot disentangle the ancillary benefits of having a CDC, especially if it effectively provides local governance where there was none before, from those of having a project at all. Second, NSP explicitly requires that grant-receiving communities contribute 10% of the total cost in the form of labor, materials, or funds.[88] Finally, NSP is administered by the Afghan government's Ministry of Rural Rehabilitation and Development (MRRD) and explicitly advertised as such. The "Afghan face" of NSP could either make their projects more or less likely to be attacked by rebels. NSP activities might be more attractive targets for rebels since they symbolize the central government's expansion into relatively untouched areas but insurgents may also be hesitant to attack projects where the receiving communities are both involved and personally invested.

The second development program is USAID's LGCD, which seeks to improve Local Governance (LG) and Community Development (CD) in insecure areas (USAID, 2010). LGCD projects are also community-initiated and driven since proposals can be brought up and approved through the local CDC, but they lack the explicit block grant funding scheme of NSP. LGCD funds a wide spectrum of development projects from infrastructure construction to equipment purchases and training programs. In contrast to NSP, LGCD is relatively new with initial projects starting in 2007. While LGCD itself is active in other regions of Afghanistan, our data are limited to just projects in the South and East regions.

The final reconstruction program for which we have data is the US military's Commander's Emergency Response Program (CERP). As its name suggests, CERP is

---

[88] The effect of explicitly imposing some of the cost on beneficiaries has not yet been closely studied since it directly ties each community to their NSP project. Unfortunately, we do not have data on compliance with this rule.

intended to allow commanders to provide “urgent, small-scale, humanitarian relief, and reconstruction projects and services that immediately assist the indigenous population” (US Army, 2009, Ch.4). However, CERP projects do not have an explicit maximum and range in size from small condolence payments to construction of major roads in Afghanistan. Since 2004, CERP has appropriated almost $2.64 billion in Afghanistan (SIGAR, 2011).

While CERP is the longest running of our three reconstruction programs in Afghanistan, we only have district identifiers and project expenditures for Fiscal Year 2010 (October 2009 - September 2010). The full time series of CERP projects only allows us to calculate project *counts* at the district-month level.[89]

Using these data on individual project locations, dates, and costs, we construct a panel of reconstruction expenditures, or project counts in the case of CERP, by uniformly spreading project spending over all days in which each project was active and then aggregating up to the district-month level. This spending measure is our main explanatory variable.[90] Uniformly distributing expenditure over each project day will likely induce measurement error in our calculated spending series since the true timing of project disbursements is likely to be much lumpier. Since this variable shows up on the right-hand side, our spending coefficients will suffer from attenuation bias and be smaller in magnitude than an unbiased estimate.

[89] Originally, the financial and operational records for CERP projects were stored separately and without enough information to link project expenditure to project location. The incompleteness of CERP records before 2009 has only recently been pointed out by the Special Inspector General for Afghanistan Reconstruction (SIGAR) as an area that needs improvement (SIGAR, 2009).

[90] Implicitly, we are using these data on project cost to proxy for government service provision. The institutional environment of Afghanistan might be such that dollars spent are poorly correlated with actual service provision, which is an issue we are currently examining in other research.

Means for our measures of spending and violence are listed in Table 3.1. From April 2002 to January 2010, the average number of SIGACTs per month is 0.016 incidents per 1000 people, or about 9.6 attacks annually in a median district of 50,000 residents. For comparison, this is about six times lower than in Iraq, which averaged about 0.098 attacks per 1000.[91] Insurgent violence is also highly skewed; out of almost 30,000 district-month observations, over 75% have no recorded events. SIGACTs also exhibit some interesting temporal and spatial patterns which will be discussed at the beginning of Section 3.5.

Like violence, spending by each of the three reconstruction programs also appears to be quite skewed in our sample. Average monthly spending by NSP is about \$0.23 per person while average LGCD and CERP spending is only half that (around \$0.10 per capita).[92] Over our entire sample of CERP, districts average around 0.003 projects per 1000 people per month, or about 1.8 projects annually for a district with 50,000 residents. All three programs seem to exhibit some significant outliers that are more than 10 standard deviations from the mean, though it is not the same district in each of the three cases.

---

[91] Average SIGACTs per 1000 in Iraq was 0.59 per half-year or 0.59/6 = 0.098 per month.

[92] For comparison, CERP spending per capita in Iraq averaged \$10.56 per half-year (about \$1.76 per person per month).

3.5 Results

Patterns of Violence in Afghanistan

Before discussing the main regression specifications, it is worth examining the time series of insurgent violence by province, which is reported in Figure 3.1 for the period from 2002 through 2010. Much of the country, including the province containing Kabul, is relatively quiet, while most of the violence occurs in provinces along the Afghanistan-Pakistan border. However, even in areas that might be predisposed to high levels of violence, SIGACTs are quite skewed geographically. While both Hilmand and Kandahar Provinces are in the traditional "heart" of Taliban territory in the South, almost one-third of total recorded SIGACTs are in Hilmand alone.

The time series plots also demonstrate two other features. First, violence is much higher since 2006. Second, violence follows a strong seasonal pattern. Both of these observations are borne out by the estimates in Table 3.2, which regress our measure of violence on temporal indicators. In Column (1), the rate of violent events has steadily increased since 2002, though the coefficients are generally much larger post-2006 compared to pre-2006. Before 2006, the average level of violence appears to be a relatively low 0.003 attacks per thousand per month. Between 2005 and 2006 however, rebel attacks more than quadrupled to approximately 0.014 attacks per 1000 residents and have continued to increase since then. This pattern in the coefficients broadly corresponds to the 2006 NATO push into the South; since our data are derived from military incident reports, a force increase into violent areas should result in more

SIGACTs –both because of increased engagement and because of improved reporting per engagement. Second, seasonality is evident in Column (2). While all three indicators for quarter are significant, the largest coefficient corresponds to Quarter 3 (July-August-September). These estimates are consistent with the insurgent "fighting season" in Afghanistan, which starts in early spring and last through early fall, accommodating climate and harvest.[93] However, year and season effects together explain only a very small fraction (approximately 5%) of the total variation in rebel violence.

The strongest predictor of current violence against government and allied forces in Afghanistan is past violence. In Column (3), lagged violence alone predicts nearly 66% of the total variance in SIGACTs between 2002 and 2010. The autoregressive coefficient is quite high at 0.9 (though still significantly different from unity --while violence is extremely persistent, it does not appear to follow a random walk). This coefficient changes little with the addition of year and quarter effects in Column (4). However, the large jump in post-2005 insurgent activity is still evident in Column (4) though the increasing trend and seasonality are somewhat muted by collinearity with lagged violence.

[93] Month indicators display this pattern clearly. Average violence increases starting in February, peaks in August, then declines monotonically through the end of the year.

Effect of Reconstruction Spending on Insurgent Violence

This section presents and discusses the empirical evidence for whether reconstruction spending by NSP, LGCD, or CERP appears to reduce violence once we control for local characteristics. The regression of interest is:

$$v_{it} = \beta \cdot g_{it} + \alpha_i + \delta' z_t + \varepsilon_{it} ,$$

where the subscript i denotes districts while t is time. The dependent variable v is rebel violence as measured by the number of SIGACTs per 1000 residents. g denotes spending per person by a particular reconstruction program. The vector z_t includes quarter and year effects to account for the seasonality in insurgent violence documented earlier. The number of districts and months covered by each reconstruction program varies, so estimation samples are limited to just those districts that ever received projects from a particular program. For example, the regressions that evaluate the effect of NSP on violence would be limited to just those 316 districts that ever had at least one NSP project.

The coefficient of interest is β, interpreted as how reconstruction spending affects violence within districts. If development aid is successful in increasing security, we would expect that CERP and LGCD would be strategically allocated to insecure places, resulting in a positive β in the cross-section. In contrast, since traditional development aid is limited to operating in areas that are sufficiently safe for NGOs and other civilians to enter, violence and program expenditure are probably negatively correlated for NSP. One way to control for districts that tend to have more insurgent activity is to include the district fixed effect α_i and look at how violence and spending are related within districts.

The general specification above is estimated using first differences (FD) as:

$$\Delta v_{it} = \beta \cdot \Delta g_{it} + \delta^{'} \Delta z_t + \Delta \varepsilon_{it}\,, \qquad (20)$$

where the district-specific effects α_i have been removed.[94] To account for size effects, both violence and spending measures are scaled by district population, and the regressions are also weighted by population. The first testable implication (H_1) is that β in Eq. (20) should be negative. Note that differencing is likely to exacerbate the noise in our SIGACT and spending data, further biasing us against finding an effect. In other words, our estimates likely underestimate the true effect of government service provision on insurgent violence.

Table 3.3 reports estimated coefficients of Eq. (20) for NSP, LGCD, and CERP separately. Starting first with the results for NSP, both the OLS and FD specifications imply that spending has no statistical effect on violence. As a traditional development program, NSP is generally absent in insecure areas where communities are reluctant to request the program. The difference between the negative OLS and positive FD specification implies a negative omitted variable bias that is consistent with that pattern: areas that tend to be prone to violence, perhaps due to norms or low costs of rebel violence, tend to receive low amounts of NSP. These economically small estimates are consistent with the findings of Beath et al. (2010), who show no

[94] As discussed earlier, violence exhibits very strong serial correlation, which is why we estimate using differences rather than fixed effects. Fixed effects regressions also resulted in statistical zeros except in the short four-month panel of CERP spending, which will be discussed below.

significant effect of project spending on measured violence in random assignment of NSP. One possible explanation is that NSP communities are in regions in which there is little violence to reduce. Another is that NSP is predominantly aimed at improving local livelihoods and is not implemented conditionally.

The next column estimates analogous regressions for LGCD. In contrast to NSP, LGCD is intended to improve stability and is specifically active in the volatile Southern and Eastern regions of Afghanistan. These features imply that (1) the cross-sectional correlation between LGCD activity and district violence should be positive and (2) positive selection bias once fixed district characteristics are removed. The OLS estimate is positive, which is consistent with the targeting that we should expect from security-improving aid programs that are being effectively delivered to insecure areas. The estimate in the FD specification is also positive though much smaller in magnitude than the OLS coefficient. Neither estimate is statistically distinguishable from zero. Together, the results for NSP and LGCD suggest patterns of selection bias in our reconstruction data that are consistent with each program's capabilities and objectives. However, neither one shows evidence of a significant effect on violence. Since neither NSP nor LGCD are implemented conditionally, this is consistent with model's prediction.

The bottom half of Table 3.3 presents results for CERP, which is a program that is both conditional and targeted at regions of active insurgency. The two columns labeled "CERP counts" uses the eight-year panel on CERP project activity while the other two columns present estimates for the FY2010 subset of CERP where we observe expenditure rather than just the project count. Looking first at the regressions on project count per capita, we see that the cross-sectional OLS coefficient is positive. CERP activity is disproportionately located in areas of high violence. In the FD specification in the next column, which controls for a district's predisposition to violence, spending and violence are still positively correlated though this coefficient is no longer statistically significant.

The results in the final columns use project expenditure rather than project counts as the main regressor. For the OLS estimate, we again see a positive cross-sectional correlation between CERP activity and rebel violence. The magnitude of this estimate is quite large compared to NSP or LGCD, though not statistically different from zero. An additional dollar per person of CERP predicts about 0.04 more SIGACTs per 1000 residents (about half a standard deviation according to Table 3.1). Once we difference out district characteristics, the spending coefficient estimate is -0.011, implying that an additional dollar per capita of CERP projects reduces violence by about 0.01 SIGACTs per 1000 residents.[95] While we cannot reject the hypothesis that CERP has no effect on violence, the magnitude of this coefficient is quite large. To put this in perspective, the increase in average violence between 2005 and 2006 (an additional 0.011 attacks per 1000 according to Column (1) of Table 3.2), could have been mitigated with an additional dollar of CERP per capita.

---

[95] This estimates for Afghanistan is remarkably similar to findings from Iraq, even though the context and environmental conditions differ somewhat between the two countries. In Iraq, the estimates for CERP were between -0.009 and -0.011 (Berman et al., 2011). The longer panel available in Iraq allowed more precise estimation of a significantly negative coefficient.

Overall, our results for CERP are qualitatively similar to those of LGCD since (1) violence and spending are positively correlated in the cross-section and (2) first differencing appears to correct for this positive selection bias. While the standard errors for NSP, LGCD, and CERP are all quite large, CERP is the only program of the three that appears to result in economically meaningful reductions of violence. This is consistent with the model's prediction that only conditional service provision has the potential to reduce violence.

Effectiveness of Small Projects

The second empirical hypothesis is that spending on small projects should be more effective at reducing violence than other spending. To test this, we classify LGCD and CERP projects as large or small based on their respective administrative guidelines, and then estimate Eq. (20) with large- and small-project spending simultaneously included:

$$\Delta v_{it} = \beta_{small} \cdot \Delta g_{it}^{small} + \beta_{large} \cdot \Delta g_{it}^{large} + \delta' \Delta z_t + \Delta \varepsilon_{it} . \qquad (21)$$

Then H_2 implies $\beta_{small} < \beta_{large}$. Funding regulations for both LGCD and CERP allow "small" projects to be authorized and implemented regionally without seeking the approval of higher-ranking (and more remote) officers. These smaller grants provide local commanders and aid officials more flexibility and responsiveness in meeting urgent community needs and, according to H_2, should be more effective at reducing rebel violence. That effect was quite strong in Iraq (Berman, Shapiro and Felter, 2011).

Since large projects are likely to require more security, large-project spending can be thought of as a proxy for the unobserved presence of US troops. This is especially true for large CERP since it is specifically a military program that, by its very nature, requires soldiers to be present before money can be spent. While LGCD does not require military forces to be present in the area, they do coordinate and work closely with the local PRT in the development and execution of projects so large LGCD projects might also proxy for unknown force allocation.

LGCD projects of $10,000 or under can be funded using "community small grants" (CSGs), which only require approval from the regional field program officer in the province and are not administered out of the Kabul.[96] CSG funds can be given directly to community actors like the CDC, rather than channeled through a non-governmental organization or other intermediary (USAID, 2008). Table 3.4 splits LGCD spending into large and small projects based on this administrative regulation and tests the hypothesis that spending via small projects is more effective at reducing violence. In Column (1), we see that higher spending on small-scale projects is associated with lower insurgent violence. Column (2) suggests that the reverse is true for large LGCD projects: higher spending predicts higher violence. The opposing signs remains even when both spending groups are included simultaneously (see Column (3)). While the coefficients on small and large spending are signed correctly, the standard errors for all three regressions are quite large and neither spending type is statistically significant. The coefficient on small spending is almost six times the

[96] USAID field program officers are stationed in each Provincial Reconstruction Team (PRT) and should be more accessible than the central office in Kabul. There are currently 27 PRTs spread through Afghanistan's 34 provinces.

magnitude of that for larger projects, which is consistent with Hypothesis H_2, but we cannot reject that they are equal ($p = 0.81$).

Similarly, "small" CERP projects also face fewer administrative hurdles and should be more effective at reducing violence than larger ones, according to H_2. Operating procedures for CERP in Afghanistan allow battalion commanders to authorize projects of $50,000 or less (USFOR-A, 2009, paragraph 5.K).[97] Table 3.5 explores H_2 using the FY2010 CERP data, splitting spending using the $50,000 cutoff. In the first difference specifications in the first three columns, both types of spending seem to reduce violence though the coefficient is much larger for small projects than small ones. The standard errors are again quite large, and the coefficients on large and small spending are not statistically different ($p = 0.35$). While we cannot reject that small-project spending has no effect on insurgent violence, the magnitude of this estimate is more than double that from the overall spending regressions in Table 3.3. The violence reductions from an additional dollar per capita in general CERP spending could be achieved by increasing small CERP by $0.38 (= 0.011/0.0290) per person.

Interestingly, the coefficient on small CERP spending is virtually unchanged between Column (1) and Column (3). Since the only difference between the two specifications is the inclusion of large spending, this suggests that unobserved location of US military units does not strongly bias estimation of β_{small}, at least in this particular four month sample.

---

[97] Also, SOP explicitly states that "Project splitting (separating procurements that are related to the same requirement in order to stay below the CERP approval thresholds) is prohibited."

The CERP results appear to be sensitive to the estimator used. When the base specification is estimated using district fixed effects rather than first differences, small CERP projects appear to be strongly violence-reducing. Fixed effect and first difference specifications both consistently estimate the same population parameter, though the latter places less weight on observations in the first and last period. In Column (4), the coefficient on small spending is -0.041; an extra $1 per person channeled through fast, high-impact projects reduces the rate of insurgent violence by 0.04 attacks per 1000 residents. In contrast, spending through large CERP projects appears to be only one-tenth as effective; an extra $1 through these projects lowers violence by 0.004 attacks per 1000 residents. We can also weakly reject the hypothesis that the effects of large and small spending are the same ($p = 0.07$).

Together, the difference in coefficients between large and small projects in both LGCD and CERP spending provide weak evidence supporting Hypothesis H_2, even though most estimates suffer from a lack of precision.

Reconstruction and Unobserved District Stability

Given the nonlinear relationship between SIGACTs and Afghan security perceptions in Figure 3.2, one possible explanation for the results so far is measurement error. The dotted lines in Figure 3.2 partition the underlying stability index into thirds. Perhaps development aid only reduces measured violence in the (stable) left third of the stability distribution where stability and measured violence are negatively correlated, while the reverse might be true in the (unstable) right third where instability is associated with less violence. Table 3.6 presents spending regressions separately for each third of the distribution. Column (1) consists of districts that appear to be relatively stable, defined as those with an index value less than -0.41 while Column (3) contains districts that are relatively unstable (index > 0.41). Districts used in Column (2) are "contested" and have an index between -0.41 and 0.41.

In the top two panels, both NSP and LGCD display a consistent zero effect across the entire stability distribution. In Panel C, CERP spending demonstrates similar statistical zeros, but the magnitude in Column (2) is over three times bigger than the overall estimate in Table 3.3 (β = -0.0110). While this estimate is not statistically different from the CERP estimates in Column (1) or Column (3), the difference in magnitudes between stable/unstable and contested districts suggests that the conditions of the theoretical model might better fit areas that are actively contested rather than those that are strongly controlled by one side. In general, splitting the estimation sample by unobserved stability does not qualitatively change the conclusion that development spending is ineffective at reducing insurgent violence.

The bottom panel reports CERP regressions for each third of the stability distribution by project size. In all three types of districts, higher small-project spending is correlated with lower insurgent violence, but the standard errors are such that we cannot reject a null hypothesis of no effect. The magnitudes across the three columns are quite similar to the estimates in Table 3.5 (β_{small} = -0.0290): there does not appear to be much heterogeneity in small-CERP's effectiveness across regions.

This is not the case for large CERP projects, which have zero effect in secure districts, a small negative effect in unstable districts, and a large but imprecise negative in contested districts. This heterogeneity across the stability distribution contrasts with the relatively constant effect of small-CERP. One explanation for this difference between large and small CERP coefficients is that rent-seeking behavior (e.g. illegal tolls, bribe payments in return for services) becomes more widespread as government presence increases. Small CERP projects, because they do not involve large amounts of money or multiple levels of contracting, are more insulated from such behavior because there are less rents to extract. While these estimates are not especially precise, they do suggest that failures in either aid conditionality or program implementation that is correlated with government control. They provide no evidence that development projects stabilize, in the sense of improving personal security for noncombatants.

Heterogeneity Across Project Sectors

Unlike traditional development programs that focus on one or a few specific interventions (e.g. digging wells, providing immunizations, buying school supplies), each reconstruction program examined here funds a wide variety of project types. While none of the three overall spending regressions appear to be strongly violence-reducing, the difference in estimates between big and small projects suggests that other dimensions of heterogeneity might be useful in guiding future aid practices or theoretical developments. In particular, one argument for why reconstruction in

Afghanistan is not “winning hearts and minds” is that rampant corruption along the implementation chain makes “dollars spent” a poor measure of “work done.” Large construction projects, in particular, could suffer severely from this issue as they involve multiple levels of contractors and subcontractors who could be colluding or otherwise acting anti-competitively. Infrastructure development like renovating the district center or building schools might improve rural livelihoods but also present new targets for rebel violence. Larger projects might require valuable construction equipment and foreign personnel, making them better targets for extortion.

Table 3.7 allows the possibility of sector-specific coefficients, splitting NSP projects by category. Of the four largest expenditure categories, only irrigation projects are statistically significant; moreover, irrigation projects are positively correlated with insurgent activity. This effect is also quite large, almost 10 times as large as the overall spending coefficient. This seemingly contradicts Hypothesis H_2 since improvements to agricultural productivity should be highly desirable in Afghanistan’s agrarian economy. One potential explanation for this positive estimate is project timing; irrigation projects might occur earlier in a district’s rehabilitation when it is still contested and vulnerable to rebel assaults whereas building construction or training programs cannot start until the district is reasonably secured against insurgents. Alternately, irrigation projects might increase the value of the land, making it more attractive to the insurgents. Spending from the other three major project categories (power, transport, and water supply/sanitation) are all statistical zeros.

We conduct an analogous sample splitting exercise using the LGCD data, though the project categories differ. In Table 3.8, irrigation again appears to be positively correlated with violence with a practically large coefficient, though this time it is not statistically different from zero. As we saw with the NSP regressions, the remaining three project sectors are far from significant. Interestingly, Road/Transport, which potentially involves a large amount of fraudulent and extortable construction activity, enters the regression with a negative coefficient.

CERP project categories (see Table 3.9) also generate a mix of coefficient signs though only the Other category is statistically different from zero. While this seems like an unusual finding, it might be due to compositional differences in the "Other" classification. Unlike NSP and LGCD, many CERP records were incomplete and did not list a category, and these missing values were treated as "Other." The coefficient on Agriculture/Irrigation spending is again large in effect and positive, though not statistically significant.

Cutting the NSP sample by the primary activity also provides an interesting insight: only construction projects positively predict insurgent attacks, and this estimate is weakly significant (see Table 3.10). Construction projects are both vulnerable to collusive bidding practices (and other forms of corruption along the implementation chain) and provide visible and obvious targets. At a cost of $224 million, construction projects are also, by far, the largest reconstruction activity among NSP projects. For comparison, the second largest activity (supply) is only $96.1

million. The other four major activities (supply, boring, basic access, and gravelling) do not appear to have any statistical effect on violence.

Overall, different categorical cuts of the spending data do not provide much insight about which project types are more effective in improving district stability. Given the lack of precision in most of the estimates so far, it is somewhat surprising that irrigation projects consistently exhibit a large, positive sign across all three programs. If building and maintaining irrigation networks is truly destabilizing, then that would oddly contradict one of the two testable implications of the theory so far. While we should be cautious in interpreting statistical zeros as strong evidence in either direction, there exploratory analysis here does suggest noticeable heterogeneity across projects types that might lead to future refinements of the current theory.

3.6 Conclusion

This paper tests two empirical hypotheses on current counterinsurgency theory in the Afghan context. Our results suggest that development aid in Afghanistan, whether it comes from the US military, USAID, or the Afghan government itself, has not been effective in reducing rebel attacks. These findings do not support the predictions of our theory. We expected to find that aid provided conditionally in the CERP program would be violence-reducing, even if unconditional aid would not. While overall service provision did not appear to reduce violence, we did find suggestive evidence that "small" CERP might be a useful tool to reduce violent insurgency.

Given the vast amount of resources, both monetary and human, that the US and other international donors have committed to rebuilding Afghanistan, a natural question to ask is: why does CERP spending not appear to be effective in reducing insurgent activity in Afghanistan when it did so in Iraq? The results suggest three potential explanations. First, the conditionality of aid is a necessary, and often overlooked condition underlying the theoretical model developed in Section 3.3. While the majority of CERP implementers in Afghanistan report practicing conditionality, a significant minority do not (Berman et al., 2011). Aid conditionality is the military's official policy for CERP, but the importance of conditionality implies that future reconstruction efforts would benefit from a greater emphasis or stronger guidelines about aid provision and community cooperation. Second, the lack violence-reduction raises questions about program effectiveness: perhaps money spent is not translating into services provided. The analysis assumed that project expenditure was a viable proxy for actual service provision, but this connection might be tenuous in an institutionally weak environment such as Afghanistan where monitoring is absent. Third, we've modeled only a static interaction. Consider a dynamic model in which noncombatants consider their future wellbeing. In that model we speculate that development would increase support for government only if it signaled a permanent shift in improved governance provision, either by signaling an institutional change or by reducing future marginal costs of governance. In the Afghan context the mismanagement of development funds might be signaling the opposite.

In research in progress we seek a longer time series of both (retrospective) CERP programs and violence, which may allow us to distinguish between these competing explanations, and provide more explicit policy recommendations.

3.7 Acknowledgement

Chapter 3 is currently being prepared for submission for publication of the material. Berman, Eli; Tiffany Chou. The dissertation author was a primary investigator and author of this material. Thanks to Michael J. Callen, Luke Nayef Condra, Michael J. Callen, Joseph Felter and Jacob Shapiro for excellent discussion and comments. We acknowledge grant #2007-ST-061-000001 by the United States Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events and by U.S. Department of Defense's Minerva Initiative (AFOSR Grant FA9550-09-1-0314). Any opinions, findings, conclusions or recommendations in this document are those of the authors and do not reflect views of the United States Department of Homeland Security or United States Department of Defense.

Table 3.1: Summary Statistics

	Observations	Number of Districts	Mean	Min	Max
Incidents per 1000 (Apr 2002 – Jan 2010)	37412	398	0.0174 (0.0867)	0	4.706
NSP spending (per capita) (Apr 2002 – Jan 2010)	29704	316	0.233 (0.399)	0	17.658
LGCD spending (per capita) (Jul 2007 – Dec 2009)	5328	144	0.109 (0.487)	0	30.319
- Large projects (>$10,000)	5184	144	0.102 (0.473)	0	26.143
- Small projects (≤ $10,000)	5184	144	0.0107 (0.0688)	0	5.649
CERP count (per 1000 pop) (Apr 2002 – Jan 2010)	32430	345	0.00329 (0.0154)	0	0.812
CERP spending (per capita) (Oct 2009 – Jan 2010)	808	202	0.0810 (0.386)	0	10.154
- Large projects (>$50,000)	808	202	0.0271 (0.298)	0	6.030
- Small projects (≤ $50,000)	808	202	0.0539 (0.230)	0	10.154

Notes: Incident records are from the CIDNE database. Means are weighted by Landscan population; standard deviations are in parentheses. LGCD is only active in the South and East. An observation is a district-month.

Table 3.2: Temporal Patterns of Violence

y = Incidents per 1000	(1)	(2)	(3)	(4)
Lagged incidents per 1000	-	-	0.901***	0.894***
			(0.0227)	(0.0234)
2003	0.000638***	0.00315***	-	0.000154
	(0.000112)	(0.000443)		(0.000122)
2004	0.00202***	0.00453***	-	0.0000793
	(0.000240)	(0.000568)		(0.000147)
2005	0.00327***	0.00578***	-	0.000523***
	(0.000416)	(0.000771)		(0.000175)
2006	0.0138***	0.0163***	-	0.00213***
	(0.00222)	(0.00261)		(0.000556)
2007	0.0237***	0.0262***	-	0.00324***
	(0.00450)	(0.00488)		(0.00861)
2008	0.0327***	0.0352***	-	0.00422***
	(0.00597)	(0.00632)		(0.000908)
2009	0.0547***	0.0572***	-	0.00725***
	(0.0111)	(0.0115)		(0.00144)
2010	0.0598***	0.0699***	-	0.0121***
	(0.0145)	(0.0161)		(0.00308)
Quarter 2 (Apr – Jun)	-	0.00606***	-	0.00124**
		(0.00100)		(0.000510)
Quarter 3 (Jul – Sep)	-	0.0145***	-	0.00186***
		(0.00269)		(0.00598)
Quarter 4 (Oct – Dec)	-	0.00951***	-	-0.00213***
		(0.00193)		(0.000627)
Constant	0.0000339**	-0.0100***	0.00232***	-0.000211
	(0.0000173)	(0.00182)	(0.000420)	(0.000393)
R^2	0.047	0.051	0.762	0.764
# obs	37412	37412	37014	37014
# districts	398	398	398	398

Notes: Standard errors clustered by district are in parentheses. Regressions are weighted by population. An observation is a district-month. The dependent variable is SIGACTs per 1000 residents. Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.3: Development Spending and Violence

y = Incidents per 1000	NSP spending		LGCD spending	
	OLS	FD	OLS	FD
Spending ($/capita)	-0.000990	0.00116	0.0164	0.000246
	(0.00342)	(0.00300)	(0.0118)	(0.00319)
Year FE		X		X
Quarter FE		X		X
R^2	0.000	0.002	0.002	0.005
# obs	29704	29388	5328	5184
# districts	316	316	144	144

y = Incidents per 1000	CERP counts		CERP spending	
	OLS	FD	OLS	FD
Activity (per capita)	0.812**	0.0587	0.0387	-0.0110
	(0.323)	(0.0366)	(0.0233)	(0.00967)
Year FE		X		
Quarter FE		X		
R^2	0.020	0.002	0.004	0.003
Observations	32430	32085	808	606
Number of districts	345	345	202	202

Notes: Robust standard errors in parentheses, clustered by district. Regressions are weighted by population. An observation is a district-month. Sample is strongly balanced to include only districts that ever have NSP or LGCD as where appropriate. CERP projects count per 1000, rather than spending per capita, is the explanatory variable in the "CERP counts" column. Quarter and year fixed effects are omitted from CERP spending regressions since there is only one year of data. Dependent variable is insurgent events per 1000 population as recorded by CIDNE. Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.4: Small vs. Large LGCD projects

y = Incidents per 1000	(1)	(2)	(3)
Spending (small)	-0.00487	-	-0.00512
	(0.0244)		(0.0243)
Spending (large)	-	0.000595	0.000714
		(0.00271)	(0.00264)
Year FE	X	X	X
Quarter FE	X	X	X
R^2	0.005	0.005	0.005
p-value for: $\beta_{small} = \beta_{large}$	-	-	0.811
Observations	5184	5184	5184
Number of districts	144	144	144

Notes: Robust standard errors in parentheses, clustered by district. Regressions are weighted by population. An observation is a district-month. Sample is strongly balanced to include only districts that ever have LGCD. "Small" projects are those that spend $10,000 or less. All regressions estimated using the first difference specification in Eq. (20). Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.5: Small vs. Large CERP projects, Oct 2009 – Apr 2010

y = Incidents per 1000	FD (1)	FD (2)	FD (3)	FE (4)
Spending (small)	-0.0291	-	-0.0290	-0.0437**
	(0.0267)		(0.0268)	(0.0206)
Spending (large)	-	-0.00363	-0.00306	-0.00356
		(0.00390)	(0.00389)	(0.00360)
R^2	0.006	0.000	0.006	0.942
p-value for: $\beta_{small} = \beta_{large}$	-	-	0.348	0.066
Observations	606	606	606	808
Number of districts	202	202	202	202

Notes: Robust standard errors in parentheses, clustered by district. Regressions are weighted by population. An observation is a district-month. Sample is strongly balanced to include only districts that ever have spending data from CERP. “Small” projects are those that spend $50,000 or less. Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.6: Development Spending and Violence, by Stability Index

y = Incidents per 1000	(1) 1st Third	(2) 2nd Third	(3) 3rd Third
A: NSP			
Spending ($/capita)	0.00413	0.00881	0.00389
	(0.00415)	(0.00815)	(0.00651)
R^2	0.002	0.002	0.004
Observations	6510	6417	5766
Number of districts	70	69	62
B: LGCD			
Spending ($/capita)	-0.00342	0.00765	-0.00427
	(0.00387)	(0.00664)	(0.0235)
R^2	0.008	0.003	0.015
Observations	900	1225	1368
Number of districts	25	34	38
C: CERP Spending (overall)			
Spending ($/capita)	-0.00637	-0.0385	-0.0149
	(0.0145)	(0.0263)	(0.00900)
R^2	0.006	0.009	0.006
Observations	108	141	144
Number of districts	36	47	48
D: CERP Spending (by size)			
Spending (small)	-0.0531	-0.0240	-0.0329
	(0.0973)	(0.0695)	(0.0261)
Spending (large)	0.00106	-0.0561	-0.00745**
	(0.00156)	(0.0385)	(0.00290)
R^2	0.047	0.010	0.010
Observations	108	141	144
Number of districts	36	47	48

Notes: Robust standard errors in parentheses, clustered by district. Regressions estimated using first difference specification and weighted by population. Regressions are estimated separately for each third of the stability index distribution (index < -0.41, between -0.41 and 0.41, index > 0.41). Stability index generated from the September 2009 wave of the ANQAR survey. Higher values of the index are interpreted as less stability. An observation is a district-month. "Small" CERP projects are those that spend $50,000 or less. Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.7: NSP Spending by Project Sector

y = Incidents per 1000	(1) Irrigation	(2) Power	(3) Transport	(4) WaterSan	(5) Other
Spending ($/capita)	0.0149*	-0.00436	0.00342	0.00131	0.00621
	(0.00790)	(0.00842)	(0.00508)	(0.00680)	(0.00512)
Year FE	X	X	X	X	X
Quarter FE	X	X	X	X	X
R^2	0.002	0.002	0.002	0.002	0.002
# obs	29388	29388	29388	29388	29388
# districts	316	316	316	316	316
Dollars spent (millions)	83.9	113	163	116	54.2

Notes: Robust standard errors in parentheses, clustered by district. Regressions are weighted by population. An observation is a district-month. Sample is strongly balanced to include only districts that ever have NSP. "Other" consists of 7 other project sectors. All regressions estimated using the first difference specification in Eq. (20). Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.8: LGCD Spending by Project Sector

y = Incidents per 1000	(1) Civil Service Capacity Building	(2) Road/Transport	(3) Irrigation	(4) Agriculture	(5) Other
Spending ($/capita)	0.00799	-0.00360	0.0133	-0.00557	-0.000431
	(0.0124)	(0.00551)	(0.0146)	(0.0175)	(0.00526)
Year FE	X	X	X	X	X
Quarter FE	X	X	X	X	X
R^2	0.005	0.005	0.005	0.005	0.005
# obs	5184	5184	5184	5184	5184
# districts	144	144	144	144	144
Dollars spent (millions)	16.5	12.6	6.80	4.84	17.1

Notes: Robust standard errors in parentheses, clustered by district. Regressions are weighted by population. An observation is a district-month. Sample is strongly balanced to include only districts that ever have LGCD. Insurgent events come from CIDNE reports. “Other” consists of 6 other project sectors. All regressions estimated using the first difference specification in Eq. (20). Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.9: CERP counts by Project Category

y = Incidents per 1000	(1) Education	(2) Energy	(3) Community Development	(4) Agriculture & Irrigation	(5) Other
Project count (#/capita)	0.0782	-0.0597	-0.121	0.250	0.109*
	(0.0857)	(0.140)	(0.156)	(0.188)	(0.0586)
Year FE	X	X	X	X	X
Quarter FE	X	X	X	X	X
R^2	0.002	0.002	0.002	0.002	0.002
# obs	32085	32085	32085	32085	32085
# districts	345	345	345	345	345

Notes: Robust standard errors in parentheses, clustered by district. Regressions are weighted by population. An observation is a district-month. Sample is strongly balanced to include only districts that ever have CERP. "Other" includes 10 other project categories. All regressions estimated using the first difference specification in Eq. (20). Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.10: NSP Spending by Project Activity

y = Incidents per 1000	(1) Construction	(2) Supply	(3) Boring	(4) Basic Access	(5) Gravelling	(6) Other
Spending ($/capita)	0.00967*	-0.00549	-0.00490	-0.00916	0.00806	0.0334**
	(0.00553)	(0.00903)	(0.00845)	(0.00652)	(0.00588)	(0.0156)
Year FE	X	X	X	X	X	X
Quarter FE	X	X	X	X	X	X
R^2	0.002	0.002	0.002	0.002	0.002	0.002
# obs	29388	29388	29388	29388	29388	29388
# districts	316	316	316	316	316	316
Dollars spent (millions)	224	96.0	59.7	61.8	57.5	31.0

Notes: Robust standard errors in parentheses, clustered by district. Regressions are weighted by population. An observation is a district-month. Sample is strongly balanced to include only districts that ever have NSP. “Other” consists of 13 other project activities. All regressions estimated using the first difference specification in Eq. (20). Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.11: NSP and LGCD spending, full results

y = Incidents per 1000	NSP			LGCD		
	(1)	(2)	(3)	(4)	(5)	(6)
Spending ($/capita)	0.00168	0.000798	0.00116	0.000653	0.000779	0.000246
	(0.00292)	(0.00300)	(0.00300)	(0.00312)	(0.00314)	(0.00319)
Year FE		X	X		X	X
Quarter FE			X			X
R^2	0.000	0.000	0.002	0.000	0.000	0.005
# obs	29388	29388	29388	5184	5184	5184
# districts	316	316	316	144	144	144

Notes: Robust standard errors in parentheses, clustered by district. Regressions are weighted by population. An observation is a district-month. Sample is strongly balanced to include only districts that ever have NSP or LGCD as appropriate. All regressions estimated using first differences. Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

Table 3.12: Small vs. Large CERP projects, Oct 2009 – Apr 2010

y = Incidents per 1000	All CERP		Large CERP		Small CERP	
	(1)	(2)	(3)	(4)	(5)	(6)
Spending ($ per capita)	0.0387	-0.0169*	-0.0100	-0.00568	0.126**	-0.0440**
	(0.0293)	(0.00991)	(0.00839)	(0.00381)	(0.0489)	(0.0205)
Includes district FE?		X		X		X
# obs	0.004	0.942	0.000	0.941	0.015	0.942
# districts	808	808	808	808	808	808
R^2	202	202	202	202	202	202

Notes: Robust standard errors in parentheses, clustered by district. Regressions are weighted by population. An observation is a district-month. Sample is strongly balanced to include only districts that ever have spending data from CERP. "Small" projects are those that spend $50,000 or less. Symbols denote coefficients significantly different from zero at 10%(*), 5%(**), and 1%(***).

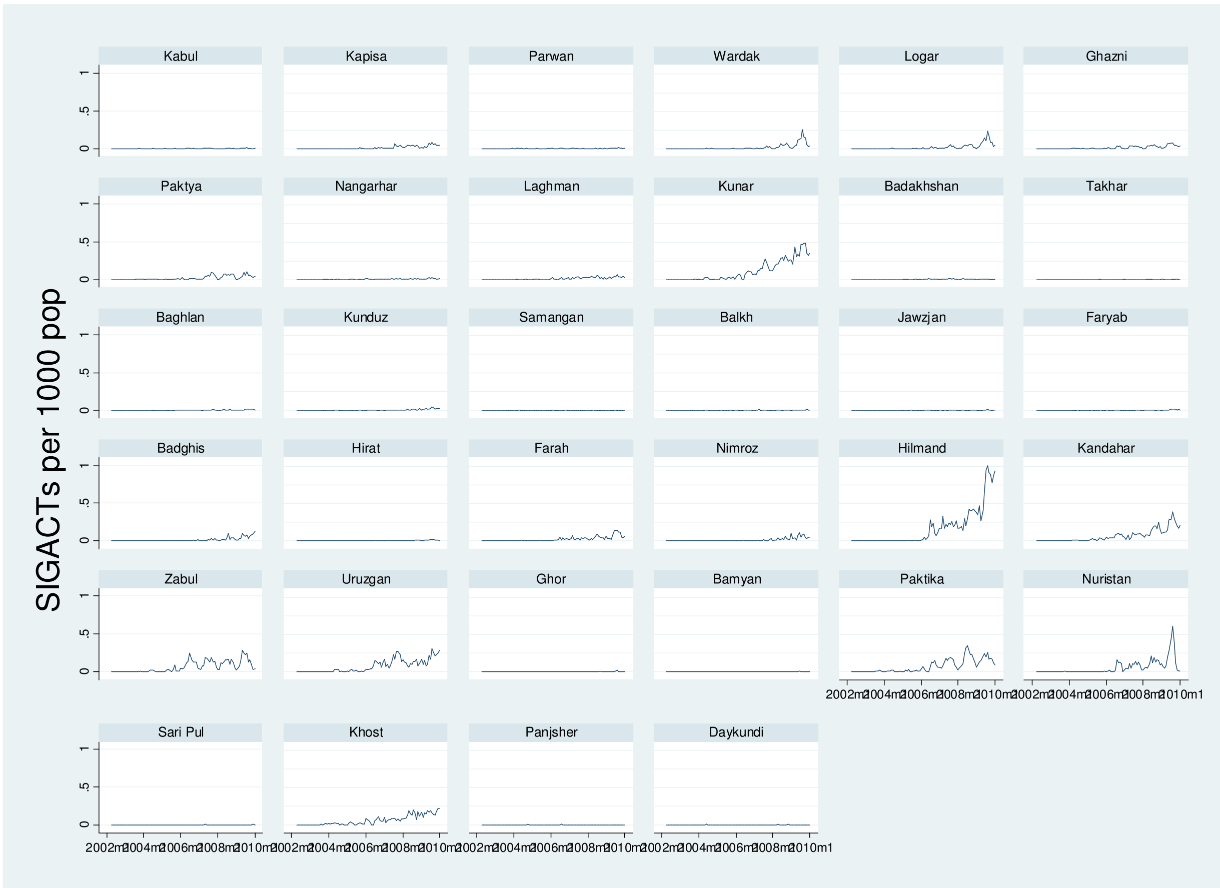

Figure 3.1: Violence in Afghanistan

Notes: Event-level SIGACT records are from the US military's CIDNE database.

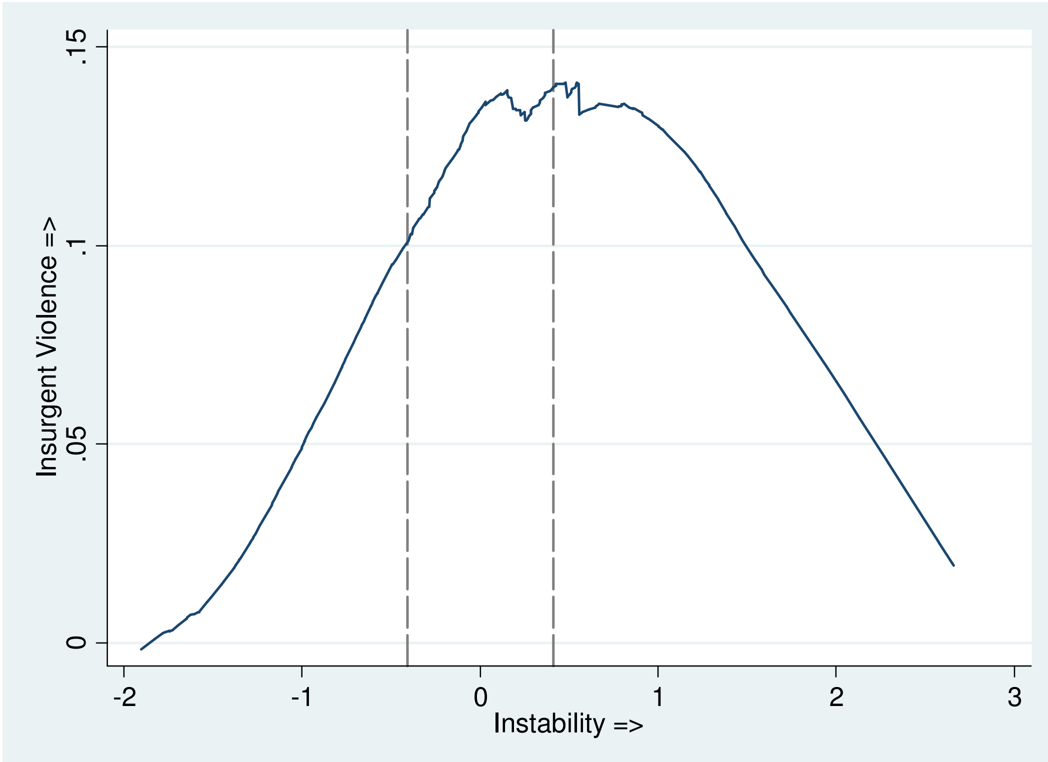

Figure 3.2: Nonlinear Relationship between Insurgent Violence and Afghan Stability Perceptions

Notes: Insurgent violence is measured as the number of SIGACTs per 1000 residents. SIGACT records are from the US military's CIDNE database. Stability perceptions are along the x-axis, with increasingly worse outcomes going to the right. Stability index is generated as a weighted sum of seven security questions from the Afghan National Quarterly Assessment Reports (ANQAR). Dotted lines at -0.41 and 0.41 denote thirds of the stability index distribution. Data are smoothed using LOWESS.

3.8 Appendix

Table 3.11 presents estimates of our first difference spending regressions for NSP and LGCD that sequentially include year and quarter effects. For both programs, neither year nor seasonal effects appear to explain much of the variation in violence.

Results for FY 2010 CERP spending are in Table 3.12. These regressions are estimated using both OLS and district fixed effects for large, small, and total spending. Note the positive coefficients in Columns (1) and (5); both total and small CERP spending are correlated with higher violence in the cross-section (small spending is significantly so). Spending on large projects still appears to have no statistical effect on violence once we control for fixed district characteristics. However, small project spending is violence-reducing, consistent with H_2.

Comparative Statics

This section formally derives the comparative statistics results cited in Section 2. First, recall that $p^*(m, g, v) = h(m) \cdot f \cdot (g + v - n_L)$.

Proposition 1: p^* is increasing in each of its three arguments.

Proof of Proposition 1:

For g and v, partially differentiating implies that:

$$\frac{\partial p^*}{\partial g} = \frac{\partial p^*}{\partial v} = h(m) f > 0 .$$

Differentiating with respect to m implies that:

$$\frac{\partial p^*}{\partial m} = h'(m) f\left(g + v - n_L\right) > 0 .$$

The inequality follows from the assumption that $h(m)$ is an increasing function and that the distribution of *n* was wide enough such that $n_L \le v + g \le n_U$. ■

Proposition 2: $v^*(m, g)$ is decreasing in both arguments.

Proof of Proposition 2:

Recall that *R*'s first-order condition is:

$$\begin{aligned}\frac{\partial}{\partial v} EU_R &= [1 - p^*]A'(v) - A(v)\frac{\partial}{\partial v} p^* - B'(v) \\ &= [1 - p^*]A'(v) - A(v)h(m)f - B'(v) = 0.\end{aligned}$$

Differentiating again with respect to *v* yields:

$$\begin{aligned}\frac{\partial^2 EU_R}{\partial v^2} &= (1 - p^*)A''(v) - \frac{\partial p^*}{\partial v} A'(v) - A'(v)h(m)f - B''(v) \\ &= (1 - p^*)A''(v) - 2A'(v)h(m)f - B''(v).\end{aligned}$$

$A''(v) < 0$ and $B''(v) > 0$ by assumption so this derivative is negative.

By the implicit function theorem:

$$\begin{aligned}\left.\frac{\partial v^*}{\partial g}\right|_m &= -\left[\frac{\partial^2 EU_R}{\partial v^2}\right]^{-1}\left[\frac{\partial^2 EU_R}{\partial g \partial v}\right] \\ &= -\left[(-)\right]^{-1}\left[- A'(v)\frac{\partial p^*}{\partial g}\right].\end{aligned}$$

By Proposition 1, the bracketed term on the right is negative. Hence $\partial v^* / \partial g < 0$.

Similarly:

$$\left.\frac{\partial v^*}{\partial m}\right|_g = -\left[\frac{\partial^2 EU_R}{\partial v^2}\right]^{-1}\left[\frac{\partial^2 EU_R}{\partial m \partial v}\right]$$
$$= -[(-)]^{-1}\left[-\frac{\partial p^*}{\partial m}A'(v) - A(v)h'(m)f\right].$$

and again by Proposition 1, the bracketed term on the right is negative. ■

Proposition 3: Effort m^* and service provision g^* are both increasing in violence.

Proof of Proposition 3:

Recall that G's first-order condition for m is:

$$\frac{\partial}{\partial m}EU_R = -A(v)\frac{\partial p^*}{\partial m} + D'(m) = 0$$
$$= -A(v)h'(m)f(g+v-n_L) + D'(m).$$

Differentiating again with respect to m yields:

$$\frac{\partial^2 EU_R}{\partial m^2} = -A(v)h''(m)f(g+v-n_L) + D''(m)\,.$$

$h''(m) < 0$ and $B''(v) > 0$ so this derivative is positive. By the implicit function theorem:

$$\frac{\partial m^*}{\partial v} = -\left[\frac{\partial^2 EU_G}{\partial m^2}\right]^{-1}\left[\frac{\partial^2 EU_G}{\partial v \partial m}\right]$$
$$= -[(+)]^{-1}\left[-A'(v)\frac{\partial p^*}{\partial m} - A(v)\frac{\partial^2 p^*}{\partial v \partial m}\right]$$
$$= -[(+)]^{-1}\left[-A'(v)\frac{\partial p^*}{\partial m} - A(v)h'(m)f\right].$$

By Proposition 1, the rightmost bracketed term is negative and the whole derivative is positive.

G's first-order condition for *g* is:

$$\frac{\partial}{\partial g} EU_R = -A(v)\frac{\partial p^*}{\partial g} + H'(g) = 0$$
$$= -A(v)h(m)f + H'(g).$$

The only place where *g* appears is as the argument for $H'(\cdot)$. Hence the second derivative with respect to *g* is also positive. By the implicit function theorem:

$$\left.\frac{\partial g^*}{\partial v}\right|_m = -\left[\frac{\partial^2 EU_G}{\partial g^2}\right]^{-1}\left[\frac{\partial^2 EU_G}{\partial v \partial g}\right],$$
$$= -[(+)]^{-1}[-A'(v)h(m)f]$$

which is also positive. ■

Construction of a Stability Index

The main dataset used to construct the stability index demonstrated in Figure 3.2 is the Afghan National Quarterly Assessment Report (ANQAR). This is a nationally representative survey of Afghans that includes questions about security and public service provision. The stability index itself is the factor score for the first principal component from a PCA of seven ANQAR security questions. The selection criteria for the included questions were relatively simplistic since we only required that they (a) were consistently asked across all ANQAR waves and (b) seemed intuitively related to underlying security perceptions among Afghans:

1. Was the security situation in your mantaqa bad?
2. Was security worse in your mantaqa compared to 6 months ago?
3. Do you feel unsafe traveling outside during the day?

4. Has your mantaqa been affected by… operations & bombings?
5. Has your mantaqa been affected by… criminality?
6. Has your mantaqa been affected by… AGE (anti-government elements) activities?
7. Do you see the police (ANP) around less than once a week?

By construction, the stability index is mean zero, with a standard deviation of one. We use the factor analysis weights from Wave 5 of ANQAR, which was conducted in September 2009.

3.9 References

Beath, Andrew, Fotini Christia, & Ruben Enikolopov. (2011). Winning Hearts and Minds? Evidence from a Field Experiment in Afghanistan. MIT mimeo.

Beath, Andrew, Fotini Christia, Ruben Enikolopov, & Shahim Ahmad Kabuli. (2010). Randomized Impact Evaluation of Phase-II of Afghanistan's National Solidarity Programme (NSP): Estimates of Interim Program Impact from First Follow-Up Survey. World Bank mimeo, July.

Becker, Gary S. (1968). Crime and Punishment: An Economic Approach. *Journal of Political Economy*, 76, 169-217.

Berman, Eli, Joseph H. Felter, Jacob N. Shapiro, & Michael Callen. (2011). Do Working Men Rebel? Insurgency and Unemployment in Afghanistan, Iraq, and the Philippines. *Journal of Conflict Resolution*. March.

Berman, Eli, Jacob N. Shapiro and Joseph H. Felter. (2011). Can Hearts and Minds be Bought? The Economics of Counterinsurgency in Iraq. April.

Condra, Luke N., Joseph H. Felter, Radha K. Iyengar, & Jacob N. Shapiro. (2010). The Effect of Civilian Casualties in Afghanistan and Iraq. NBER Working Paper 16152. July.

Crost, Benjamin, Joseph H. Felter, & Patrick Johnston. (2010). Community Driven Development and Civil Conflict – Evidence from the Philippines. Stanford mimeo.

Fearon, James. (2008). Economic Development, Insurgency, and Civil War. In E. Helpman (Ed.), *Institutions and Economic Performance*. Cambridge: Harvard University Press.

Fearon, James D. & David D. Laitin. (2003). Ethnicity, Insurgency, and Civil War. *American Political Science Review*, 97, 75-90.

Galula, David. (1964). *Counterinsurgency Warfare: Theory and Practice*. New York: Praeger.

Grossman, Herschel I. (1991). A General Equilibrium Model of Insurrections. *American Economic Review*, 81(4), 912-921.

Kalyvas, Stathis. (2006). *The Logic of Violence in Civil War*. Cambridge: Cambridge University Press.

Ministry of Rural Reconstruction and Development. (2007). National Solidarity Programme Operational Manual, Version V. http://www.nspafghanistan.org/default.aspx?sel=16

Popkin, Samuel L. (1979). *The Rational Peasant: The Political Economy of Rural Society in Vietnam*. Berkeley: University of California Press.

Rashid, Ahmed. (2008). *Descent into Chaos: The United States and the Failure of Nation Building in Pakistan, Afghanistan, and Central Asia*. New York: Viking.

Skaperdas, Stergios. (1996). Contest success functions. *Economic Theory*, 283-290.

Special Inspector General for Afghanistan Reconstruction. (2009). Increased Visibility, Monitoring, and Planning Needed for Commander's Emergency Response Program in Afghanistan. SIGAR Audit-09-5. September.

Special Inspector General for Afghanistan Reconstruction. (2011). Commander's Emergency Response Program in Laghman Province Provided Some Benefits, but Oversight Weaknesses and Sustainment Concerns Led to Questionable Outcomes and Potential Waste. SIGAR Audit-11-7. January.

United States Forces – Afghanistan. (2009). Money As A Weapon System Afghanistan (MAAWS-A), Commander's Emergency Response Program, Standard Operating Procedures. USFOR-A Pub 1-06. December.

United States Agency for International Development. (2008). Local Governance and Community Development South and Eastern Regions Annual Report October 2007 – October 2008.

United States Agency for International Development. (2010). Local Governance and Community Development (LGCD) Fact Sheet. June. http://afghanistan.usaid.gov/documents/document/document/1020

United States Army. (2006). Counterinsurgency. FM 3-24. December.

United States Army. (2009). *Commander's Guide to Money as a Weapons System*. April. http://usacac.army.mil/cac2/call/docs/09-27/toc.asp

Wilder, Andrew. (2009). Losing Hearts and Minds in Afghanistan. *Afghanistan, 1979-2009: In the Grip of Conflict*. Middle East Institute. December.