This is to certify that the

dissertation entitled

TWO ESSAYS IN LABOR AND GENDER ECONOMICS AND ONE
IN THE RETURNS TO EXPERIENCE IN AVIATION SAFETY

presented by

Yuri S. D. Soares

has been accepted towards fulfillment
of the requirements for

Ph.D. degree in Economics

Major professor

Date June 19 2002

TWO ESSAYS IN LABOR AND GENDER ECONOMICS AND ONE IN THE RETURNS TO EXPERIENCE IN AVIATION SAFETY

By

Yuri S. D. Soares

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

Department of Economics

2002

ABSTRACT

TWO ESSAYS IN LABOR AND GENDER ECONOMICS AND ONE IN THE RETURNS TO EXPERIENCE IN AVIATION SAFETY

By

Yuri S. D. Soares

This dissertation contains three essays. The first essay looks at household consumption and labor supply response to the gender composition of household children in Brazil. Although we find that consumption is largely unresponsive to child gender, mother's labor supply is remarkably responsive. Mothers are more likely to supply labor the more girls there are in the household relative to boys. This effect is found for children ages 8-16. We find evidence that girls are substituting for mother's labor in home production, thus enabling mothers to work in the labor market. We find no evidence that the labor supply responses are due to either investment motives or to parental preferences for boys.

The second chapter looks at the evolution of female labor supply in Brazil, throughout the 1980s and 1990s, from a cohort perspective. We find that married women's labor supply accounted for the vast majority of increases in participation. We also find that although cohorts explain the majority of the change in participation, increases in schooling and decreases in fertility are also important factors. For both married and single women, we found that women have increasingly entered informal and self-employment versus formal wage work. The evidence also suggests that although women were more likely to work informally, they held a sizeable qualifications advantage in highly male-concentrated occupations. This advantage

decreased over time, as occupations became less segregated. Lastly, we attempt to separate cohort, period and age effects in women's employment. To the extent that it is possible to identify these effects individually, we find that cohort changes are more important than period changes in explaining increased female participation and employment.

In the last chapter we determine the importance of pilot experience and training in mitigating pernicious aviation safety outcomes. We find that in all aviation flight categories total flight hours is key in reducing the incidence of pilot error in accidents. It is also important, in varying degrees, in reducing both the fatality rate of accidents and the extent of aircraft damage. Other experience-related variables that are found to mitigate harmful and deadly safety outcomes are the proportion of flight time as pilot in command and the proportion of flight time in the make and model of aircraft being flown. In addition, the evidence shows strong non-linearities in the returns to experience. For general aviation flights almost all gains are realized after the first thousand hours of flight. For trunk carriers returns are steep for the first three or four thousand hours of flight.

ACKNOWLEDGMENTS

A number of people have helped, in one way or another, to produce this dissertation. John Strauss spent numerous hours reading and commenting on all drafts of all three essays. His contribution can not be overstated. I also thank John for his encouragement and patience. Without his help I would not have completed this project. David Neumark had several key insights, particularly in the second and third essays. As always, his comments are thoughtful and to the point. I also thank Jeff Biddle for his involvement, particularly his interest in the rather exoteric topic of aviation safety and his suggestions for further work in this field. I thank all three of my original committee members, as well as John Giles, for serving on my committee, reading the text and providing comments. I also thank Sergei Soares for comments on the first two papers, as well as for help with the PNAD datasets. Other people who provided useful comments were Ricardo Paes e Barros, the participants of the two seminars at Michigan State University and Marina Farias.

I must also thank those who provided me with the data. Here I thank Ricardo Paes e Barros and IPEA for providing access to the PNADs and Stan Smith at the NTSB for the AIDS database.

Lastly, I would like to thank my family and friends for the encouragement and their love, and for helping me through the rough times. In particular, I thank my wife Marina for her endless patience, love and dedication.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Ch

I.

allo

emp

allo

gen

evi

per

doc

Sch

fem

ant

Sub

girl

allo

and

of t

dev

[1] See
evide

Chapter 1: Gender bias in consumption and labor supply in Brazil

I. Introduction

Recently there has been increased interest in the issue of intra-household resource allocation and the role of children in that allocation. More specifically, both theoretic and empirical work has been done regarding the role that a child's gender plays in the allocation of resources within the household—how does a household respond to the gender of children within the household? This work, to some extent, grew out of evidence of gender bias in South Asia. Sen (1990, 1992) and later Coale (1991) showed perhaps the most striking evidence of differential treatment of boys and girls in their documentation of excess female mortality in South Asia. Work by Rosenzweig and Schultz (1982) also found evidence of gender bias, through differential survival rates for females and males. Other evidence of gender bias, in the form of differences in anthropometrics and nutrition for boys and girls (almost exclusively from the Indian Subcontinent), is plentiful, although many studies found no differences between boys and girls.[1]

Perhaps another reason for the renewed interest in issues of intra-household allocation is a greater availability of data on household consumption, income and savings, and just as importantly, theoretical developments which allow for a more careful analysis of these data. Two recent developments are prominent in this field. The first is the development of bargaining theories of household behavior, which present an attractive

[1] See Drèze and Sen (1989) and Harris (1990) for a review of the nutrition, health and anthopometric evidence.

alte

met

con

wil

resu

con

gen

exa

In s

des

pre

II.

hou

that

may

of d

gen

not

[2] See
(199
[3] See

alternative to the unitary (or dictatorial) model of household behavior.[2] The second is the methodology developed for deducing household resource sharing rules from data containing only household-level information on consumption and savings.[3] Both of these will be discussed in section II, as they are important in interpreting possible gender bias results.

In this paper we will address two issues. The first is the effect of child gender on consumption, as well as on labor force participation in Brazil. Is there evidence of gender bias in Brazil? If bias is found, why do we observe it? In other words, we will examine which hypotheses of household behavior are consistent with the findings here. In section II, theory and previous results are discussed. In section III the data is described. In section IV the empirical model is described. In section V results are presented. Section VI concludes the paper.

II. Theory and review of previous results

The interpretation of gender bias results depends on which hypothesis of household behavior is assumed. If one assumes that household preferences are unitary—that is, that all household members share the same preferences—then a gender bias result may originate because of different net returns associated with boys and girls, or because of different household preferences for boys and girls. However, if one fails to find gender bias under the unitary model, it is not necessarily true that household members do not prefer boys to girls or vice versa. It may be that, in the context of a bargaining model,

[2] See Lundberg and Pollak (1993), Manser and Brown (1980), McElroy and Horney (1981), and McElroy (1990).

[3] See Browning et al. (1994), Bourguignon and Chiappori (1992), and Chiappori (1988, 1992).

mo

out

hou

the

alr

hou

ag

fur

as

bec

its

uni

by

hou

util

pro

ind

eva

unit

supp

4 The

mothers attempt to favor girls and fathers attempt to favor boys, but that this effect nets out when one observes overall consumption.

This paper is not an attempt to discriminate between competing theories of household behavior. However, it is still instructive to briefly review the competing theories of household behavior, as later results will refer back to them. The reader who is already familiar with them may safely skip the remainder of this section.

The first class of models are those based on the so-called "unitary" view of households.[4] The unitary framework maintains that households behave as though one agent made all household decisions.[5] Thus, the household maximizes a given utility function, subject to production and budget constraints. One can think of the unitary view as one in which all household members share the same preferences. This framework has been the workhorse of microeconomics of household decisions, and it is hard to overstate its contribution to the field. A number of reasons underlie its popularity. First, the unitary framework's simplicity presents an appealing way to model economic behavior by households, allowing economists to analyze economic phenomenon by treating the household as the economic unit, without having to worry about how the household's utility function came about. Also, since one can typically observe characteristics such as production, consumption, and savings behavior at the household level, but not at the individual level, it is easier to directly test hypotheses about household behavior and evaluate the effects of policy on households by treating the household as the economic unit. For example, by using the unitary framework it is possible to derive demand and supply functions for households, and to recover household preferences by using

[4] The unitary model is also referred to as the common preference or dictatorial model.

hous

assur

of ho

deter

and t

intro

depe

wife

utili

barg

hete

each

barg

barg

beca

spec

outs

a co

max

func

[5] See

household level data. Of course this is only possible because the unitary model is assumed to be true.

Becker's (1973) pioneering work used this framework to first discuss allocation of household resources in the context of a marriage model. In his model, allocations are determined by maximizing the household dictator's utility function, subject to income and time constraints, and to an exogenous sharing rule. In a later paper Becker also introduces interdependent utilities or the 'caring' type. That is, the dictator's utility depends on the spouse's utility also. By doing this, allocations are not such that the wife's utility is just equal to her reservation utility; instead the wife will be outside her utility frontier outside the household because the husband cares about her utility.

An alternative and more flexible way to model household behavior is the bargaining framework. In the bargaining framework households are composed of heterogeneous agents, who each have their own set of preferences, and who bargain with each other for control of household resources. Household allocations are the result of this bargaining process.[6] Manser and Brown (1981) apply a Nash-bargained solution to the bargaining model of a two-member household. In their model individuals may marry because of shared goods and because of greater utility associated with marriage. They specify a model in which members face a reservation utility corresponding to the utility outside the marriage. The resulting allocations of the bargaining model are then given by a cooperative solution in which the product of the gain function for each member is maximized, subject to budget and time constraints. The resulting demands are then functions of husband and wife's incomes, prices, as well as factors that may shift the

---

[5] See Behrman, Pollak and Taubman (1982) for an exposition of a unitary pure investment model.

rese

that

simi

hou

McE

para

pare

the

char

wou

coop

this.

whe

solu

The

solu

deve

previ

they

alloc

[6] For a

reservation (or threat-point) utility levels of each party. Manser and Brown also show that the unitary model is nested inside this framework. McElroy and Horney (1981) use a similar framework but derive generalized comparative statics for the Nash-bargained household solution, for which the unitary model's comparative statics is a specific case. McElroy (1992) also discusses the comparative statics of extrahousehold environmental parameters, such as factors affecting the competitiveness of the marriage markets, parent's incomes, and others. These parameters determine the reservation utility outside the household, and therefore are functions of household demands. Her paper looks at the changes in household demands due to shifts in these parameters.

The papers above treated as threat points the utility that household members would receive outside the household, but this is not the only possible threat point. A non-cooperative Nash solution is a possible threat point. Lundberg and Pollak (1995) discuss this, and they point out that household members may retreat to non-cooperative strategies where each member is maximizing his utility subject to the other agent's strategy. This solution is not Pareto optimal, and both members would be better off by cooperating. The shared good, as would be expected, is also underprovided in the non-cooperative solution.

In papers by Chiappori (1992), and Bourguignon et al. (1993, 1994), the authors develop another model of household allocation. Their premise is different from the previous two models in that they do not explicitly specify a bargaining structure. Instead, they derive individual demands when all that is assumed is that household members' allocations are efficient. As it turns out, this framework is equivalent to assuming that

[6] For a review see Browning (1992).

househ
maxim
the sha
unitary

advanta
househ
income
betwee
slopes
one goo
to testi
exampl
ratio of
the true
assumin
The cha
change
second
allocatio
due to a
income

household members divide income according to a sharing rule in a first stage and then maximize their individual utilities subject to the sharing rule in a second stage. In effect the sharing rule is providing each member's weight in a household utility function. The unitary model is, again, nested in this model.

The framework developed by Bourguignon and his colleagues has an additional advantage: it allows recovery of some information about the sharing rule with only household-level data on consumption. With only household level information and income from household members, it is possible to derive the marginal rate of substitution between members' sharing rules. It is also possible to recover the ratio of Engle curve slopes of both members. Furthermore if one can observe individual consumption of just one good, the sharing rule can be recovered up to a constant.[7] The model also lends itself to testing both the validity of the unitary models and to testing its own validity. For example, if one can observe mother and father's incomes separately, one can compute the ratio of the demand of private goods with respect to each income. If the unitary model is the true model, this ratio should be unity because the unitary model is equivalent to assuming that household incomes are pooled—who owns the income should not matter. The change in demand due to a change in mother's income should be the same as the change in demand due to a change in father's income, up to sampling variability. The second test that can be performed is that of the pareto-efficient model itself. If household allocations are indeed efficient, then it must be true that the ratio of the change in demand due to a change in mother's income to the change in demand due to a change in father's income must be the same for all private goods. The framework described above has all

the

rea

for

doe

whi

---

[7] Alt
obse
[8] Th
whic
incor
incor
predi
be pr
Thes

the va
share
types
ratio
consu
strong

evider
the ch
incom
but th
model
pool in

educat
Thoma
same f
pooling
They t
likely t
outcom
in deter
these di

efficien
controll

allowed several tests of both the unitary model, through tests of income pooling, and of the efficient model.[8]

In the context of the unitary model, gender bias may originate for a number of reasons. The first reason is that households may have different preferences for boys than for girls. It is important to make a distinction here. Under this hypothesis, a household does not prefer a boy because he may bring the household future benefits (or lower costs) which a girl would not. The preference here is because of differences in how boys and

---

[7] Alternatively, if consumption of two excludable goods (goods that only one individual consumes) is observed, again, the sharing rule can be recovered up to a constant.

[8] Thomas (1992) uses Brazilian household data to test if unearned income of different household members, which is treated as exogenous, has the same effect on health outcomes. He finds that it does not, and that income under a mother's control affect child survival probabilities and weight for height much more than income under the control of the father. This is evidence against the unitary view of the households, which predicts that household income be pooled. However, the treatment of unearned income as exogenous may be problematic, to the extent that it represents past labor supply and human capital investment decisions. These decisions, in turn, may be driving part of the difference in the paternal and maternal effect.

Thomas and Chen (1994) use data from Taiwan's Personal Survey of income distribution to test the validity of the unitary model. They estimate Engle curves and reject the hypothesis that each member's share in income does not affect demand (i.e. they reject income pooling). Their results are robust to the types of income used (both earned and unearned). Furthermore, they cannot reject the hypothesis that the ratio of the change in the Engle curve due to changes in husband and wife's incomes is not the same across consumption groups. This is evidence in favor of the efficient model. Their findings in this respect are stronger for urban areas than for rural areas.

Bouguignon et al (1992) were the first to outline and test the collective model. They also find evidence against the unitary model, but reject conditions for their efficient model to hold. They find that the change in household expenditures on consumption items due to changes in household member's incomes is not the same for each household member (a necessary condition for the unitary model to hold), but that these differences are not the same across consumption items (a necessary condition for the efficient model to hold). This, again, is inconsistent with household members pooling their income. If members did pool incomes, these elasticities should be the same, regardless of the income's source.

In their article, Thomas, Schoeni and Strauss (1997) look at the effects of mother and father's education on sons and daughters. They also test for income pooling, and thus homogeneous preferences. Thomas et al first test if the difference in the effects of parental education on sons versus daughters is the same for mothers and fathers. They find that it is not. This is consistent with a rejection of income pooling, but is not a test of income pooling, as boys and girl's production functions may be parent specific. They then use a similar approach as above, but use parent's income, not education, since income is not likely to be in a production function. Here they also find that the difference in the parent's income in the outcomes of boys versus girls, or "difference-in-differences" is not zero. Income "ownership" is important in determining the relative effect of income on boys' and girls' educational outcomes. They also find that these differences are more pronounced for younger cohorts and for less educated households.

Udry (1996) uses data from agricultural production in Sub-saharan Africa to test for pareto efficiency. He finds that plots controlled by women in the household have lower yields than plots controlled by men. In a Pareto-efficient model plot control should not determine yield.

gir

to p

hav

rep

rea

co

hy

sh

co

ho

fu

inc

res

the

ma

dau

wit

exp

exp

[9] For
daug
it is p

girls enter the household's utility function. If this is the case, households will be expected to provide greater resources to a boy than to a girl.

If boys are expected to earn higher income than girls do, the birth of a boy will have two effects on the household. The first is an income effect. The birth of a boy represents higher future incomes for the household, and the household can be expected to react by increasing consumption. To the extent that households seek to smooth consumption over time, as is predicted by the Life-cycle and Permanent Income hypothesis, they will increase consumption immediately in response to a positive "gender shock". Also, to the extent that leisure is a normal good, household members may also consume more leisure and reduce their labor supply, although this may well depend on how each parent's time enters into their sons and daughters' human capital production functions.[9] It is important to note, however, that even if households cannot appropriate incomes from their offspring, if they are altruistic they may also decide to allocate resources according to the gender of a newborn. For example, if households care for their children's well being after they leave the household, and if there are indeed higher market returns to investing in a boy, households may want to make more transfers to daughters than to sons. This would arise because the higher marginal utility associated with lower incomes for a daughter would prompt parents to make unequal transfers. In expectation of these higher transfers they may increase labor supply and decrease expenditures.

[9] For example, if mothers are better at transmitting human capital to daughters, perhaps the birth of a daughter will lead the mother to spend less time working and more investing in her daughter. In this case, it is possible that the marginal product of her time will compensate for the income effect on her leisure.

is a s

hous

This

hous

It ma

into

boys

retur

done

marr

there

from

barg

diffe

alloc

mem

vers

Brow

[10] See

The second effect one can expect if boys face greater earnings in the labor market is a substitution effect. Boys may represent a productive asset for a household, so households may substitute from other assets into boys more than they would into girls. This will typically take the form of more schooling, health care, or in the case of households near poverty—more food provided to a boy than would be provided to a girl. It may also involve more time spent with children, to the extent that parent's time enters into the child's human capital production function.

A third way in which gender bias may originate is through differential costs for boys and girls. This is parallel to the income effect associated with differential market returns for boys and girls. Recent studies on gender bias and differential costs have been done in India and have focused on marriage costs, which are higher when daughters are married off.[10] However, since there is no social institution such as a Dowry in Brazil, there is little reason to believe that girls will represent higher costs to households, apart from an income effect due to differential returns to girls.

Even if the true model of household behavior is not the unitary model, but, say, a bargaining model, many of the results described above will still carry through. The key difference is that in an individualistic model, resource control will have an affect on the allocation of goods within a household after the birth of a boy or a girl. And if household members do have different preferences, possibly including different preferences for boys versus girls, resource control should affect any possible gender bias.

In the context of the Indian ICRISAT data, both Deolalikar and Rose (1995) and Browning and Subramaniam (1995) have analyzed the effect of a "gender shock" on

---

[10] See Browning and Subramaniam (1995).

hou

as th

boy

a de

time

on p

imp

and

The

farn

betv

and

ente

girl

that

abo

The

shou

cons

only

ICR

retur

household consumption and production decisions. The "gender shock" can be thought of as the realization of the sex of a child. In Deolalikar and Rose's analysis the birth of a boy, relative to that of a girl, is associated with an increase in household consumption and a decrease in production income. Since ICRISAT follows rural Indian villages across time, the authors can model savings, consumption and production income as depending on present and past fertility decisions. The authors' results, although estimated imprecisely, suggest that households increase consumption the year following the birth, and that they decrease production in the year of the birth, and for subsequent years. These effects are more significant for large landowners than for the landless or for small farm owners.

Also using ICRISAT data, Browning and Subramaniam attempt to distinguish between three possible motivations for the observed relationship between consumption and income (and thus savings), and the birth of a boy, relative to that of a girl. They entertain three underlying motivations that explain why households save more when a girl is born. The first is that rural Indian households prefer boys to girls, the second is that there are higher net returns associated with having a boy (these two discussed above), and the third is that there are higher marriage costs associated with having a girl. They show that if households have preferences for boys, consumption of adult goods should fall after a boy is born, as households substitute to goods that the boy will consume. Overall consumption should also increase. Because consumption with adult-only goods rises with the birth of a boy, the authors conclude that the evidence from the ICRISAT villages is consistent with both the marriage cost motive and the higher net returns motive. If boys represent a more productive asset, households would have higher po

co

hi

po

re

lo

bo

go

th

ex

tha

the

red

a b

equ

but

sign

How

here

[11] Th
numb

permanent income after the birth of a boy, and thus would be expected to increase consumption of all goods, including adult goods. But this is also true if girls represent higher costs than boys. Then, if a boy is born, households again will have higher permanent income because they will not face a Dowry payment in the future. These results do not hold for households with small land holdings.

Deaton (1989) takes a different approach to the issues of gender bias. Instead of looking at the birth of the last born, Deaton looks at the relative effects of an additional boy versus an additional girl on household overall consumption and consumption of adult goods. Using cross-section data he constructs outlay equivalence ratios, which are just the ratio of a composition elasticity to an income elasticity.[11] Treating children as exogenous, he shows that an additional child in the household brings in new needs, so that the additional child acts like a reduction in income. In the absence of gender bias, the ratio of the reduction in the expenditure on adult goods due to a new child to the reduction in expenditure due to a change in income should be the same, weather a girl or a boy is born. He breaks the groups down into age-gender cells and compares the outlay equivalent ratios using the 1985 Living Standards Survey of Côte d'Ivoire and Thailand, but finds that the differences in outlay equivalent ratios for boys and girls are not significant for any of the age groups.[12]

One of the frameworks used in this paper follows much of Deaton's framework. However, whereas previous work in the area has looked only at consumption and income, here we also look for labor supply responses to child gender. First, we will test for the

[11] The composition elasticity is just the percentage change in expenditure due to a percentage change in the number of individuals in a particular group (boys vs. girls in this case).

relative

at the ag

to test fo

part of t

boys and

the obse

analysis

III.

1999 Pe

surveys.

labor su

are only

is also d

results.

well as

such as

covers t

individu

resident

[12] The En
household
[13] Note th

relative effects of an additional boy or girl on household consumption. This will be done at the aggregate level and by age and education groups. We use Deaton's methodology to test for appropriate adult goods and to test for demographic separability. In the second part of the paper we look at labor supply. In particular, we test for differential effects of boys and girls on father and mother's labor supply. Hypotheses that are consistent with the observed patterns are then discussed. All of the analysis will be done by regression analysis, the data and methodology being described in detail below.

III. Data

We use the 1987-88 *Pesquisa de Orçamentos Familiares* (POF) and the *1981-1999 Pesquisa Nacional por Amostra de Domiíilios* (PNAD), both Brazilian household surveys, to measure the effect of the gender of a newborn on household consumption and labor supply decisions.[13] Although the PNAD cross sections span nineteen years, there are only seventeen surveys, since surveys were not conducted in 1991 or 1994. Analysis is also done with the 1988 PNAD alone, for greater comparability with the 1987-88 POF results. The POF contains detailed information on household expenditure, income, as well as some personal and demographic characteristics pertaining to household members, such as age and education. The survey was conducted between 1987 and 1988, and covers the 11 largest Brazilian metropolitan areas. It contains observations on 56,569 individuals from 12,959 households (for an average of 4.4 people per household). Every resident in selected households is interviewed.

---

[12] The Engle curve specified is an extension of Working's (1943) functional form, changed to include household demographics.
[13] Note that PNADs were not fielded in 1991 and 1994.

Consumption information is available on a number of categories. We use consumption with three categories: education, health, and adult goods. The first two are of particular interest because they represent, in part, investments in human capital. The last is of interest because it may be an assignable good, consumed only by the adults of the households. We also compute overall household consumption.[14] Consumption is measured by two methods. The first is a "notebook" in which the household annotates consumption of every-day items. The second is consumption measured by four- and six-week recall. These are for items purchased less frequently. All monetary values are deflated to October 1987 Cruzados (Brazil's currency at the time).

The PNADs are surveys of much larger scale. Whereas the POF only covers the 11 largest metropolitan areas, the PNAD is a representative sample of the entire country, both rural and urban. The only areas that are not covered are selected rural areas of the North. Since POF does not have questions about labor supply, we use the PNADs for our labor supply analysis and the POF for our consumption analysis. In both cases our interest is in households with both a mother and a father present. All observations with missing age or education values were discarded, as were all households with single parents. This leaves a total of 9,314 households in the POF dataset. For the seventeen cross sections of the PNAD the sample sizes vary, the smallest survey being 1986 (65,446), the largest besing 1989 (119,651). Even though we use "mother" and "father" throughout, couples without kids are also part of our sample.

[14] Since some households do not rent (they may own, or perhaps live in a relative's house), or do not report any rent, a value is imputed for those households. The rent is predicted by regressing the rent from those who do rent on the number of rooms in the household, the number of bedrooms, dummy variables for the type of water supply (no water, well water, water duct, city water), dummy variables for the type of sewer (no sewer, city sewer, county sewer), as well as a dummy variable for the interviewer's classification of the house (rustic, temporary, permanent no floor, permanent cement floor).

t

a

c

h

t

p

w

T

th

A

co

pe

ac

en

av

wi

15

qua

gra

col

Table 1.1 presents summary statistics of the variables used in the analysis, using the 1987 POF and 1988 PNAD. It is important to notice that even though the two analyses are done using household survey data, these two populations are quite dissimilar. The urban POF population is more educated: 19 percent of the POF husbands have at least one year of college (13 percent for the wives), whereas only 12 percent of the PNAD population do (9 for wives). On the other side of the spectrum we see that 22 percent of the husbands in the PNAD report having completed less than primary school, whereas this number is 10 percent for the POF.[15] The results are similar for mothers. This is consistent with the POF being a survey that covers only the main urban areas of the country, which are typically wealthier and more educated.

We can also see in Table 1.1 the breakdown of consumption, by category. Amounts are given on the lower left column and percentages on the right. Food consumption accounts for 21 percent of total consumption, healthcare for only six percent, and education for even less: three percent. We also see that expenditures with adult-exclusive goods account for roughly 8 ½ percent of all household consumption.

From the PNAD we see that 93 percent of the husbands in the sample are employed, as opposed to only 41 percent of the wives. Mothers who work also work on average eight fewer hours per week than fathers do. Finally, we see that the proportion of working mothers that is self-employed is similar to that of fathers.

---

[15] Some of this difference may be due to different education questionnaires. The POF education questionnaire asks if the respondent completed or not primary middle, high school, or if he has a college or graduate degree. The PNAD questionnaire asks the highest grade and school (primary, middle, high, college) completed for both those in school and those not in school.

IV.

Lab

asso

gen

hou

oth

by

wo

ob

en

in

po

ol

so

e

f

e

r

i

i

IV. Model

Labor market differences

One of the strongest arguments for differential gender effects is that a boy is associated with a higher expected future income. It is thus useful to first describe any gender differences in the Brazilian labor market. If they do exist and are prominent, then households may indeed expect to have an "income" effect when a boy is born. On the other hand, if they do not exist, then the "income" effect argument loses validity.

Table 1.2 describes basic gender differences in Brazilian labor markets for 1988, by age and education level. In the right panels one finds employment rates for men and women; in the left, average monthly earnings (in 1987 Cruzados). These figures are obtained from the same 1988 PNAD that is used later in the analysis. We clearly see that employment rates are higher for men than for women. The difference tends to slowly increase, from a low of 30 percentage point difference to a high of about 50 percentage points for the 55-64 age group, before falling to 25 percentage points for those 65 and older. If we consider only lesser-schooled individuals—those with less than primary schooling—this increase is not so pronounced. However, for this same group employment rate differences are even larger (about 10-15 percentage points larger than for the full sample). Some of the increase in the difference in male to female employment rates is because women tend to retire earlier;[16] therefore one will have many more retiring women in the 45-54 and the 55-64 age groups. This effect is then reversed in the 65 and older group, as the men retire. Another possible explanation is that the increase in the difference of male to female employment rates is largely capturing cohort

[16] In fact, for purposes of social security, Brazilian law allows women to retire five years earlier than men.

effects, women entering the labor force in the seventies and eighties. Since only one year is shown, one cannot separate the cohort from the age effect.[17]

Table 1.2 also shows average earnings for men and women. We see that for all age groups men's earnings are substantially higher than women's, ranging from 25 to 60 percent. We also see this difference increasing monotonically with age. Again, the analysis is replicated for those with less than primary schooling. The gap for the lesser-educated is larger, except for the oldest groups, and shows a similar increase with age group. The table suggests that differences in labor market participation and earnings favor men, and that these differences increase with cohort (or with age), the older cohorts (older men) having the largest differences. We also see that these differences are largest among the lesser-educated.

There are two questions of interest here. First, what is the effect of a child's gender on household consumption and labor supply decisions, and does this effect vary with household (and individual) characteristics and previous fertility outcomes? Second, why do we observe this effect?

Household behavior may be different when faced with the birth of a boy versus a girl for a number of reasons. In the context of a utility maximization model households may allocate resources differently because they prefer one gender to the other (say, boys to girls) or because one gender represents a higher net return. In terms of net returns, a boy may be associated with a higher expected income, because of differences in labor force participation rates, differences in hours worked, or differences in wages—this

[17] With a panel one could separate age and cohort effects (albeit not age, cohort and period effects). With a single cross-section one can only identify age or cohort effects, as older people will also belong to older cohorts.

woul
beca
for b
pref
hou

adu
inc
mc
lat
in

m
H
th
c
r
c

would be an income effect. Alternatively, households may allocate resources differently because the returns to education, health, or even nutrition investments may be different for boys than for girls. Furthermore, as was mentioned previously, different "household" preferences for boys and for girls may reflect different preferences of individual household members

If households simply prefer boys to girls, one would expect consumption with adult goods to decrease after the birth of a boy and consumption with other goods to increase. We would also expect labor supply to increase, as household members work more hours to increase the boy's consumption. However, households may also decrease labor supply and increase leisure, to the extent that their gender bias may be through increased enjoyment of leisure when a boy is born.

If households respond to the gender of a child simply through the pure investment model, one would expect a boy to increase the household's expected future income. Households would then be expected to react to the birth by increasing consumption (since their permanent income is now higher), including consumption with adult goods and decreasing labor supply (if leisure is a normal good). However, since a boy may represent a more productive asset, households may also increase expenditures with education and health care. Health care expenditures may be expected to increase immediately. However, expenditures with education may only be increased when a child reaches school age.[18]

[18] It is important to note that school fees, when present, tend to be minimal in Brazil, and almost all public urban schools have no fees at all. Similarly, public universities are free, although the entry requirements can be stringent. Thus, expenditures may be in the form of outlays with private (higher quality) schools and more and better supplies. However, costs of school attendance which are not tabulated include the opportunity cost of the child not working, for poor families, as well as the parental costs associated with schoolwork.

di
ef
da
ar
m
re
di
di
te
pr
Th
de

Ou

ma
De
san
me
acc
way
cha

If husbands and wives bargain for control of household resources, and they have different preferences for boys and girls, one would expect that some of the gender bias effect would depend on household members' individual incomes (in so far as that will determine their relative bargaining position in the household). Factors such as education and age could also determine relative bargaining positions. Therefore, more educated mothers may be able to affect the sharing rule; for example, it might increase her reservation utility outside the household. If this is true, and if husbands and wives have different preferences, households with more educated mothers may allocate resources differently after the birth of a boy and the birth of a girl. This is, however, by no means a test of the unitary model. It is perfectly possible for husbands with different child gender preferences to also have different preferences for the education levels of their spouses. This could explain a correlation between gender bias and education levels (or any other demographics, for that matter).

Outlay equivalency ratios

In this section we test for the validity of adult goods, which will be used in the main analysis. We will also test for gender bias using the methodology developed by Deaton (1989). The approach is to find a group of goods, all of which respond in the same way to the addition of a child. That is, when a child enters the household that member will bring about a change in consumption of different types of goods. An acceptable set of adult goods is one in which expenditures on each good react in the same way to an additional child as they would to a decrease in income. In this way, the only change in consumption of that good brought about by the child is through an income effect—there are no substitution effects. Goods which satisfy this property are said to be *demographically separable* from children. Deaton, Ruiz-Castillo and Thomas (1989) develop preferences which are consistent with demographic separability using a unitary model framework. Here we will be somewhat more general and, following Deaton (1997), show demographic separability using the efficient bargaining model developed by Chiappori, Bourguignon and Browning.[19]

Bourguignon and his colleagues derive a system of demand equations from a pareto-efficient bargaining outcome in which each party maximizes his utility subject to a sharing rule. For example, suppose there are N agents. Agent j will

(1) $\max\ U^j(x,g)\ s.t.\ p^j x = \theta^j(p, p_g, y, n^C)$

where x is a vector of goods consumed by agent j, g is the optimally provided public good, and p is the price vector for all goods. Theta is the sharing rule, which allocates each agent's share of total outlays. For our purposes, there are two types of groups, children (j=C) and parents (j=P). Let n^C represent the number of children. Note that we can trivially redefine parents as adults, and include the number of adults in the sharing rule, however, for simplicity we are only allowing the number of children to vary, and thus the number of parents fixed. We are also assuming that all income accrues to parents. The solution to this maximization problem yields a set of demand equations for parents:

(2) $x_i^P = f_i\left(\theta^P(p, p_g, y, n^C), p, g, n^C\right)$

The budget constraint in (1) implies that

(3) $\theta^C(p, p_g, y, n^C) = y - p_g g - pq^P$

---

[19] See Chiappori (1992) and Bourguignon et al. (1994).

Now define a group of adult goods, q^A ,such that these goods satisfy demographic separability with respect to the number of children. This implies that the number of children only has an income effect on adult goods (through the sharing rule), and thus is not an explicit price effect in the demand function. Therefore, omitting the public goods for simplicity, we have adult goods demands:

$$(4) \qquad x_i^A = f_i^A\left(\theta^P(p, y, n^C), p\right)$$

For each adult good demand function we can calculate two partial derivatives, on with respect to the number of children and the other with respect to income:

$$(5) \qquad \frac{\partial x_i^A}{\partial n^C} = \frac{\partial x_i^A}{\partial f_i}\frac{\partial f_i}{\partial \theta^P}\frac{\partial \theta^P}{\partial n^C}, \quad \frac{\partial x_i^A}{\partial y} = \frac{\partial x_i^A}{\partial f_i}\frac{\partial f_i}{\partial \theta^P}\frac{\partial \theta^P}{\partial y}.$$

Notice that the ratio of these partial derivatives,

$$(6) \qquad \frac{\dfrac{\partial x_i^A}{\partial n^C}}{\dfrac{\partial x_i^A}{\partial y}} = \frac{\dfrac{\partial \theta^P}{\partial n^C}}{\dfrac{\partial \theta^P}{\partial y}},$$

<u>does not</u> depend on i and therefore is the same for all adult goods. This allows us to test for an appropriate (i.e. demographically separable) set of adult goods: goods in which the ratio of demographic effects to income effects are the same.

In the preceding discussion we dealt with demand equations, which means they had prices as their arguments. In a cross-section there is only negligible variability in prices, so the results above will be applied holding prices fixed. This does not change any of our results since the ratios in (6) do not depend on prices. In fact, if we estimate expenditure shares rather than demand functions (which is what we do), the results in (6) still l

canc

expe

(7)

x_i is

com

adul

num

grov

acr

onl

(8)

De

(9

Sc

(1

Ir

T

still hold, since the prices will appear in both the denominator and the numerator and will cancel out.

Redefine x_i as expenditures with good i. Now consider a function which relates expenditures with good i to total expenditures, and household characteristics.

(7) $$x_i = f(x, n, z)$$

x_i is expenditure with good i, x is total outlays, n denotes a vector of demographic composition, and z a vector of household characteristics. According to the definition of adult goods above, the ratio of the change in expenditure in good i due to a change in the number of members in the demographic group r (here r can index both sex and age groups combinations), n^r to that good's marginal propensity to spend should be the same across all adult goods. That is, a group of goods $x_1, x_2, \ldots x_a$ is a set of adult goods if and only if:

(8) $$\frac{\frac{\partial x_1}{\partial n^r}}{\frac{\partial x_1}{\partial x}} = \frac{\frac{\partial x_2}{\partial n^r}}{\frac{\partial x_2}{\partial x}} = \ldots = \frac{\frac{\partial x_a}{\partial n^r}}{\frac{\partial x_a}{\partial x}} \quad \text{for all children categories.}$$

Deaton defines an outlay equivalency ratio (*oer*) as

(9) $$oer_{ir=} \left(\frac{\frac{\partial x_1}{\partial n^r}}{\frac{\partial x_1}{\partial x}} \right) \frac{n}{x}$$

So that the definition of adult goods can be rewritten in terms of outlay equivalency ratios

(10) $$oer_{1r} = oer_{2r} = \ldots = oer_{gr} \quad \text{for all child categories.}$$

In order to estimate an *oer*, we follow Deaton in our specification of an Engle curve. This function specifies a linear relationship between the share of expenditure with each good and th

by Working

(11) $\frac{x_i}{x}$

Here n den

Note that c

adults. Th

13-15. Fro

(12) oe

The stand

goods fro

The first

these thre

goods, w

needs fo

The esti

Thailanc

Deaton'

clothing

for the e

good and the logarithm of total expenditures. It is the same type of Engle curve proposed by Working (1943).

$$(11) \qquad \frac{x_i}{x} = \alpha_i + \beta_i \ln\left(\frac{x}{n}\right) + \eta_i \ln(n) + \sum_{1}^{J-1} \gamma_{ij}\left(\frac{n_j}{n}\right) + \delta_i z + u_i$$

Here n denotes total household size and n_j denotes the number of people in category j. Note that only J-1 groups are estimated. The omitted group, group J, is the group of adults. The six other groups are boys ages 0-5, 6-12, and 13-15; girls ages 0-5, 6-12, and 13-15. From (11) above one can compute *oers* as follows

$$(12) \qquad oer_{ir} = \frac{\eta_i - \beta_i + \gamma_{ir} - \sum_{1}^{J-1} \gamma_{ij}\left(\frac{n_j}{n}\right)}{\beta_i + \frac{x_i}{x}}.$$

The standard errors for (6) are computed using the delta method.[20] Three groups of goods from the POF were tested as possible adult goods: tobacco, beer and adult clothing. The first part of Table 1.3 presents the estimated outlay equivalency ratios for each of these three goods and for each of the children categories. The *oers* are negative for all goods, which is what one would expect since the additional child would bring in new needs for the family, and since we do not expect children to consume these adult goods. The estimated *oers* are in general larger than those estimated by Deaton using data from Thailand and Côte d'Ivoire, and present less variation between demographic groups than Deaton's. The *oer*s for tobacco and beer are very similar, but not the *oer* for adult clothing. Two Wald tests for the equality of the *oers* are performed. In the first we test for the equality of all three adult goods, the second we test for the equality of tobacco and beer. We reject the equality of all three *oers* in all but one of the demographic groups, but fail to reject the equality of beer and tobacco *oers* in all but one demographic category.[21] Hence, tobacco and beer are a more suitable group of adult goods than all three goods combined.

In the absence of gender preferences, an additional child should affect the consumption of adult goods to the extent that said child acts as a reduction in income. So, one way to test for gender bias is to test the equality of *oers* for boys and girls for the different categories of adult goods. Note however that a rejection of their equality does not necessarily mean that households have different gender preferences, as boys may bring different needs than girls.

The bottom part of Table 1.3 presents the difference in *oers* for boys and girls, as well as the standard error for this statistic. In all but one case the addition of a boy acts like a larger income shock than the addition of a girl. That is, households reduce expenditures with adult goods more when a boy is born than when a boy is girl is born. However, these differences are never significantly different than zero, even at the 10% confidence level. Hence, we can conclude that there is no evidence of gender bias in consumption.

---

[20] See Deaton, Ruiz-Castillo and Thomas (1989) for a description of the formulas used to calculate the standard errors and for those used to calculate the Wald statistic.

[21] Note that we cannot disqualify adult clothing as an adult good. It may be that adult clothing is the adult good and tobacco and beer are not. And in general, strictly speaking this definition of adult goods does not necessarily imply that an additional child brings only an income effect. It is consistent with goods with similar substitution effects also.

Model for the main analysis

In this section we describe the model for the main analysis. One possible empirical specification would be the following:

(13) $\left(E(y_{it+j,\,boy})-E(y_{it,\,no\ child})\right)-\left(E(y_{it+j,\,girl})-E(y_{it,no\ child})\right)$, *for j=1...T*

Were *y* denotes consumption and labor supply, *i* denotes the household, and *t* denotes the year. This can be thought of as an experiment in which households are chosen to have children and randomly assigned a boy or a girl. Their consumption and labor supply is observed before the child is born and then compared 1,2,...,J years after the child is born. However, estimation of this model is not possible. First, we do not have a panel data, and therefore cannot observe the same household over time. Second, we cannot control a household's decision to have a child, much the less how many children to have. This is a selection problem. The households who do have children may not be a random sample of our population of candidate households. Lastly, even if the decision to have a child were exogenous, the decision of when to stop having children may not be. That is, if households have preferences for boys, their probability of having an additional child may be closely related to whether or not they already have a boy. The same could hold, for example, if households prefer a "gender balanced" household, and are more likely to have another child if it has only boys or only girls. Therefore, we cannot guarantee that observed gender effects hold for households who did not decide to have a child. However, if fertility decisions are correlated with observable household characteristics, such as age and education variables, then it may be possible to match households who decided to have children with those who did not. That is, we can instead attempt to estimate

(14) $E\left(y_{it+j,boy} \mid Z\right) - E\left(y_{it+j,girl} \mid Z\right)$, *for j=1...T*

In equation (14) we are controlling for the level effects of demographic variables Z. In addition, we may compare slope and level effects by estimating equation (14) by education and age groups. That much said, we can never be certain that all covariates related to fertility decisions have been taken into account—one must still exercise care in generalizing gender bias results (if found) to the population as a whole.[22]

The method used to estimate (14) above is to estimate consumption and labor supply equations:

(15) $y = \alpha + \sum_{j=1}^{5} \beta_j nboy_j + \sum_{j=1}^{5} \delta_j ngirl_j + Z\gamma$

Here $nboy_j$ and $ngirl_j$ are the number of boys and girls in category *j*, respectively, where *j*={0-5 years old, 6-9 years old, 10-11 years old, or 12-13 years old and 14-15 years old}. For the labor supply results j={0-9 years old, 10-11 years old, 12-13 years old, or 14-15 years old}.[23] *Z* is a vector of mother's attributes, including schooling, age, age squared,

[22] For the consumption equations, five education categories were created: less than primary school, completed primary but less than completed middle, completed middle but less than completed high, completed high, more than high. Less than primary was omitted in the estimation. For the labor supply equations a fully flexible dummy variable specification was used, with "no schooling" omitted. Other demographic controls include father and mother's age, father and mother's age squared, household size, and a dummy for whether the household is headed by the father (dummy equals 1) or the mother (dummy equals 0).

[23] In earlier estimations the two youngest groups were estimated separately, but results were very similar (and insignificant). Therefore they were concatenated into one category.

and household size.[24] Our measure of gender bias is then $\beta_i - \delta_i$. Their equality is tested for with either an F-test (for OLS equations) or a Chi-squared test (for probit regressions). Here we should note that, even if number of boys and number of girls is endogenous, and therefore β_i and δ_i may be biased, most of this bias is likely common to both boys and girls, and will disappear in the difference in the coefficients.

OLS equations are estimated for consumption's share of total expenditures, education's share of consumption, health care's share of consumption, and adult goods' share of consumption. That is, consumption divided by total expenditures, education expenditures divided by total consumption, etc. These equations are also estimated separately for the entire sample and for the lesser-educated sub-sample, where lesser-educated is defined as those with less than primary school (21 percent of mothers in our 1999 sample). Results are presented in Tables 1.4 and 1.5.

The employment equations were estimated with a probit regression, and the marginal effects are reported on Tables 1.6 and 1.7. These estimates are done for the 1988 PNAD, for comparability with the 1987 consumption results, and with the full set of the PNADs. Since one has reason to believe that households in the poorer regions of Brazil (north and northeast) may be quite different than households in the richer regions of Brazil (south, southeast), region dummies were included in the 1988 PNAD results. In the full sample of the PNADs, we use both region and year dummies, since heterogeneity

[24] Since the POF has only credential schooling information, we are unable to construct a completed years of schooling variable. Therefore, for the consumption regressions, we use as schooling categories: less than elementary (omitted dummy), completed elementary, some middle, completed middle, some high school, completed high school, more than high school. One should note that the schooling system gradually changed in Brazil, starting in the late sixties, in which a three-tiered schooling system (elementary, middle, high), was replaced with a two-tiered system. The first tier absorbed what used to be first eight years of schooling and became the "1^0 grau." The "2^0 grau" then became roughly equivalent to

may arise between households of different survey years. This also accounts for changes in macroeconomic conditions over years that might also affect labor supply.

Six labor supply models are estimated for mothers, all of them probit models. In all models the labor supply variable is whether or not the mother reported work as her main activity during the reference week. In the first labor supply model employment is regressed on mother's education variables, mother's age, age squared, household size, and child status variables. Since there were no differences in the age categories 0-5 and 6-9, these were combined into one category. A dummy variable, *child*, for whether the household had children was also included.

Our first probit model is simply the probit of equation (15) above, allowing for a different intercept for households without children.

$$(16) \qquad p(y=1) = \Phi\left(\alpha + \sum_{j=1}^{4} \beta_j nboy_j + \sum_{j=1}^{4} \delta_j ngirl_j + \theta child + Z\gamma\right)$$

Boys and girls may have different employment rates and different school enrolment rates. If wages for boys are significantly higher than wages for girls, one can expect boys to participate in the labor market to a greater extent as they enter into their teens. They may also enter the labor market earlier than girls. This in itself may affect mother's labor supply as household may substitute into child labor and away from mother's labor, the more boys there are in the household. Therefore, in the second probit model we include the proportion of children at work as a control variable.[25] In the third model we include

---

the old high school and encompassed the remaining three years (grades 9-11). For the labor supply regressions we use a flexible functional form or seventeen completed years of schooling dummies.

[25] We also use alternative definitions of child work, including children who reported working ten or more hours per week. In addition, we estimate the model by using only the respondent's own children in the proportion. In all cases estimates were similar. All child labor proportions are for children ten years or older since no labor supply information is asked of young children.

the proportion of children who reported to have "worked at home" as their primary activity in the reference week. If girls help in household chores to a greater extent than boys do (as they get older), mothers may substitute away from home care, the more girls there are in the household. In the last model we include both proportions. Thus, *Model 2 – Model 4* can be summarized:

$$(17) \qquad p(y=1) = \Phi\left(\alpha + \sum_{j=1}^{4}\beta_j nboy_j + \sum_{j=1}^{4}\delta_j ngirl_j + \theta_0 child + \eta_k p_k + \theta_k p_k x child + Z\gamma\right),$$

where $k=1$ for proportion of children at work (*Model 2*), and $k=2$ for the proportion of children doing housework (*Model 3*). These proportions are endogenous, since boys and girls are not randomly assigned to work. However, since there are no good instruments to deal with this endogeneity, we will present the results and interpret them with this in mind. In any case, even if η_i and θ_i are biased (and therefore β_i and δ_i will probably be biased also) there is no *a priori* reason to believe that this bias will be present in the difference in boy-girl coefficients, $\beta_i - \delta_i$. In Tables 1.6 and 1.7 child status variables are reported, together with the child activities proportion variables discussed above. Note that *Models 1-4* are estimated with regional dummies.

A second issue concerns expected future employment of boys and girls. Households may react not to the current higher labor force participation of boys relative to girls, but to expected future labor participation and earnings. We will not postulate that households make rational expectations about future employment and earnings levels. Instead, a more likely expectations mechanism is that households make expectations about the next five or ten years based on current observed levels of employment and wages. To account for this in *Model 5* we computed the average ratio of boy to girl wages for children 10-15, and for young adults, ages 18-24, by region (north, northeast, central east, south, southeast), urbanization (major urban, minor urban, rural) and year. The hypothesis is that if households do react to expected future labor force differences, we may expect to see larger labor force reactions for households in cells whose *current* differences are large. In effect, we are comparing the reaction of households in high difference cells to those in low difference cells. We do this by defining a high difference cell dummy variable, equal to one if a household is in a high difference cell and zero otherwise. We then interact this dummy variable with the child composition-gender variables. Three thresholds are used for the high-difference dummy variable, both based on the percentile distribution of the ratio of boy to girl wage. The first threshold corresponds to the 50th percentile (i.e., household is in high difference cell if that cell's percentile is at or above the 50th), the second to the 70th, the third to the 90th percentile. Results presented in Table 1.8 correspond to the first threshold, but were similar for all three.

$$(18) \qquad P(y=1)=\Phi\begin{pmatrix} \alpha+\sum_{j=1}^{4}\beta_j nboy_j+\sum_{j=1}^{4}\delta_j ngirl_j+\theta_0 child+\eta hidif \\ +\sum_{j=1}^{4}\kappa_j nboy_j hidif+\sum_{j=1}^{4}\lambda_j ngirl_j hidif+Z\gamma \end{pmatrix}$$

From this model we can compute a few statistics of interest. The first is the gender bias statistic, controlling for current labor market differentials,

$E\left(\frac{\partial\Phi}{\partial nboy}-\frac{\partial\Phi}{\partial ngirl}\right)=\left((\beta_j-\delta_j)-(\kappa_j-\lambda_j)\times\overline{hidif}\right)\, x\ \phi(.)$, where the high difference dummy can be evaluated at the sample mean (either 0.5, 0.7 or 0.9, depending on the threshold used). $\phi(.)$ is the density function evaluated at the sample mean. The second is the difference in this statistic between high and low differences cells,

$$E\left(\frac{\partial\Phi}{\partial nboy}-\frac{\partial\Phi}{\partial ngirl}, highdif=1\right)-E\left(\frac{\partial\Phi}{\partial nboy}-\frac{\partial\Phi}{\partial ngirl}, highdif=0\right)=(\kappa_j-\lambda_j)\,x\,\phi(.).$$

The last model to be estimated, labeled *Model 6* in Table 1.9, is an attempt to identify the role of child care availability in gender bias results. In particular, if girls do substitute for mother's household activities, including child-rearing, one would expect gender bias results to particularly strong in households where there are no older relatives to take care of children, and attenuated in households where older relatives can take care of children. In *Model 6* we compare gender bias results for both households using a difference-in-differences estimator similar to the one in *Model 5*.

$$(19)\qquad P(y=1)=\Phi\begin{pmatrix}\alpha+\sum_{j=1}^{4}\beta_j nboy_j+\sum_{j=1}^{4}\delta_j ngirl_j+\theta_0 child+\eta childcare\\ +\sum_{j=1}^{4}\kappa_j nboy_j childcare+\sum_{j=1}^{4}\lambda_j ngirl_j childcare+Z\gamma\end{pmatrix}$$

Here *childcare* is a dummy variable for the presence of older female relatives, ages 20-65, in the household.[26] The first difference and the difference-in-difference statistics are computed in the same way as (18).

In addition to the tabular results discussed at length below, we present predicted employment rates from a probit model for different child age categories—and for boys and girls separately. The specifications used for Figures 1.1 – 1.4 closely follow *Model*

[26] We also used as controls in alternative specifications (a) the presence of older female relatives 55-65 years old and (b) older unemployed females ages 20-65. The magnitudes of the bias effects in (b) were similar and more significant. The magnitudes of the bias effects for (a) were smaller, but still present.

$1 -$ M

are e

hous

IV.

First

resu

expe

bein

asso

con

one

larg

hyp

join

is n

par

27 C
mod
boy
nun
than

1 – Model 6.[27] However, for the difference-in-difference models separate regressions are estimated for each of the differences sub-samples (high and low wage areas and households with and without available childcare).

IV. Results

Results are highlighted in this section, and presented in full in the appendix tables. First consumption results, found in Tables 1.4 and 1.5, will be discussed. Labor supply results (Tables 1.6, 1.7 and 1.8, 1.9) will follow.

In Table 1.4 we find that controlling for household size, consumption's share of expenditures decreases with both the number of boys and the number of girls, this result being larger for younger children. Higher father and mother's education tend to be associated with higher overall consumption shares. This result holds for the share of consumption categories also. This suggests that children consume less than adults, which one would expect. Older households (measured by the age of the father) also consume larger shares than younger households do. This result does not hold for women's age.

What is absent from Table 1.4 is any discernable gender effect. The null hypothesis for the equality of boy and girl variables is rejected in every single case. The joint test for significance also rejects gender bias in consumption in all instances. There is no evidence of any convincing gender effect in the consumption tables. And in particular, we find no evidence whatsoever of any change in household consumption of

---

[27] One should note that the empirical specification used for the graphs is slightly different than the ones for *model 1 – model 6.* In particular, the model is more flexible, as fifteen different age categories are used for boys and for girls. However, for the graphs, these age categories are actually dummy variables, and not the number of boys or girls in each category. In most cases this is irrelevant as few households will have more than one boy or girl of the same age.

adult goods. If households prefer boys to girls (or vice versa) there should be a difference in the adult goods category for boys and for girls. In addition, if boys or girls represent net wealth gains associated with future incomes, there should be a difference in the effect of boys and girls on overall consumption category. None of this is evidenced in Tables 1.4, or Table 1.5.

Looking at Table 1.6 we can compare the labor supply results using both the 1988 PNAD and the full sample of 1981-1999 cross sections. The specification used for this table was *Model 1*. The results for both time periods are similar, if more precisely estimated for the full sample (as expected). Although not reported, effects are not present for fathers, who have inelastic labor supply. For mothers, labor force participation increases almost monotonically with schooling level, for the first years of schooling. This holds true for the first two school levels—primary and middle—but not for high school. Mothers who complete high school (grade 11 in Brazil) are eleven percent more likely to work. These large returns to higher levels of education are widely documented in Brazil. The more young children in the household (ages nine and under), the less the mother's probability of being in the labor force. Presumably those mothers would more often than not be needed at home to help with child-rearing activities. For the older age groups the result reverses itself: mothers now are more likely to be in the labor market in households with more children ages 10-11, 12-13 or 14-15. As we will see below this reversal is much more significant for girls than for boys. It is plausible that these children substitute for mother's labor efforts in the house. Mother's age also affects employment in the usual hump-shape. Again, results are similar for both time periods.

gender

work (

boy ag

boys r

previo

highe

Mode

work

work

the c

three

reac

moth

than

Here

coef

help

to tw

type

hous

[28] Wa

If the first set of labor supply results were somewhat expected, the ones relating to gender bias were not. In Table 1.7, for *Model 1*, we see that mothers are less likely to work (1.2 percent for every additional boy ages 10 to 11, 1.9 percent for every additional boy ages 12 to 13 and 3.1 for every additional boy ages 14-15) in households with more boys relative to girls. This result is significant at any conventional level. As mentioned previously, it is sensible to believe that some of this difference could be attributed to higher labor force participation rates among boys and young teens than among girls. *Model 2* controls for this by adding to the regression the proportion of children who are at work. If a higher proportion of boys in the labor force drive the gender bias result at work, we should see the coefficients reduced significantly. This is not the case. In fact, the coefficient on the difference between boys and girls increases in magnitude for all three of the older age groups.[28] In other words it does not seem that mothers were reacting to higher earnings from boys by reducing their labor supply. On the contrary, mothers seem to be employed to a greater extent when boys are working in the household than when girls are. In *Model 3* we include the proportion of children doing housework. Here we see the exact opposite of what we say in *Model 2*. The gender difference coefficients are significantly smaller when we control for the proportion of children helping in the household. Although we still see gender differences, they are now one half to two thirds of what they are in *Model 1*. In *Model 4* we control for both employment types, and we see that the difference in gender coefficients due to the inclusion of the housework variable, conditional on proportion of children at work, is similar to the

[28] We also use the proportion who work at least 10 hours a week. Results were similar.

unconditional difference. This is what we would expect since the housework and market labor variables are mutually exclusive in the PNAD labor questionnaire.

Lastly, a few comments on the proportion variables are in order. The estimates of the proportion variables are somewhat puzzling, and difficult to interpret due to their endogeneity. The more children at work, the more the mother herself will work. If child labor were endogenous, one would expect an income effect leading to a reduction of mother's labor supply. However, households with more children working may in fact be responding to a negative income shock, and thus may be adjusting by increasing labor supply of both children and mothers.[29] Child labor may also be correlated with other unobservables which also increase mother's labor supply. For the proportion of children reporting housework during the reference week, we expect that the more children help out in the household and in child-rearing—the more mothers will be able to work outside the home. This seems to the case, although the magnitude is certainly small: a fifty percentage point increase in the percentage of children doing housework is associated with a 2.6 percentage point increase in a mother's work probability.

Tables 1.8 and 1.9 present the difference-in-differences results: *Model 5* and *Model 6*. For the high versus low relative wage experiment we find little evidence that differences in child wages—and thus no evidence that child employment among girls and boys—is driving mother's labor supply results. For the first differences we see the same results as in Table 1.7: large and statistically significant gender bias differentials. When we look at the difference-in-differences estimator, these differences vanish. Mother's in high relative wage areas do not exhibit greater labor supply biases.

In Table 1.9 the results are markedly different. Here we expected to see lower bias results for mothers in households who have older female relatives who can engage in child-rearing and other household activities. This would be consistent with the hypothesis that older girls substitute for the mother in these activities and therefore allow the mother engage in market employment. The results support this hypothesis. We see that when no childcare is available the employment bias is substantially greater for both the 12-13 age groups and the 14-15 age group. Furthermore, the results are significant at the five and one percent levels for each group, respectively.

The graphical results closely parallel the tabular results presented above. Figure 1.1 presents predicted mother's employment for several child age categories, and for boys and girls. To illustrate, the value for boys age one and under is given by

$$(20)\quad \begin{aligned} &E\Phi\left(\alpha+\sum_{j=1}^{15}\beta_j boy_j+\sum_{j=1}^{15}\delta_j girl_j+\theta child+Z\gamma \mid boy_1=1\right) \\ &=\Phi\left(\alpha+\beta_1-\sum_{j=2}^{15}\beta_j\overline{boy}_j+\sum_{j=1}^{15}\delta_j\overline{girl}_j+\theta\overline{child}+\overline{Z}\gamma\right) \end{aligned}.$$

We see immediately that there is a marked employment bias for the older ages, beginning in most cases with age 11. In Figures 1.2 and 1.3 results are presented for high and low relative wage areas (Figure 1.2) and for households with and without childcare availability (Figure 1.3). In addition, the bottom panels present the 95 percent confidence interval for the difference in the predicted boy-girl employment for each child age group. As with *Model 6*, we see that the gender bias is strong in households without the availability of childcare, and not significant in households with available childcare.

---

[29] Lam, Duryea and Levison (2000) find evidence of an "added worker effect" for both children and mothers.

The c
anyth
consi
whic
signi
of in

pape
the n
diffe
more
num
boys
supp
seco
diff
chil
som
stro
cons
to a
than

The results of this paper are lukewarm in one respect but remarkable in another. The consumption results were not significant or strong enough for us to be able to say anything about the nature of the gender bias evidence. No significant results were consistently found, and in particular we found no gender bias results in adult goods, which suggests the absence of differences in preferences for boys or girls. Since no significant bias results were found in overall consumption, we likewise find no evidence of income effects associated with boys versus girls.

The strong evidence of gender bias in labor supply is the striking result of this paper. We see that the more boys relative to girls are present in the household, the lower the mother's probability of being in the labor force. These results are robust to several different specifications. The results hold for both the lesser-educated households and the more educated households, although they are stronger for the former. We hypothesized a number of explanations for this effect, including that mothers in households with more boys were reacting to greater present and future household earnings by reducing labor supply. Both the empirical specifications designed to test this hypothesis rejected it. A second explanation espoused is that girls enter into the household production function differently than do boys. In particular, they are more close substitutes for mothers in child-rearing and other household activities. Although our levels specification did find some support for this hypothesis (*Model 3*), the difference-in-differences estimator results strongly support it. One must keep in mind, however, that there are other hypotheses consistent with this explanation. In particular, it is possible that this gender effect is due to a belief that girls can "take care" of themselves better and thus require less supervision than boys. This would certainly be in line with most of the results found. The common thread throughout this analysis is that women seem to be responding less to income and substitution effects generated by child gender, but primarily by different roles that boys and girls have in the household.

V. Conclusion

Brazil is a country that lacks social institutions that heavily favor the existence of higher net market returns for male workers. Furthermore, the wage gap between men and women is not as large as in other countries where gender bias has been found, and has reduced considerably over the last decade. Differences in educational outcomes between men and women in Brazil have also narrowed drastically over the past decades. And in recent cohorts, women's educational attainment has surpassed that of men. There are certainly gender differences in labor market outcomes, but all things considered Brazil is an unlikely setting for one to find gender bias. The evidence presented here finds little support for the type of income and substitution driven gender effects economists have found elsewhere. Gender effects on consumption are very subtle or altogether absent. Nevertheless this same evidence suggests extremely large differences in mother's labor supply in response to the number of boys (relative to girls) in the household. This is important in two ways. First, it is the first gender bias result outside of Asia (indeed, outside of the Indian subcontinent). But more importantly, it extends gender bias studies to encompass labor supply effects. And in doing so this study finds evidence that labor supply bias effects, in Brazil, are due largely to differences in household production and not to conventional income and substitution effects associated with boys and girls.

Beck
Beck
Beck
Beh
Bro
Bro
Bo
B
C
C
I
I
I

REFERENCES

Becker, Gary S. *A treatise on the Family*, 1981, Harvard University Press.

Becker Gary S. "A Theory of Marriage: Part I." *Journal of Political Economy*, 1973, 81, pp. 813-846.

Becker Gary S. "A Theory of Marriage: Part II." *Journal of Political Economy*, 1974, 81, pp. S11-S26.

Behrman, Jere R. Robert Pollak and Paul Taubman. "Parental Preferences and Provision for Progeny." *Journal of Political Economy*, 1982, 90(1), pp. 52-73.

Browning, Martin. "Children and Household Economic Behavior." *Journal of Economic Literature,* September 1992, 30(3), pp. 1434-75.

Browning, Martin and Subramaniam, Ramesh. "Gender Bias in India: Parental Preferences, Differential Returns or Marriage Costs?" Mimeo, Yale University and McMaster University, October 1995.

Bourguignon, F., Browning, M., Chiappori P.-A., and Lechene, V. "Income and Outcomes: A Structural Model of Intrahousehold Allocation." *Journal of Political Economy*, 1994, 102, pp. 1067-1096.

Bourguignon, F., Browning, M., Chiappori P.-A., and Lechene, V. "Intrahousehold Allocation of Consumption: Some Evidence from French Data." *Annales d'Économie e de Statistique*, 1994, 29, pp. 137-56.

Chiappori, Pierre-André. "Collective Labor Supply and Welfare." *Journal of Political Economy*, 1992, 100, pp. 437-467.

Coale, Ansley J. "Excess Female Mortality and the Balance of the Sexes in the Population: An Estimate of the Number of "Missing Females"." *Population and Development Review*, September 1991, 17(3), pp. 517-523.

Deolalikar, Anil and Rose, Elaina. "Gender, Savings and Production in Rural India." *Journal of Population Economics*, 1998, 11(4), pp.453-470.

Deaton, Angus. "Looking for Boy-Girl Discrimination in Household Expenditure Data." *World Bank Economic Review,* 1989, 3(1), pp. 1-15.

Deaton, Angus, Javier Ruiz-Castillo, and Duncan Thomas. "The Influence of Household Composition on Household Expenditure Patterns: Theory and Spanish Evidence." *Journal of Political Economy,* 1989, 97(1), pp.179-200.

Deato

Drèze

Harri

Lam,

Lund

Man

McE

McE

Rose

Sen,

Tho

Tho

Tho

Tho

Deaton, Angus. *The Analysis of Household Surveys: a Microeconometric Approach to Development Policy.* World Bank, July 1997.

Drèze, Jean and Amartya Sen. *Hunger an Public Action*, 1989, Oxford.

Harris, Barbara. "The Intrafamily Distribution of Hunger in South Asia." In Jean Drèze and Amartya Sen, eds. *The Political Economy of Hunger*, 1990, Oxford.

Lam, David et al. "Effects of Economic Shocks on Children's Employment in Brazil", mimeo, May 2000.

Lundberg, Shelly and Robert Pollak. "Separate Spheres Bargaining and the Marriage Market." *Journal of Political Economy*, 1993, 101(6), pp. 988-1010.

Manser, Marilyn and Murray Brown. "Marriage and Household Decision-making: A Bargaining Analysis". *International Economic Review*, February 1980, 21(1), pp. 31-44.

McElroy, Marjorie B. and Mary Jean Horney. "Nash-bargained Household Decisions: Toward a Generalization of the Theory of Demand." *International Economic Review*, June 1981, 22(2), pp. 333-349.

McElroy, Marjorie B. "The Empirical content of Nash-bargained Household Behavior." *Journal of Human Resources*, 1990, 25, 559-83.

Rosenzweig, Mark and Schultz, T. Paul. "Market Opportunities, Genetic Endowments, and Intrafamily Resource Distribution: Child Survival in Rural India." *American Economic Review*, September 1992, 72(4), pp. 803-815.

Sen, Amartya. "Women Survival as a Development Problem." *Bulletin of the American Academy of Arts and Sciences*, 1989, 42(2).

Thomas, Duncan. "The Distribution of Income and Expenditure within the Household." *Annales d'Économie e de Statistique*, 1993, 0(29), pp.109-35.

Thomas, Duncan and Chien-Liang Chen. "Income Shares and Shares of Income: Empirical Tests of Models of Household Resource Allocations." Rand Working Paper Series 94-08, April 1994.

Thomas, Duncan. "Like Father like Son, like Mother like Daughter; Parental Resources and Child Height." *Journal of Human Resources*, 1994, 29, pp. 950-88.

Thomas, Duncan, Schoeni, Robert F., and Strauss, John. "Parental Investments in Schooling: Gender and Household Resource Allocation in Urban Brazil." Mimeo, Department of Economics, Michigan State University, January 1997.

Thomas, Duncan. "Intra-Household Resource Allocation: An Inferential Approach." *Journal of Human Resources*, 1990, 25(4), pp.635-664.

Udry, Christopher. "Gender, Agricultural Production, and the Theory of the Household." *Journal of Political Economy*, 1996, 104(5), pp.1010-1046.

Chap

I.

ubiq

of th

We

(yea

othe

sect

of a

ben

legi

larg

ear

19

An

par

the

par

per

[1] Fo

Chapter 2: Trends in female employment in Brazil

I. Introduction

One of the most striking changes in recent labor supply trends has been the ubiquitous increased labor participation and employment of women during the latter part of the twentieth century. In this paper we explain this increased employment in Brazil. We focus on the role of life-cycle (age), cohort and—to the extent possible—period (year) effects. We also focus on the extent to which these effects are in fact proxies for other covariates, such as education and child status. We will then discuss changes in the sectoral distribution of female employment over time, and explain these changes in terms of age and cohort effects. Lastly, we use a 1988 constitutional change in mandated benefits for working women to determine the effect of more generous women's labor legislation on both female employment and sectoral choice.

In the United States the literature on changes in female labor participation is quite large. This literature shows that female labor force participation started to increase as early as 1930.[1] In 1930 the participation rate for American women was below 0.25; by 1980 over half of all women ages sixteen and above were in the labor force. In Latin America there is a smaller but growing literature dealing with female labor force participation. This literature points to significant heterogeneity in women's entrance into the labor market. Caribbean countries such as Jamaica, Barbados and Haiti had female participation rates over 0.5 as early as the 1950. On the other hand, during this same period countries such as Brazil, Ecuador and Honduras had female participation rates in

[1] For a review see Goldin (1990).

the te

1970s

perce

empl

men

nine

in fe

in th

marr

seve

parti

poin

mar

the

poir

mar

pres

the

---

[2] Sea

[3] We

[4] For

refer

PNA

daug

do no

the teens and low twenties.[2] They only experienced growth in participation rates in the 1970s, 80s, and 90s.

In Brazil female labor force participation increased from 17 percent in 1950 to 41 percent in 1999. Figure 2.1 details the more recent change in male and female employment rates throughout the 1980s and 1990s.[3] We see a considerable decrease in men's employment, from around 0.80 in the mid eighties to under 0.75 in the late nineties. For women the change is even more pronounced: a ten percentage point increase in female employment over the two decades. However, not unlike the American pattern in the sixties and seventies, Table 2.1 shows that this increase has been mostly from married women. In the table we see three years of means—1981, 1990, 1999—for several variables of interest for married and single women, ages 15 to 70.[4] The change in participation rates for single women was of 5.7 percentage points and 2.7 percentage points for the 1980s and 1990s decade, respectively. Compare that to the change for married women: a full ten percentage points for each of the two periods. The figures for the employment for those two decades tell a similar story: a six and three percentage point gain for the single women group and a ten and six percentage point increase for the married group. Closely related to Table 2.1 are the patterns we see in Figure 2.2, which presents the age-employment profile for selected years, from 1981-1999. Here we see the classic hump shaped life-cycle employment profile, with employment rates increasing

2 See Psacharopoulos and Tzannatis (1993) for a compendium on female employment in Latin America.

3 We define the employment rate as the proportion employed, as opposed to proportion in the labor force.

4 For the purposes of this study, "married" is defined as household heads with spouses and their spouses. It refers to currently married and not to ever-married. There is no information on past marriages in the PNAD, so ever-married determination is not possible. In a small but significant number of households daughters and sons of the household head may also be married. These are treated as "single", since the data do not allow for their identification.

rapid

and f

large

sma

seco

mu

eff

The

inc

(ag

be

w

pr

ar

m

re

I'

e

se

rapidly as women reach their late teens and early twenties, leveling off though the thirties and forties, and then dropping off in the late forties and early fifties. We also see that the largest increases in employment rates occur in the first decade (the eighties), and to a smaller extent in the first half of the second decade (early nineties). The latter half of the second decade shows very little increase in employment rates.

Our approach is to model changes in female employment with both probit and multinomial logit models. The probit model is more appropriate for estimating the effects of education, child status, and household composition on female employment. The multinomial logit model allows us to determine which sector cohorts chose as they increasingly entered the labor market. We also attempt to identify all three time effects (age, period, cohort), although there are several difficulties with this endeavor, described below.

Beyond explaining the extent of female participation, we also would like to know what sectors and occupations these women entered. The multinomial logit model provides a crude measure of sector choice; however, we also present a mainly descriptive analysis of changes in occupational structure over the nineteen year period, by using measures of gender concentration.

The remainder of the paper is structured as follows. Section II provides a brief review of the literature. Section III describes the problem at hand and the data. In section IV the models are described and results are discussed. In this section, we also discuss the effect of changes in labor legislation on women's employment and sector choice. In section V we describe occupational changes in women and men's employment. We also discuss

paper.

II.

quite l

descrip

force p

have re

the inc

factors

force

incre

wag

im

h

hy

hyp

discuss changes in the degree of gender segregation. In section VII we conclude the paper.

II. Review of literature

The literature dealing with female employment and wages in the United States is quite large and varied. Goldin (1990) provides the most comprehensive historical description of American women's work. She points out that almost all the gains in labor force participation have come from married women. Participation rates for single women have remained relatively unchanged since the early nineteen hundreds. Explanations for the increased participation of women can generally be cast into either supply or demand factors. Smith and Ward (1984) favor demand-side explanations for the increased labor force participation. They use Mincer's (1962) technique and find that 60% of the increase in female labor force participation from 1900 to 1980 can be explained by rising wages. Some of this increase is through a direct wage effect, some of it though wage's impact on fertility. However, the majority of research on changes in female labor supply has emphasized supply-type factors. These have more often than not used variants of methods pioneered by Mincer (1962) and Ben-Porath (1973) to estimate labor supply functions. They emphasize the role of education and other forms of human capital in explaining labor supply.

In newly industrialized (or non-industrialized) settings the so-called "U"-hypothesis is commonly invoked to explain changes in participation rates.[5] The "U"-hypothesis maintains that during the early phases of industrialization female participation falls.

in pa

jobs

avai

ferti

Psac

amo

size

dyna

de L

mod

asso

She

Thi

inve

Pau

rece

low

Are

fron

[5] Nor
exper

falls. This fall occurs, in part, as home production gives way to industrialized goods, and in part as agricultural employment declines. Women are unable to find compensating jobs in industry, these jobs going mainly to men. Participation then increases with the availability of more administrative and secretarial jobs, and as women reduce their fertility and increase their schooling. This hypothesis has found support in Psacharopoulos and Taznnatos (1989), Goldin (1995) and Mammen and Paxson (2000), among others.

The literature concerning female participation and employment in Brazil is sizeable, but consists mostly of cross-sectional analysis, with little or no mention of a dynamic component. Tiefenthaler (1992) uses the 1989 *Pesquisa Nacional por Amostra de Domícilios* (PNAD), a Brazilian household survey, to estimate a multinomial logit model of female employment. She finds that higher levels of education are strongly associated with formal employment (versus informal, self-employed or no employment). She also found that small children had the largest negative effect on formal employment. This large disemployment effect is also found in Soares (2000). Humphrey (1996) investigates the cyclical nature of children and female employment in the state of São Paulo. He found that the quality of jobs held by women actually increased during recessions, a trend that did not carry over to men. Men's jobs became more informal and lower paying during recessions. This is evidence against the "added worker effect". Arends-Keunning (1997) looks at changes in female labor force participation in Brazil, from 1976-1990. Arends-Keunning, using the PNADs, applies the methodology

---

[5] Note that the U-hypothesis is also applied in industrialized countries to explain their female participation experience. See Goldin (1995).

dev

fem

eff

dec

the

per

ad

or

er

pa

P

a

w

t

I

developed by Heckman and McCurdy (1980) to evaluate the effect of income changes on female labor force participation. She creates cohort x education cells and uses fixed cell-effects to distinguish between temporary and permanent income changes. She finds that decreases in family incomes during the early eighties account for almost fifty percent of the observed increase in female labor participation among single women and twenty percent among married women. She argues that this is strong evidence in favor of the added-worker effect.

Sedlacek and Santos (1991) estimate a labor supply function for women, focusing on the role of husband's earnings. They find that a husband's earnings affect the wife's employment in a U-shaped manner. Among the very poor families the wife's participation is higher, as it is among wealthy families. Wajnman and Rios-Neto use the PNADs to make projections about female labor force participation. Based on a stationary assumption regarding future cohort participation rates, they find that participation rates will increase simply due to the age composition of present and future cohorts. Although the focus of their paper is to outline secular changes in the wage gap in Brazil, Barros, Ramos and Santos (1993) describe changes in female labor participation during the 1980s. They find that between 1981 and 1989 the proportion of workers who were women increased from 29.8 to 35 percent. They find that women are not over-represented in low-paying occupations or sectors. Rather, they are over-represented in higher-paying occupations. Also, in their analysis they find that in higher-paying occupations women have higher levels of schooling than men.

Lam and Duryea find a strong link between increased school attainment and decreases in fertility in Brazil, using data from the 1984 PNAD. However, they find highly significant, if weak (in magnitude) link between fertility and labor supply. They discard the explanation that a decrease in fertility led to an increase in women's labor force participation. Instead, Lam and Duryea favor the hypothesis, based on work by Willis (1973) and Becker and Lewis (1973), which maintains that higher levels of schooling raised the shadow price of children and thus led to a substitution into quality versus quantity. Lastly, Soares and Izaki (2001) do a multivariate analysis of female labor force participation in Brazil using five PNAD cross-sections. They find that, as in the United States, increases in female labor force participation in Brazil came mainly from married women. In their model education explains 50 percent of the increase in labor force participation. Fewer children in the household were also associated with an increase in participation, but these and all other variables accounted for only three percent of the increase over the 19-year period spanned by their cross sections. Neither cohort nor year effects were explicitly modeled in their multivariate analysis.

III. Changes in cohort employment and the data

A number of factors likely contribute to the changes in female labor supply. Marriage, the decision to have children—and how many children to have—and retirement are among the major factors that vary over a woman's lifecycle. These covariates will lead to different employment rates for women at different ages, creating an age profile for female labor supply. In addition, the labor supply, income and wages of other household members will also affect female labor supply through income and substitution effects. If we look at leisure time as a normal good, if other members of the household face a higher wage, women would be less likely to participate in the labor market, *ceteris paribus*.

Secular changes may also be associated with changes in labor supply. In particular, cyclical factors, such as recessions, will affect the demand and supply of labor. During recessions the demand for labor will contract, and employment and wages will fall. Also, during recessions women may delay their entry into the labor market because of lower wages or high unemployment, deciding instead to invest in more schooling or training. Women may also change their fertility decisions based on cyclical factors. In recessionary periods women may delay entry into the labor market (or exit labor markets) in favor of starting a family. Secular changes also reflect both more progressive attitudes toward women's work as well as changes in the degree and nature of discrimination in labor markets, although these effects may arguably be driven by inter-cohort changes rather than by inter-year changes.

Changes between cohorts play a key role in determining labor supply. Schooling levels, and the quality of schooling, will likely vary between birth-year cohorts, with latter-year cohorts being more schooled and having access to better quality schooling. Beyond schooling, younger cohorts will also have different attitudes toward work: current generations may favor careers for their daughters to a greater extent than their parents did. Closely related to this, social norms may change from cohort to cohort, increasingly favoring the entry of women, and married women in particular, into the labor force. Lastly, to the extent that younger cohorts have fewer children (and perhaps higher quality children), and thus have fewer household commitments, they may be more likely to enter the labor force, and less likely to exit the labor force for child-bearing.

cov
cha
of
the
wo
tha
Fo
ex
hi
In
Th
H
co
w
c
T
a
a
c
l
r
r

In describing age, period and cohort effects we have described all the major covariates commonly associated with labor supply: education, child status, marriage, changes in social perceptions regarding women in the workplace, and wages and income of other family members (husbands in particular). If we knew for certain that these were the only factors affecting labor supply—and if we could measure them perfectly—we would not need age, period or cohort effects. However, most of these measures, when they are available, may be measured with error, and others may not be measured at all. For example, although one can obtain a respondent's completed years of schooling, it is exceedingly difficult to measure the quality of schooling. Also, a woman's fertility history, although not difficult to measure, is not routinely asked in household surveys. Instead what is usually measured is the age of children currently living in the household. This may be a poor proxy for children ever born, especially for older households. Husband's wage, although easy to obtain, may be endogenous and actually reflect other covariates. In particular, successful men may marry highly schooled and career-oriented women, and vice versa. Finally, changes in social attitudes and perceptions are hard to conceptualize, let alone measure—and any measurement would be ambiguous at best. Therefore, the use of age, period and cohort effects is justified as a way to take into account omitted or poorly measured/modeled variables. In this paper we primarily answer two questions related to time effects and female labor supply (i) what role do changes in education, child status and other household variables play in the changing labor supply of Brazilian women and (ii) when these effects are taken into account, how much do age and cohort (and year) variables shape labor supply. For identification reasons detailed below we focus on age and cohort effects.

ef

co

ex

or

co

lo

an

lo

dy

lif

Br

en

m

co

ob

is

co

no

wi

col

[6] Th

As the discussion above makes clear one can expect, age, period and cohort effects to be large. However, there are a number of difficulties in using age, period and cohort effects. First, these correlations are devoid of any causal interpretation, to the extent that said effects largely measure the degree to which we are unable to account for or quantify economic and social variables that determine labor supply.[6] Second, true cohorts require that one follow the same households from survey to survey. Due to the logistical problems of following people over time there are few surveys with this design; and these suffer from serious attrition problems, as people move, die, or simply are no longer willing to participate in the survey.[7]

That much said, there are several advantages to using these effects. First, a dynamic labor supply picture is necessary to eliminate bias associated with confounding life-cycle effects with cohort effects. In a growing economy (such as is the case in Brazil), the life-cycle profile taken from a single cross-section will overstate the employment at younger ages, and understate the employment at older ages, as older members of the cross-section will also be those belonging to older cohorts and older cohorts will tend to have lower labor employment rates. In effect, the life-cycle effects obtained from a single cross-section will be picking up cohort effects as well. Second, it is reasonable to proxy omitted and poorly measured covariates with age, period and cohort effects. In a sense these estimates will tell us the magnitude of what we either do not know or cannot measure. Also, to the extent that these latent variables are correlated with the variables we <u>can</u> measure, such as schooling or child status, omitting age and cohort effects may engender serious omitted variable bias. Including the age and cohort

[6] This point is elegantly developed in Heckman and Robb (1985).

effe

deta

with

esti

onc

yea

sup

(1)

whe

cha

dur

prol

c ye

effe

om

also

and

(2)

[7] In a
[8] For

effects can mitigate this bias or, in some instances eliminate it altogether.[8] Furthermore, detailing the life cycle and inter-cohort patterns in employment is a laudable goal in itself with its own policy-making implications.

The task then becomes one of modeling age, period and cohort effects and estimating them empirically. The main problem with modeling all three effects is that, once you know an individual's birthyear and his age in the survey, you will know the year of the survey. This presents a serious identification problem. Consider the labor supply function:

(1) $$y_i = \alpha + X_i'\beta + f(A,C,T) + \varepsilon_i,$$

where y is a measure of labor supply, X is a vector of household and individual characteristics, A are age dummies, C are cohort dummies and T are period (year) dummies. The problem then becomes how to specify and estimate $f(A,C,T)$. This is problematic due to the fact that a person observed at year t, and born on year c, must be t-c years old. In other words, $a=t-c$. This makes identification of year, cohort and age effects simultaneously impossible without *ad hoc* restrictions. On the other hand, omitting either of these effects will likely bias the other two effects, and possibly bias β also (to the extent that individual and household covariates are correlated with age, period and cohort effects). Consider estimation of a simple linear-effects model:

(2) $$y_i = \alpha + X_i'\beta + \eta a + \kappa t + \lambda c + \varepsilon_i,$$

[7] In addition, for these surveys to be useful they must follow families for numerous years.
[8] For example, as Deaton (1997) points out cohorts may be used to create fixed-effects estimators.

wh

esti

sub

(3)

We

tha

coh

ma

cas

(4)

wh

ca

ide

wh

du

lin

ad

be

do

ag

where *a, t* and *c* are age, period and cohort trends. We know that this model is not estimable, due to the *a=t-c* identity. If we omit one of the effects, say time, we obtain by substituting *a*+c for *t*

$$\begin{aligned} &y_i = \alpha + X_i'\beta + \eta a + \kappa(a+c) + \lambda c + \varepsilon_i \\ &y_i = \alpha + X_i'\beta + (\eta+\kappa)a + (\lambda+\kappa)c + \varepsilon_i \end{aligned} \tag{3}$$

We can see that the coefficient on both age and cohort will be biased; and to the extent that household and individual covariates increase or decrease with time, conditional on cohort and age, the coefficient on β may be biased. The same type of argument can be made if effects are modeled flexibly, with age, cohort and year dummies.

There are a few solutions proposed for this identification problem. Consider the case were dummies are used. Then (2) becomes

$$y_i = \alpha + X_i'\beta + A\mu + T\varsigma + C\psi + \varepsilon_i, \tag{4}$$

where *A,T,* and *C* are matrices of rowsize equal to the number of age, period, and cohort categories and μ, ς, ψ are vectors of *a ,t* and *c* dummies. In this case we have the identity $As_a = Cs_c + Ts_t$, where s_a, s_c, s_t are trend vectors. For example, $s_a = \{1, 2, 3, ..., a\}$, where *a* denotes the number of age categories. Again, this implies that even after one dummy is omitted in each of the effects matrices, we will still have vectors which are linear combinations of each other. The only way to estimate the model is to impose additional restrictions on the parameter values. The most parsimonious restriction would be to drop an additional dummy category from one of the effects matrices. If this were done for the first age group we would have the restriction $\mu_1 = \alpha$, in effect restricting the age effect to be the same for the first two age categories. One could also impose stronger fu

im

a t

or

to

si

In

ve

ex

m

be

ye

re

ar

m

fu

fl

su

fe

[9] D
an
tre
du
an
ca
lab
du

functional form restrictions. If one believes that year effects increase linearly, one could impose the restriction that $\varsigma_2 - \varsigma_1 = \varsigma_3 - \varsigma_2 = ... = \varsigma_t - \varsigma_{t-1}$. This is equivalent to estimating a time trend instead of the more flexible dummy configuration.[9] Lastly, one could simply omit one of the effects, most notably the period effects. If period effects are orthogonal to the β parameters, the only induced bias would be on the age and cohort effects—and since these coefficients have no economic interpretation, arguably this bias is innocuous. In all of the identification strategies described, interpretation of the effects parameters is very tenuous. One cannot know to what extent one is picking up true effects and to what extent one is picking up an inappropriate identification strategy. For this reason, for the most part we omit period (year) effects, and are mindful that cohort and age effects will be biased. We also attempt to identify all three effects based on restrictions placed on the year effects. In particular, we use both (a) differences in functional form and (b) restrictions on dummy variables to identify effects. The differences in functional form are between the treatment of age and the treatment of period and cohort effects. Age is modeled with a quartic polynomial, while period and cohort effects are treated with a fully-flexible dummy variable specification. Although the quartic specification is not as flexible as a semi-parametric dummy variable specification, a preponderance of labor supply evidence indicate that age-employment profiles are 'hump-shaped'; therefore we feel that the specification is flexible enough. Second, we also restrict level year effects to

[9] Deaton (1997) adds a time trend to the year effects matrix, and subtracts the same time trend from the age and cohort effect matrices. He then imposes the restriction that the year effects be orthogonal to the time trend. In our specification we rely both on functional form differences and restrictions on year and cohort dummies to identify all three effects. In particular, age is modeled with a quartic polynomial, while cohort and year effects are modeled with dummies. The reason age is modeled with a quartic is that, *a priori*, we can be confident that age-employment profiles will follow the hump-shaped pattern found in virtually all labor supply studies, and therefore their functional form does not require the semi-parametric flexibility of dummy variables.

be the same for a range of years. In both Figure 2.1 (which presents a time trend progression of female employment) and Figure 2.7 (which presents a time-series picture of informal employment), we see that there is very little difference in the employment pictures between 1996 and 1999. Therefore, in our identification strategy we restrict the year effects to be the same for the last four survey years. We do this by modeling year effects with dummies and omitting the last four dummies in the estimation. We also present results for alternative sets of restrictions, in order to better understand how sensitive age, period and cohort identification is to alternative specifications.

Before proceeding further it is necessary to detail the methodology used to construct cohorts. The data used, which will be described in detail below, consist of several cross-sections. That is, we do not have a true panel, since the same households are not re-interviewed every year. Rather, a new sample is drawn in each successive year. With this design one cannot follow cohorts, since the households will only be in the survey on one particular year. Synthetic birth cohorts (which we will refer to simply as cohorts) are then constructed by matching individuals of age a in year t with individuals of age $a+1$ in year $t+1$, individuals of age $a+2$ in year $t+2$, and so on. It is important to note that this method actually has a number of advantages (and disadvantages) over true cohorts. The main advantage is that survey attrition is not a problem, since we are not attempting to match the same individuals from one survey to the next. Second—and related to the first point—cohorts can be constructed for unlimited number of cross-sections. In panel data, the longer the panel length, the less representative successive rotations will be of the original sample. On the other hand, cohorts can tell us nothing about the dynamics within a cohort, a key advantage of panel data. Also, cohorts implicitly assume that the sample of individuals (and households) of age $a+n$ in sample year $t+n$ is drawn from the same sample as individuals of age a in sample year t, n years latter. This may not hold for older cohorts, since many of them will have died, or if there is migration from populations not sampled in year t (e.g. other countries, areas not sampled, etc.). A more serious problem can occur if household heads are replaced by their progeny. In this scenario, the household heads in year t may be reclassified as father of household head in year $t+n$.[10]

The question we would like to answer is exactly what accounts for the increase in employment rates among women, and among married women in particular. Among the factors that customarily enter into a labor supply function, education and child status (fertility) are prominent. More educated women are more likely to work outside the home than less educated women, an empirical regularity observed in countless settings. Child status is also important, as couples will, to varying degrees, determine family size and labor supply jointly. Regional differences are important in Brazil, since there is substantially observed (and unobserved) heterogeneity between different regions, particularly between the industrialized Southeast and the economically depressed Northeast. A likely additional source of heterogeneity arises from deep cultural differences (which may partly determine attitudes toward female work), which although we cannot quantify, may vary to a large extent from region to region. Beyond regional differences, the dichotomy between urban and rural setting is also important—indeed, one can argue that the urban centers of the more developed Southeast have more in

[10] See Deaton (1997), Deaton and Paxson (1994) and McKenzie (2001) for a detailed discussion of the potential problems involved in drawing inferences from synthetic-cohorts.

common with the urban centers in the Northeast than they do with the sparsely populated rural settings of the Southeast.

Even when accounting for the covariates above, both observed and unobserved factors common to birth cohorts are likely explanations for increased employment rates. Although we may not observe these factors, we can certainly observe their effect, by analyzing the employment rate of each cohort separately. Figure 2.3 presents a visual age profile of the employment rate of eight cohorts, for both married and single women. First notice that although the employment rates for married women have increased, they are still significantly below that of single women. Single women's employment rate peaks at 0.75, while the employment rate of the highest employed married cohort barely reaches 0.5 at around 35-40 years of age. Second, one immediately notices that the cohort lines for single women are close to each other. In fact, there is no discernable trend. This is characteristic of a stationary state cohort, as the age-employment profile of each successive cohort basically mirrors the previous cohort's. The same cannot be said of the married group. Here we see that each successive cohort of married women has an employment rate markedly above its predecessor's. We also see the largest cohort gains between the 1937-41 and the 1942-46 birth-year cohort, and between the 1942-46 and the 1947-51 birth-year cohort. The gains taper off for younger cohorts, and are modest (albeit very discernable) between the youngest two cohorts.

Returning to Table 2.1, we would like to know to what extent these cohort changes are driven by changes in variables of the women's labor supply function, such as school attainment, child status, and urbanization. Although the table is ordered by year and not by cohort, one can still see that all the covariate changes commonly associated with higher employment rates are present in both the top (married) and bottom (single) panels. The proportion of women in rural areas declined in both periods (1981 to 1990 and 1990 to 1999), consistent with the observed migration to urban centers throughout the two decades. In fact the increase in urbanization accelerated during the second period. In both panels we also see a marked increase in educational attainment, mostly driven by increases at lower levels of education. It bears noting that the increase is not symmetric with respect to gender. During the 19-year period covered by our analysis girl's schooling equaled and eventually surpassed schooling levels for boys. One should also note that education levels, as measured by completed years of schooling, are biased downward, to the extent that many of those still in school are also in the sample. The last set of changes is those with respect to child status. They show that the number of boys and girls declined dramatically in both periods. This is in line with the observed concurrent decline in total fertility rate by 50 percent in the first decade.

The data used in Table 2.1 and throughout the analysis are from the 1981-1999 PNADs.[11] The PNAD is an annual household survey conducted by the Brazilian Statistics Ministry (*Instituto Brasileiro de Geografia e Estatística*, IBGE). It is a representative sample from the entire country, except for a few rural regions in the north, which are not surveyed. The PNADs include detailed information on labor supply asked of all household members at or above the age of ten. This information includes type and sector of employment, formal employment (employees who work under the auspices of the government), informal employment (employees who work outside the legislated employment framework), hours worked, earnings, tenure (later years only), etc. Every

[11] Since the survey was not conducted in 1991 and 1994, data from those years are not available.

household member is surveyed, and information is collected on schooling, age, race (later years only). All relationships are measured *vis à vis* the household head. Therefore it is possible to identify the household head's spouse and children, but not stepchildren. This poses a problem since our variable for sons and daughters will include all the household head's sons and daughters living in the household (from all marriages), but will exclude both sons and daughters of the spouse from another union, or children no longer in the household. Since household heads are predominantly male (80 percent male), this will potentially affect son and daughter information for 80 percent of the wives in the sample.

We use two samples in our analysis, the married sample and the single sample. Both samples are limited to women ages 15-70. Girls ages 10-14 were omitted since the labor supply of children is likely very different than that of teens and women. On the other hand, since many adolescents enter the labor market after 8 years of schooling (primary school) and after 11 years of schooling (secondary school), 15 years old would encompass most of the secondary-schooled and all of the primary-schooled. For the married women we dropped all observations for which any of the covariates was missing for either the husband or the wife (43,828 observations), leaving a sample size of 963,102 observations. For the single women sample the total number of observations was 684,514. The husband's wage was deflated to constant 1999 Reais.[12] Note that for all models data is pooled across years.

The multinomial logit model will try to account for the different employment choices made by women. The choices were defined as follows: (0) not at work, (1)

[12] The exchange rate in October of 1999, the month of the survey, was 1.7 Reais to the U.S. Dollar. Since all PNADs deflate monetary values to October of the survey year, the October price index is used to deflate values across years.

formal worker ,(2) informal and (3) self-employed.[13] Informal wage workers tend to be in very low-paying, low quality, jobs, characterized by high turnover. Self-employed workers will be small business owners, usually with no employees, as well as self-employed professionals. We did not stratify the self-employed by number of employees, although the data permit us to do so. They were treated as one category, irrespective of establishment size.

As mentioned previously, defined benefits introduced in 1989 considerably increased the firing costs of all formal workers and increased the labor costs of women in child-bearing age, to the extent that these women now had increased maternity leave and were protected from being fired while pregnant, or immediately following a birth. In particular, the second legislative innovation allowed women to take unpaid maternity leave during their pregnancy followed by up to four months of paid maternity leave after birth. To test the hypothesis that these legislative innovations caused employers to increasingly hire "high-risk" women informally, or perhaps not hire them at all—we will compare the employment of women in child-bearing ages (20-35) to that of older women (40-55). The sample of older women will in effect serve as a control group, to the extent that their employment is not affected by the legislative change to the same extent as the employment of younger women is affected. We will also use a second control group, single women 20-35. We compare this group to the group of married women 20-35, who are arguably at the highest risk for child bearing. This reflects the fact that the majority of

[13] Formal workers carry a "worker's notebook", which the employer must sign. This entitles the worker to paid vacation, as well as numerous other worker benefits. It also means both the worker and the firm must pay income tax on any earnings. Informal workers do not have their notebook signed, and thus the government does not know of their employment. Although this is technically illegal, in practice there are few consequences for either the firm or the worker.

births are in wedlock, and employers can expect single women to have lower fertility rates than married women.[14]

IV. Main model, results and analysis

Four models are estimated. The first is a probit model for employment. This is the focus of the analysis. The second is a multinomial logit model. We use the multinomial logit to allow coefficients to differ across employment choices. The third and fourth models are probit and multinomial logit models for informal employment and employment type, to determine the effect of the changes in labor legislation mentioned above. In all of our model estimates it is important to keep in mind that education, child status, and other covariates are not exogenously changed from individual to individual. These are all variables that are likely correlated with a number of other variables, most of them unobserved. Education is certainly correlated with ability, and with other forms of human capital. Child status is likely to be endogenous, as couples by-and-large decide when to have kids, how many, and the spacing between them. These decisions are jointly made with labor supply decisions. The husband's wage may also be endogenous—at least to a certain extent—due to the fact that career-driven and high 'ability' husbands may chose wives with certain characteristics, such as high schooling levels. However, since we have information on husband's schooling and age, two covariates which wives may also select on, the endogeneity of husband wages will not be as pernicious as the endogeneity of child status. There are certainly methods to try to deal with labor supply

[14] This control group strategy also rests on the assumption that employers can observe marital status at time of employment.

endogeneity, however, our attempt to deal with the problem here will be limited to controlling for most observed sources of possible heterogeneity.[15]

For the married women four probit models are estimated, to allow for the incremental inclusion of covariates. They are as follows:

(M1) $p(emp) = \Phi\left(\sum_i \alpha_i R_i + \sum_l \varsigma_l A_l + \sum_j \beta_j C_j\right)$

(M2) $p(emp) = \Phi\left(\sum_i \alpha_i R_i + \sum_l \varsigma_l A_l + \sum_k \lambda_k S_k + \sum_j \beta_j C_j\right)$

(M3) $p(emp) = \Phi\left(\sum_i \alpha_i R_i + \sum_l \varsigma_l A_l + \sum_k \lambda_k S_k + \sum_m \eta_m H_m + \sum_j \beta_j C_j\right)$

(M4) $p(emp) = \Phi\left(\sum_i \alpha_i R_i + \sum_l \varsigma_l A_l + \sum_k \lambda_k S_k + \sum_m \eta_m H_m + \sum_n \mu_n F_n + \sum_j \beta_j C_j\right)$

R_i are six regional dummies and one macroeconomic variable, GDP growth. A_l are four age polynomials (age, age squared, age cubic, age quartic). S_k are seventeen education dummies. H_m are four husband variables (husband's predicted wage, wage squared, years of schooling and age). F_n are two sets of six child status dummies and household size.[16] C_j are fifty cohort dummies. Note that the husband's education enters into the husband's wage and into the wife's employment equations. The motivation behind this is that apart from increasing wages, husband's schooling may also affect the degree of assortive mating, child quality, and general attitudes toward work. To the extent that more educated husbands may prefer women who work (do not work), and to the extent that professional women may prefer well educated (less educated) husbands, it is appropriate

[15] There is a vast literature on instrumental variable techniques, fixed effect models, as well as studies of twins and siblings, among others. Other approaches try to find exogenous sources of variation in covariates, or "natural experiments". If there are no good instruments, or no exogenous sources of variation, one can try to control for all known correlations, in an attempt to capture most of the omitted variable effects and thus mitigate the bias involved.

to control for husband's schooling beyond its effect on wages. Furthermore, since schooling and time at home enter into the child quality production function, conditional on husband's wage, women with more educated husbands may want to spend less (more) time with their children. The exact direction of the effect cannot be determined *a priori* by economic theory and will depend on how the husband's education enters into the production function.

Table 2.2 presents the probit results for the employment of married women. Marginal effects are reported. Since there are 53 birth year cohorts (1921-73), only every fifth cohort is reported, starting with the 1925 cohort. All cohort results are relative to the omitted category, the 1921 cohort. Controlling for age and region only, we see that the cohort coefficients are very large, the youngest cohort being more than 30 percent more likely to work than the oldest. However, the marginal change in employment probabilities is not as strong in the older cohorts as Figure 2.3 suggested. In fact, the changes between cohorts are small until the 1930 cohort and then distributed evenly. We also see that even the age quartic is significant at conventional levels. This result motivated us to keep the four degree polynomial format for age.

Seventeen education dummies are added for model M2. As expected, higher educational attainment is associated with higher employment probabilities. We also see that these "returns" are highly non-linear. The returns to the first ten years of schooling amount to about eleven percent. The return to an addition year adds a full ten percentage points to the employment probability. This is not unexpected since the 11th grade is the

---

[16] One daughter 0-4 years old, two daughters 0-4 years old, three or more daughters 0-4 years old, one son 0-4 years old, two sons 0-4 years old, three or more sons 0-4 years old. Same for daughters and sons 5-9 years old.

last grade of the Brazilian high school. If labor markets are susceptible to returns to credentialism, then the high school diploma should affect wages (and thus employment) beyond the marginal effects on human capital of an additional year of schooling. Furthermore, to the extent that students who repeat years are more likely to drop out of school, holding a high school diploma will also be a signal of those characteristics associated with non-repetition.

Adding the education dummies (M2) reduced the cohort coefficients by about one third. This indicates that a good portion of the increase in married women's employment rates across cohorts was due to more and more educated cohorts. In M3 husband's covariates are added in, using a two stage approach. Wages are predicted in a first stage, based on seventeen husband education dummies, age and age squared, as well as region dummies. These predicted values are for working husbands only; results are reported in the Appendix Table 2.13.[17] The predicted values are then used in the wife's employment probit in the second stage, and standard errors are adjusted. We see that even after accounting for age and education in the husband's wage and wage squared, schooling and age are significant, more schooled husbands decreasing a wife's employment probability. Both polynomials of the husband's wage are significant, with the linear term positive and the squared term negative.[18] There is no obvious interpretation for this result. One would expect substitution and income effects to move in the opposite direction, to the

---

[17] Note that husband's schooling and age are entered into the woman's employment probit also. Non-linearities in returns to one's own education in Brazil have been widely documented, motivating the flexible functional form in the first stage estimation. However, we have no reason *a priori* to expect that husband's education will affect wife's supply in a largely non-linear way, conditional on husband's wage. We maintain that the effect of education and age on men's wages is going to be significantly more non-linear than the effect of these covariates in the wife's employment probit (second stage).

extent that leisure is a normal good. It is possible that wages are still capturing omitted variables associated with assortive mating. In any case, this is contrary to the results found by Sedlacek and Santos. Also note the inclusion of husband's covariates had very little impact on the cohort effects.

The last probit model (M4) introduces child status and household composition variables. Household size and six child status variables for daughters and six for boys are included, although the 5-9 year old dummies are not reported in the table. We see that additional young children are associated with large decreases in employment probabilities, and these decreases are relatively linear in the number of children. We also see that larger households will more likely have the wife employed. The inclusion of the child status and household controls again reduced the magnitude on the cohort coefficients. Although the reduction was not as large as the one associated with schooling, it is in the direction expected, and is around two to three percentage points for all but the oldest cohorts.

So far we have estimated a labor supply function, allowing different level effects for each cohort. However, it is possible that the underlying supply function may have changed from one cohort to another. One way to determine this would be to estimate different supply functions for each cohort. Figure 2.4 presents predicted employment rates for different levels of schooling, based on cohort-specific supply functions. Predictions are made at sample means, except for age, which is held constant at 35 years of age. The model used to make predictions was the one based on the M4 probit model,

---

[18] Higher degree polynomials were also modeled for wages, as was a partially linear model. In both cases the result was similar to the level and quadratic formulation described above. Furthermore, when model M4 was estimated for each cohort separately, the wage functions were similar to the pooled model.

only with cohort effects omitted, since equations were estimated by cohort. We hold age constant to net out any age effect associated with cohorts having different average ages in different cohorts. Although there is some evidence of a more concave education effect for the older cohorts, this effect is not very pronounced. In general the education effect seems to be similar across cohorts, with older cohorts having slightly lower returns.

In Figure 2.5 results are presented for the same set of cohort-specific equations with respect to the child status variables. The estimates are for married women. Here it bears repeating that fertility and labor supply, at least to some extent, are jointly determined. Therefore, caution is warranted when interpreting the child status coefficients as causal. The panel presents predicted employment rates for households with one, two, and three or more boys ages 0-4. The estimates for ages 5-9, as well as the estimates for girls are not graphed.[19] We clearly see that the effect is not the same across cohorts. In particular, the effect tends to be greater for younger cohorts than for older cohorts. For the 1962-66 cohort the effect of one boy 0-4 is associated with a full ten percentage point reduction in employment rates of the wife. The disemployment effects are also nonlinear, with the first child being larger than the second, and the marginal effect of the second larger than the third. The inter-cohort differences are also evident, with older cohorts seeing smaller effects than younger cohorts. This may be due to the fact that older cohorts with young children in the household may be unrepresentative of a

[19] Results for the 5-9 year old group are significantly smaller, and are smaller for older cohorts than for younger cohorts.

typical working family. The same may also be true of the youngest cohort, but to a lesser extent.[20]

Although the focus of this paper is on married women, as this is the group that has experienced the most dramatic increases in employment and participation rates, labor supply functions are also estimated for single women. In their case, child status variables are not reported, as they are of small magnitudes and often not significant at conventional levels. Table 2.3 reports the probit results for single women. As expected, there is little inter-cohort change in employment rates. This is the same result found in Figure 2.3, and confirms the absence of cohort effects for single women. It is also interesting to note that for single women, the association between education and employment is not monotonic. In fact, in the first ten years of schooling the parameter estimates vary almost randomly from one year to the next. The strong positive effect for higher education is still present, and of similar magnitudes.[21]

Identifying age, period and cohort effects

In this section we present the results from out attempt to identify all three temporal effects: age, period and cohort. We have already discussed that any attempt at identifying all three effects necessarily hinges on some identification restriction, and that the results are typically quite sensitive to that restriction. Our analysis is no exception. The identification restriction imposed was that the period effects for the 1996-1999 years

[20] Results for girls, although not reported, were similar to those for boys. The only difference is that for the older group more boys are associated with larger disemployment effects. This is similar to the result found by Soares (2000).

be the same. Imposing this restriction was trivial: we simply omitted the last four year dummies from our labor supply equation. We also should point out that since age effects are modeled with polynomials and not with age dummies, part of the identification rests on the disparity in functional form, although it is not clear exactly how this would impact the period and cohort dummy coefficients (if at all). \

Table 2.4 shows age, cohort and period coefficients for model M4, when year effects are added in. As mentioned before, only dummy variables for 1981-1995 are included. The last four period dummies being excluded in order to identify the effects separately. First we see, as expected that the period and cohort coefficients are imprecisely estimated due to the inherent colinearity between the two effects. In particular, most of the period effects are not significant when jointly estimated with cohort effects. The latter cohort effects are still significant at conventional levels. We also see that when estimated with cohort effects, the year effects bear little resemblance to their values in the model without cohort effects. This is not true for the cohort effects, which are now 20-30 percent attenuated *vis-à-vis* the model without year effects. This provides some evidence that cohort effects are more robust to inclusion of period effects than vice-versa—at least with the identification restrictions that we adopted.

In order to determine how sensitive our period and cohort effects are to identification restrictions, we estimated the model in Table 2.2 with three other identification strategies. Therefore, the first identification strategy (corresponding to the last column of Table 2.2, only with year effects added in) was to restrict the last four

[21] Although not reported in the table, female employment is procyclical, even if the small magnitude on the GDP growth variable indicates small effects. For the married women a ten percent increase in GDP is associated with a 0.8 percentage point increase in employment probability. For single women the effect is about twice as big. In both cases the effect is robust to model specification.

year effects were restricted to be equal (i.e. dummies for the last four years were dropped). In the second and third specifications the year effects for 1997, 1998 and 1999 (for the second specification) and 1998 and 1999 (for the third specification) were omitted. The fourth specification omitted only the 1999 dummy, relying purely on functional form restrictions to identify all three effects. The inclusion of four specifications helps to demonstrate the sensitivity of identification to different identification restrictions.

Figure 2.6 presents predicted age, period and cohort effects based on the model M4, only with the GDP growth variable omitted and with the year dummies corresponding to four different specifications added in. The life-cycle profile is as we would expect, with women entering the labor force in their late teens and early twenties and exiting in their fifties and sixties. We also see—and this is true for all effects estimated—that the first and second identification schemes behaved very similarly, while the last two schemes behaved rather erratically. Looking at the cohort effects, we see that the specification relying purely on functional form results in nonsensical values: an employment profile increasing with birthyear. However, for the three other specifications, *ceteris paribus*, changes in cohort birthyear account for a seventeen percentage point (first and second specifications) and a thirty percentage point increase (third specification) in employment from the oldest to the youngest cohort. These changes are also concentrated in the older cohorts. For the year effects we see a similar story: for the first two specifications the change in employment is rather modest, at about five percentage point, concentrated in the first decade. For the third specification there is no increase and even a modest decrease in over the years. For the last specification we see a dramatic secular increase, from 24 percent to 37 percent of married women.

At first glance a simpleminded explanation would be that cohorts account for much more variation in employment than do year effects. One must, however, keep in mind that in our sample we have over fifty elapsed cohorts and only nineteen elapsed years. Indeed, the average change per cohort using either of the last two specifications is 0.3 percentage points. For the same specifications, the average year effects are 0.23 percentage points, all of which occur during the first decade. Indeed, although the average marginal cohort effects are still larger, the difference is not as large as Figure 2.6 suggests.

Counterfactual simulations

In order to obtain an estimate of exactly what percentage of the secular change in employment is explained by the covariates, what is explained by the cohorts and what is unexplained, we have simulated employment rates from model M4, using the 1981, 1990 and 1999 cross sections. In effect, we have estimated the employment rate at the sample mean for each of these years, using the coefficient estimates obtained from M4. We can obtain different employment estimates for different covariate subgroups by successively using the mean of each subgroup from period $t+1$, while holding the other subgroup means constant at period t. In other words, consider four subgroups, R, S, H and C. R consists of age, region and urbanization. S consists of the education dummies. H consists of the husband variables, as well as child status variables, and C represents the fifty-two cohort dummies. The predicted change in women's employment from 1981 to 1990 is given by

(5) $$d9081 = \Phi(R_{90}'\alpha + S_{90}'\lambda + H_{90}'\eta + C_{90}'\beta) - \Phi(R_{81}'\alpha + S_{81}'\lambda + H_{81}'\eta + C_{81}'\beta).$$

If we want to know how much of this change is due to a change in region, we would compute

(6) $$\Phi(R_{90}'\alpha + S_{81}'\lambda + H_{81}'\eta + C_{81}'\beta) - \Phi(R_{81}'\alpha + S_{81}'\lambda + H_{81}'\eta + C_{81}'\beta).$$

We can then add in the schooling means from 1990 to obtain the change due to school attainment improvements, and so forth.[22] Table 2.5 summarizes the results from this simulation. First, we see that our model is actually not very good at predicting changes in employment over years. The model predicted a 6.1 percentage point increase in the first decade, while the change was actually 8.2 percent. In the second decade the model performed somewhat better. The model predicted a 5.9 percentage point change, when the actual change was 6.6 points. In other words, in the first decade 74 percent of the change was explained and in the second decade a full 90 percent of the variation was explained. The breakdown by covariate group confirms what we saw in the probit tables. Schooling explains 28 and 48 percent of the change between 90-81 and 99-90, respectively. Our 52 cohort dummies explained most of the change—61 percent for the first decade and 55 percent for the second. Also, child status, household size, and husband's variables explain a large portion of the change in the first decade (11 percent), but actually contribute to lower the magnitude of the change in the second decade. This

[22] Note that since the normal cumulative distribution is not linear, the employment of a mean is not equal to the mean of individual employments. This means that changes do not necessarily add up to the total change. However, a similar exercise was done with the linear probability model and results were comparable.

is in line with the dramatic decrease in fertility, whose effects were mostly in the 1980s and early nineties.

In the bottom panel of Table 2.5 we see the counterfactual simulation applied to changes over cohort. Here we are trying to explain the change in average employment for married women between the 1942 and 1952 cohorts and between the 1952 and the 1962 cohorts. We restricted the cohort range to be twenty years—so that we can compare the magnitudes with those of the counterfactual period simulations. Unfortunately the results are not very illuminating, mainly due to a small change in average year over the period, i.e. the median year of the 1942 cohort was 1987 and 1990 for the 1962 cohort—only three years of difference. In contrast, the change in median age is of a full 18 years. This is, of course, because most members of early cohorts are old and most members of late cohorts are young. If we had more period variation in the sample, we would have a more complete age-employment profile for early (and late) cohorts.

The counterfactual cohort simulations behave completely disparately for the two periods. In the early cohorts we have increase of employment of 5.8 percentage points. In the later cohorts we have a decrease of 3.4 percentage points. For purposes of the table, in the first period a negative value on the 'percentage explained' signifies a percentage <u>decrease</u>, in the second period it signifies a percentage <u>increase</u>. For the early cohorts (1952-1942), we see that age accounts for a large share of the increase in employment (51 percent), while changes in child status do not contribute to an increase in employment—to the contrary—they actually act to decrease employment. This is likely due to the fact that fertility declines only started for latter cohorts. The only effect that is similar to that of the 1990-1981 period change in the top panel is education. Here we see that schooling did act to increase employment of married women from the 1942 to the 1952 cohort. However, this effect mitigated for the later cohorts' change. For the 1962-1952 birthyear cohort change, education plays a smaller role in the employment change: it acts to increase employment, its magnitude being 17 percent of the total change. In fact, due to the larger variation in median age, age is the greatest predictor of the decrease in employment across cohorts. This is due to the fact that in later cohorts most members are still on the upward sloping portion of the age-employment profile.

Sectoral choice results

The last set of results to be discussed is the sectoral choice logit results. However, before detailing them it is useful to graphically motivate what changes took place over the period in the distribution of employment between the formal, informal wage, and informal self-employed sectors. Figure 2.7 presents a time series of the share of employment in the two sectors dubbed "informal". The first characteristic worthy of notice is a dramatic increase in female participation in the informal wage sector starting in 1990. One possible explanation for this is that the PNAD questionnaire changed in 1992. However, *a priori* there is no reason to believe that the change would have significantly affected the quality of the response to the employment sector variables.[23] Another explanation may be related to the constitutional change in 1988 (which went into effect in 1989). The new constitution increased firm's firing costs of formal workers, and therefore perhaps increased the attractiveness of informal employment. Theoretically, this change would impact both men and women, and indeed we do see a decrease in formal employment for both groups during this period. However, because of the additional benefits accruing to women, it is possible that employers would hire more women informally, or perhaps be less likely to hire them at all. This hypothesis is tested and will be discussed in the next section.

Whereas the secular increase in female employment of married women can largely be attributed to changes between cohorts, the same cannot be said about sector choice. Figure 2.8 details the proportion of employed women working informally, both for self-employment and for informal wage work, by cohort. Unlike the clear cohort pattern visible in the earlier cohort graph, the cohort pattern here appears to be staggered five years apart. For our graphs the cohorts themselves are defined by a span of five years. Therefore we would expect period effects to be seen every five years of age. This is exactly what we see. In other words, the age profile for the 1957-1961 birth year cohort at 30 years of age is similar to the age profile of the 1962-1966 birth year cohort at age 25. This is so because at those loci the two cohorts are actually in the same period (year). This pattern is most clear for the informal wage workers, but is also present in the

[23] Prior to 1992 military and so-called "statutory" government workers were classified as workers without a "worker's notebook", since working for the government one does not technically hold a notebook, as government workers have their own social security system. Therefore, these workers were identified based on industry and occupation codes and reclassified as workers WITH a notebook for the purposes of this study. Starting with the 1992 PNAD, a separate question was added to explicitly identify government workers, who were then also reclassified as WITH a notebook. It should be noted that prior to 1992 the occupation and industry codes were sufficient to identify an estimated 98 percent of "statutory" government workers. Even if some discrepancies arose from the questionnaire change, it is unclear why this would affect women and not men.

self-employed panel. In the sector choice model below we attempt to identify both period and cohort effects.

The multinomial logit model is important because the three types of employment are very distinct, and are characterized by different types of workers. Self-employed workers tend to be poorly educated and mostly urban, and tend to be older.[24] Informal workers tend to be urban and rural, also poorly educated, but tend to be younger than self-employed. Formal workers, on the other hand, would be the group most similar to labor markets in the United States. They are more educated and span every industry and most occupations. Informal and self-employed workers tend to be clustered in particular industries and occupations. With these differences in mind it is important to know which of these groups the married women's cohorts were more likely to enter. Did their employment patterns across the groups reinforce the patterns existing before the increase in participation and employment or were they different? The multinomial logit will help us answer these questions. The variables in the multinomial logit are the same as those in the probit model M4. In Table 2.6 the effect of a one unit change in the explanatory variable is given by the relative risk ratio. Coefficients smaller than one indicate a decrease in participation probability vis-à-vis the omitted category (not at work); greater than one indicate an increase. The first interesting result is that low levels of schooling tend to decrease a married woman's probability of being an informal wage worker, relative to not being at work. Only at higher schooling levels does this trend reverse itself. A similar trend is seen for the self-employed, although not as dramatic; a middle-schooled woman being 40 percent more likely to be self-employed relative to not at work than a woman with no schooling. For formal workers the education coefficients are as expected. A high-schooled woman is 13 times more likely to work as a formal worker relative to not being at work than is a woman with no schooling. For higher schooling levels the increases are even more prominent: for college-educated women relative to high school-educated the odds ratio is 11.

All the child status results in the same direction, as one would expect. However, the magnitudes vary from category to category. Young boys and girls have the largest disemployment effect for formal workers, followed closely by informal workers, and lastly by self-employed. This may be due to a greater flexibility of work schedule for informal workers and even greater for self-employed individuals. The effects also tend to be slightly non-linear in the number of children, as they were in the probit model.

The cohort coefficients tell an interesting story. For both formal and self-employed employees the results are similar, if somewhat larger for the latter group. The effects are larger for the younger cohorts. The results for informal wage workers are also positive, but are significantly larger than for the other two categories; and again the results are larger for younger cohorts. This indicates that (i) as with the probit results, younger cohorts are more likely to work, and (ii) this work is more likely to be in the informal wage sector than in either self-employment of in formal employment. However, Figure 2.8 suggests that we may also be picking up period effects, and not just cohort effects. In order to address this we also estimate the multinomial probit model with both year and cohort effects, omitting the last four year effects, in the same manner that was

[24] Self-employed also includes a small percentage of well-educated professionals such as doctors, dentists and lawyers. This group is very different from the typical Brazilian self-employed.

done with the employment models. Again, we should restate that identification of all three effects is problematic and, as we have seen, interpretation of the effects is quite sensitive to identification restrictions. Figure 2.9 presents the relative risk coefficients for the year and cohort effects—age effects were essentially the same as in the model with only cohort effects, and are therefore omitted. For period effects, relative risk ratios are relative to the four omitted years, 1996-1999. For cohort effects, the ratios are relative to the omitted 1921 birthyear cohort. The results are also presented in tabular form in Table 2.7. Figure 2.9 clearly shows that while there seem to be some period effects over the two decades, their magnitude is smaller than the cohort effects. Rather, we see that both informal wage employment and self-employment increased monotonically with cohort, relative to no employment. The cohort effects on informal wage employment are particularly large. There are only modest cohort effects on formal employment. We also see that most of the period effects for informal and self-employment are concentrated between 1990 and 1993—even if the effects are in opposite directions. For formal employment, the period effects are present mostly in the first decade, with an increase of roughly 30 percentage points in the probability of formal employment versus no employment. Therefore, although Figure 2.8 suggests that a large portion of the increase in informal employment is due to period effects, we find qualified evidence (qualified because of the tentative identification strategy) that cohorts may in fact be more important that period effects.

1988 policy change

The last two results discussed in this section pertain to the effect of the 1988 constitutional change on employment. Since formally employed women had access to one month's maternity leave, and had assured job security after the maternity leave, one could expect shifting from formal to informal employment following the change. This is so because the law is only binding on formally employed workers. The informal probit models the probability of informal employment, conditional on employment. The multinomial logit model extends this analysis to include shifting into other forms of employment or into no employment. We chose this structure to empirically test the hypothesis that more generous labor legislation may have led to the increase in the informal employment of women. To test this we define two control groups. The first is the group of women no longer likely to bear children, those ages 40-55. The second is the group of single women ages 20-35. In the first specification the treatment group would be women ages 20-35 and in the second it would be married women ages 20-35. The model is then specified as follows:

(7) $$p(\inf) = \alpha + X'\beta + \gamma control + \eta year9099 + \lambda control \times year9099 + \varepsilon$$

where "control" is the dummy for the control group, "year9099" is the dummy for post-policy years. The X vector is a vector of labor supply controls, equivalent to those in S3, but without age or cohort effects. The results of this model are presented in Table 2.8. The two questions of interest are (i) how did women's informal employment change after 1989? and (ii) how did this change differ for the control and for the treatment groups. Formally, we want to know

(i) $$E(y \mid year9099 = 1) - E(y \mid year9099 = 0) = \mu + \lambda \times control$$

(ii) $$\begin{aligned}&E(y \mid year9099=1, control=0)-E(y \mid year9099=0, control=0)-\\&\left(E(y \mid year9099=1, control=1)-E(y \mid year9099=0, control=1)\right)=-\eta\end{aligned}$$

Item (ii) is the difference-in-differences estimator. The parameter on the 1990-99 dummy, multiplied by the control dummy, evaluated at the sample mean (i.e. the proportion of control observations in the sample, 0.31 for the first control group and 0.60 for the second), answers the first question; and the <u>negative</u> of the parameter on the interaction of the year dummy and control dummy answers the second.[25] In both cases we see that there is a substantial post-1989 increase in informal employment of five (0.164 x 0.31) and fifteen (0.2585 x 0.60) percentage points for the first and second control group, respectively. We also see that the difference-in-differences estimator suggests a 5.4 percent treatment effect in the first specification and a 4.2 percent treatment effect in the second. This is comforting, for it suggest that both of our control groups behaved similarly. It also suggests substantial shifting to informal employment following the policy change.

The probit, however, does not model shifting away from employment. It only models shifting to informal versus formal employment. It is possible that women at high risk for child bearing may have shifted from formal employment into no employment, or even into self-employment. In order to capture all four employment alternatives—no work, informal wage work, self employed and formal wage work—we also test the constitutional change with a multinomial logit model. Here the omitted category is <u>formal employment</u>, since most interesting comparisons will be with respect to this group. The dependent variables included are the same as in the probit in (7). The results

[25] This is because we defined the control group as equal to one, the treatment group equal to zero.

are presented in Table 2.9. First note that the "Control x 1990-1999" parameter estimate measures the relative risk ratio for the control group, not the treatment group. To obtain the relative risk ratio for the treatment group (i.e. the treatment effect) one must take the reciprocal of the control group parameter estimate. Therefore, as reported, relative risk ratios greater than one actually correspond to a decrease in the category's probability vis-à-vis the omitted formal employment category. If we look at the treatment effect in the two identification strategies we see that the probability of informal employment increased relative to formal employment for both control groups (by 39 and 25 percent respectively). However, we also see that the probability of self-employment decreased (by 18 and 12 percent). If women in the treatment group indeed find it difficult to find formal employment, it is not clear why they would shift away from self-employment. If any effect were to be found vis-à-vis formal employment, it would be a shifting to self-employment.[26] In closing, both the probit and the multinomial logit results are consistent with a shifting away from formal employment for women at high risk for child bearing following the 1988 constitutional change. The only result at odds with this interpretation is the observed shifting out of self-employment for the treatment group.

[26] One should also point out a possible problem with the married women control group. The assumption necessary for this group to act as a true control group is that its employment patterns resemble those of single women throughout the period, absent the policy change, and that the policy changes not affect the group's employment patterns. The 19-year period in question witnessed dramatic changes in married women's employment patterns, so using this group as a control for employment patterns of single women

V. Occupational changes

The PNADs provide information on three-digit occupation and three-digit industry. In this paper we focus on the two-digit classification.[27] The main reason for this is that occupational and industry classifications are fraught with reporting error. Many respondents will not know exactly what occupation they belong to, and many others will not belong to any of the available occupations, despite the large number of occupations available. Also, at the three-digit level it becomes difficult to separate occupation from industry. In addition, respondent's self-reported occupational and industry classification is likely to change over time and even between older and younger generations. These difficulties are mitigated at the two-digit level. In the case of reporting error, the respondent may likely mistake an "office receptionist" for an "office assistant"—two different but neighboring three-digit occupations—but will probably not mistake "administrative assistants" for "managers"—two neighboring two-digit occupations. Nevertheless, since most of the occupational change studies done use three-digit occupations, for comparability purposes, we compute statistics for both two and three digit occupations.

One question we wish to answer in this section is where did women go as they entered the labor force? Did they reinforce previous patterns of feminization or did they enter occupations traditionally held by men? One way to address this is to look at some measure of the distance between the proportion of men in an occupation and the proportion of women. One such measure is the index of dissimilarity, which is one half

---

may be inappropriate. Nevertheless, we are comforted by the fact that both control groups yielded similar results.

[27] After 1992 the occupational and, to some extent, the industry classifications changed. All years were translated back to pre-1992 definitions, but in any case, the changes were only at the three-digit level.

of the sum of the absolute difference between men's share in each occupation (share of total male employment in said occupation) and women's share in the corresponding occupation:

(8) $$D = \frac{1}{2}\sum_i \left|\frac{F_i}{F} - \frac{M_i}{M}\right| = \frac{1}{2}\sum_i s_i \left|\frac{f_i}{f} - \frac{m_i}{m}\right| ,$$

where $f_i(m_i)$ is the proportion of women (men) employed in group i, $f(m)$ is women's share (men's share) of total employment, and s_i is occupation i's share of total employment (occupation size). Higher index numbers correspond to a greater degree of segregation. The ½ is necessary to bind the index between zero and one.[28] This index has the advantage of being neutral to scalar transformations in employment rates.[29] It is also simple to interpret. The index measures the percentage of men and women who need to be reshuffled into different occupations in order to achieve perfect occupational parity. The overall change in the index can be written as:

(9) $$\frac{dD}{dt} = \frac{1}{2}\sum_i s_i^t \left|\frac{f_i^t}{f^t} - \frac{m_i^t}{m^t}\right| - \frac{1}{2}\sum_i s_i^{t-1} \left|\frac{f_i^{t-1}}{f^{t-1}} - \frac{m_i^{t-1}}{m^{t-1}}\right|$$

We can further decompose this change into two parts, dc and ds:

(10) $$dc = \frac{1}{2}\sum_i s_i^{t-1} \left|\frac{f_i^t}{f^t} - \frac{m_i^t}{m^t}\right| - \frac{1}{2}\sum_i s_i^{t-1} \left|\frac{f_i^{t-1}}{f^{t-1}} - \frac{m_i^{t-1}}{m^{t-1}}\right| \quad ds = \frac{dD}{dt} - dc$$

dc is the change due to occupational cell sizes and ds is the change due to the change in the gender makeup of each occupation.

[28] Suppose there are two groups, and in one group women hold all jobs and in the second men hold all jobs, then the index will be ½ x (|1-0|+|0-1|) = 1. If on the other hand, employment is perfectly equitable, we have ½ x (|½-½|+|½-½|) = 0.

[29] Watts (1998) argues that the dissimilarity index is not invariant to changes in the relative size of occupation cells, and disfavors it for this reason. He also reviews a number of other indexes.

Table 2.10 presents the time-series evolution of the dissimilarity index. It also breaks down the index into changes due to changes in the relative size of the occupation cells (dc) and changes due to the gender composition of members of each occupation (ds). A cursory look at the table reveals that for three-digit occupations the index varies from the low 60s to the mid 60s and for two-digit it varies from the mid 50s to the low 60s. In both cases we see a clear decrease of five points in the index over the twenty year period. We also note that most of the change (80-90 percent) is from changes in the composition of cells, not from changes in cell sizes. Although the author knows of no other time-series of occupational (dis)similarity in Latin America, there are a number of studies in the United States. Most of the United States evidence (also using three-digit occupations) points to a stagnant index for the first half of the century (at around 65), followed by a decline of 5-7 points, with most of the decline concentrated since 1970. This is actually quite similar to what we see in Table 2.10.

Table 2.11 we seek to answer a different question. Here we have a multi-layered picture of the distribution of women in two-digit occupations. We attempt to match women and men's qualifications—schooling and age (which is a proxy for training) to deciles of occupations, where the deciles are over the degree of feminization in the base year of 1981. In other words, the first decile contains the most feminized occupations and the last decile contains the least feminized. We would like to know how much human capital women bring into occupations, and if they need to bring more human capital to occupations that are most segregated. First note that women tend to have higher levels of schooling and tend to be younger than men in the workforce. This is a regularity of Brazilian labor markets pointed out by a number of labor and development economists. The comparisons between deciles reveal what we suspected. In the least feminized occupations, women tend to have the largest educational advantage, although no discernable pattern emerges with respect to age. Furthermore, we see that this pattern is reinforced as women enter the most male dominated occupations (i.e. as the occupations become more feminized) and, surprisingly, the same is true as men enter occupations dominated by women. Indeed, the two most feminized occupation deciles are the only ones in 1999 in which men's educational attainment surpasses that of women. And the top three are the ones in which the educational gap between men and women is the largest. This pattern suggests that women need a qualification premium in order to enter male-dominated occupations (as do men entering female dominated occupations). This premium may be due to discrimination on the part of employers, but more likely reflects either coworker discrimination, possible lower productivity of women in a male dominated environment, or even latent costs associated with hiring women in a male dominated environment. The same argument can be made with respect to men in a female dominated environment.

The last result we consider is in Table 2.12. In Table 2.12 we would like to know to what extent occupational growth affects the degree of concentration of men and women in occupations. In the discussion above we saw that the degree of occupational concentration did decrease over the last two decades. We also saw that women (men) with relatively higher levels of human capital were the ones who entered male-dominated (female-dominated) occupations. Anecdotal evidence would suggest that these entrants would be young workers joining the labor force rather than older workers changing occupations. Indeed, one would expect initial entrants, from younger cohorts, to be the driving force behind the reduction in occupational concentration. In order to test this hypothesis, we define an occupation in a particular year as being *highly segregated* if at least 95 percent of its members are men or 95 percent of its members are women. We then regress this measure of concentration on the change in an occupation's share of total employment (a measure of occupational growth) and on the change in the proportion of young in an occupation, a measure of an occupation's ability to draw new entrants. Since our interest is on the effect of growth in occupational employment on occupation concentration, we also control for the occupation's initial status as *highly segregated* in 1981. We estimate this model by using occupational growth and lagged values of occupation growth, lagged up to four years, to the extent that women and men may switch occupations in favor of a more unfavorably concentrated occupation, only when there is some evidence of growth. Included in the regression as a dependent variable is women's share of overall employment in each year. Since most of the segregated occupations are male-dominated, increasing women's employment will reduce male-dominated segregation to a greater extent than it will increase female-dominated segregation, therefore controlling for women's share of employment across all occupations is appropriate. We also estimate a model with both current-year growth and all lagged growth variables. Note that the results in Table 2.12 are weighted by the size of each occupation, so larger occupations have a larger impact on the results. There were 77 two-digit occupations and 17 years in the analysis, with 1981 being the first year. The model is, thus

$$(11) \qquad p(highseg) = \Phi\left(\textstyle\sum_i \beta_i \, dshare_{t-i} + \sum_i \delta_i \, dyoung_{t-i} + \gamma \, highseg81 + \eta \, femshare\right),$$

where *highseg* is the dummy variable for highly segregated occupation; *dshare* is the change in women's share of employment in a particular occupation; *highseg81* is the dummy variable for the initial (1981) values of segregation; *femshare* is women's share of overall employment.

While many results in Table 2.12 are expected, some are not. As expected, fast-growing occupations are also the ones becoming less segregated. Furthermore, occupational concentration is more sensitive to recent growth than it is to growth further back. An increase of 0.01 in an occupation's share of total employment is associated with a 4.2 percentage point decrease in the probability of that occupation being highly segregated. This effect decreased over time: 2.6, 2.3, 2.2 and 2.1 percentage point effects for lagged growth. Unfortunately, growth rates from year-to-year are highly correlated, therefore it is not possible to provide a meaningful analysis by including <u>both</u> current period growth and lagged growth in the same specification. This is attempted and reported under the column "all effects". One can see that the multicolinearity only allows us to identify growth effects lagged four periods. All others are not significant. In Table 2.12 we also see that, as expected, the best predictor of a highly segregated occupation is its segregation status in 1981. The unexpected part of the table is the bottom panel, which presents the results for the growth in the proportion of young (defined as 16-25 years old) in an occupation. Controlling for overall occupational growth, we see that an increase in the proportion of young in an occupation does decrease the incidence of high segregation, but not significantly. Furthermore, inclusion of the change in the occupational size does not appreciably affect the coefficient on the proportion young—and in both cases it is not significant. This provides evidence against the hypothesis that new entrants, from recent cohorts, are the driving force behind the observed decrease in the level of occupational gender concentration. In fact, it appears that occupational growth does a better job at explaining changes in segregation over time. However, further inspection will reveal that this is consistent with older, married women, driving changes in occupational segregation. In other words, the evidence suggests that as married women enter the labor market, they are the ones who enter male-dominated occupations. A final note of caution is in order. We have thus far interpreted occupational growth as exogenous. It is certainly possible that we are picking up other factors (beyond growth in proportion young) that affect occupational growth. It is also possible—albeit improbable—that gender-diverse occupations are actually better able to attract new members, and therefore the causality in Table 2.12 would be backwards. Since in this paper we are not attempting to identify occupational growth (although it may be possible to do so) beyond the inclusion of the change in proportion young, the reader is duly cautioned about this possible problem.

VI. Conclusion.

Like almost all countries, Brazil has seen a dramatic increase in the participation and employment of women during the postwar era. As with the United States and other first-world countries, we have shown that this increase has come, primarily, from married women. We have also shown that the correlations in the Brazilian women's labor supply functions tend to conform to that of other countries. Education is associated with higher employment, more young children associated with lower employment of the mother. We saw that in Brazil's case education has highly non-linear effects on employment, a result in line with cross-sectional studies of Brazilian labor supply, and in line with studies of other labor outcomes in the Brazilian labor market, which seem to be non-linearly correlated with education. However, the focus of this paper has been a cohort analysis of changes in labor supply, with an emphasis on married women. Here we have seen that improvements in education and reductions in fertility rates have explained much of the inter-cohorts differences in employment. Notwithstanding, cohorts still explain a large portion of explained changes in employment. Furthermore, our sector choice analysis showed that most of the changes in married women's participation have been from women entering self-employment and, above all, informal work. We also tested the hypothesis that increases in the cost of formal employment of women in child-bearing age would engender a shifting from formal employment to informal employment. By using two different identification strategies we conclude that there is evidence that this shift did indeed take place, although we can never be sure that we are not picking up other macroeconomic effects that may have affected our treatment and control groups disparately.

Our occupational analysis revealed that Brazilian occupations have become less segregated over the last two decades, a result in line with patterns observed in the United States. We also showed that the degree of segregation is highly correlated with the schooling advantage that women have over men. In male-dominated occupations women's schooling advantage over men is the largest, and in highly feminized occupations men's schooling deficits are the lowest. These results are consistent with a number of discrimination hypotheses, but may also be due to different employment costs or differences in productivity due to gender concentration in the work environment.

REFERENCES

Arends-Kuenning, Mary. "Using repeated cross sectional data to estimate female labor supply in Brazil, 1976 to 1990" University of Michigan Population Studies Working Paper, 1997.

Barros, Ricardo Paes, Lauro Ramos and Eleonora Santos. "Gender differences in Brazilian labor markets," in T. Paul Schultz, ed. *Investment in women's human capital.* University of Chicago Press, 1995.

Becker, Gary and Gregg Lewis. "On the interaction between quantity and quality of children." *Journal of Political Economy* 81(2), 1973: S279-88.

Beller, Andrea H. "Changes in the sex composition of U.S. occupations, 1960-81." *Journal of Human Resources* 20(2), 1985: 236-50.

Ben-Porth, Yoram. "Labor force participation rates and the supply of labor." *Journal of Political Economy* 81(May-June), 1973: 697-704.

Cotter, David A., et al. "Occupational gender desegregation in the 1980s." *Work and Occupations* 22(1), 1995: 3-21.

Deaton, Angus. *The analysis of household surveys: a microeconomic approach to development policy*, World Bank, 1997.

Deaton, Angus and Christina H. Paxson. "Patterns of aging in Thailand and Côte d'Ivoire," in David A. Wise, ed. *Topics in the economics of aging*, Chicago University Press, 1994.

Goldin, Claudia. *Understanding the gender gap: an economic history of American women.* Oxford University Press, 1990.

Goldin, Claudia. "The U-shaped female labor force function in economic development and economic history" in T. Paul Schultz, ed. *Investment in women's human capital and economic development.* University of Chicago Press, 1995.

Heckman, James and T. MaCurdy "A Life Cycle Model of Female Labor Supply," *Review of Economic Studies* 47, 1980: 47-74.

Heckman, James and Richard Robb. "Using longitudinal data to estimate age, period, and cohort effects in earnings equations," in William M. Mason and Stephen E. Fienberg, eds. *Cohort analysis in social research: beyond the identification problem.* Springer-Verlag, 1985.

Hall, Robert E. "The measurement of quality changes from vintage price data," in Zvi Griliches, ed. *Price indexes and quality change*. Harvard University Press, 1971.

Hill, Anne. "Female labor force participation in developing and developed countries—consideration of the informal sector." *Review of Economics and Statistics* 65, 1983: 459-467.

Humphrey, John. "Response to recession and restructuring: employment trends in the São Paulo metropolitan region, 1979-87." *Journal of Development Studies* 33(1), 1996: 40-62.

McKenzie, David J. "Consumption growth in a booming economy: Taiwan 1976-96." Mimeo, 2000.

Lam, David and Suzanne Duryea. "Effects of schooling on fertility, labor supply and investment in children, with evidence from Brazil." *Journal of Human Resources* 34(1), 1999: 160-92.

Mincer, Jacob. "Labor force participation of married women: a study of labor supply," in H. Gregg Lewis, ed. *Aspects of labor economics*. Princeton University Press, 1962.

Polachek, Solomon William. "Occupational segregation and the gender gap." *Population Research and Policy Review* 6, 1987: 47-67.

Psacharopoulos, George and Zafiris Tzannatos, eds. *Case studies on women's employment and pay in Latin America*. World Bank, 1992.

Sedlacek, Guilherme Luis and Eleonora Cruz Santos. "A mulher cônjuge no mercado de trabalho como estragégia de geração de renda familiar." Discussion paper no. 209. IPEA, 1990.

Smith, James P. and Michael P. Ward. "Time-series growth in the female labor force." *Journal of Labor Economics* 3(1, pt.2), 1985:S59-S90.

Soares, Sergei, Rejane Sayuri Izaki. "A participação feminina no mercado de trabalho." mimeo, IPEA, 2001.

Soares, Yuri S.D. "Consumption, savings, labor supply and gender bias in Brazil." mimeo, Michigan State University, 2000.

Stelcner, M. et al. "Labor force behavior and earnings of Brazilian women and men, 1980" in Psacharopoulos, George and Zafiris Tzannatos, eds. *Case studies on women's employment and pay in Latin America*. World Bank, 1992.

Tiefenthaler, Jill. "Female labor force participation and wage determination in Brazil, 1989" in Psacharopoulos, George and Zafiris Tzannatos, eds. *Case studies on women's employment and pay in Latin America*. World Bank, 1992

Wajnman, Simone and Eduardo Rios-Neto. "Quantas serão as mulheres: cenários para a atividade feminina." in Balatar, Maria Isabel, ed. *Trabalho e genero: mudanças permanências e desafios*. Editora34, 2000.

Watts, Martin. "Occupational gender segregation: index measurement and econometric modeling." *Demography* 35(4), 1998: 489-496.

Willis, Robert. "A new approach to the economic theory of fertility behavior." *Journal of Political Economy* 81(2), 1973: S14-64.

Chapter 3: The role of pilot training and experience in aviation safety

The purpose of this study is to determine the effect of pilot characteristics, and pilot training in particular, on aviation safety outcomes, such as pilot error, accident severity, and accident rates. Although there is an enormous literature detailing the returns to training and experience in the workplace, very little exists regarding non-market returns to experience. This is partly due to the fact that training and experience, even when general, is believed to have few benefits outside the labor market. One notable exception to this is the role of experience in the determination of accident outcomes. This paper focuses on the role of experience in aviation safety outcomes.

This work contributes to aviation safety literature in four important ways. First, it models the probability of pilot error in aviation accident using microdata, by taking advantage of the richness of the National Transportation Safety Board's (NTSB) Aviation Accident/Incident Database System (AIDS). Although aviation outcomes have been studied using aggregate airline data, this is one of very few attempts at modeling similar outcomes using pilot-level microdata—and the first attempt for general aviation. Second, it models accident severity, both in terms of injury severity and in terms of extent of damage. These estimates are first calculated conditional on accident occurrence, by using the full sample of accidents. An attempt is then made to calculate severity and extend of damage estimates by using a sub-sample that more closely matches the universe of pilots. Although these estimates cannot be considerd unconditional, *per se*, they may more closely match a random draw from the population of pilots. Thirdly, we attempt to model accident probability by comparing sub-samples that more closely resemble the population of both accident and non-accident pilots to the actual sample of accidents in the AIDS.

The fourth contribution of this paper relates to the scope of analysis. In this paper we model the safety outcomes mentioned above for private flights in addition to commercial flights. The private or "general aviation" group of flights has long been ignored in the safety and aviation literature.

The remainder of the paper is organized as follows. In section I the topic is discussed, and a review of the existing literature is presented. In section II the model is developed. In section III the data requirements are outlined and the dataset is described. Section IV presents the empirical results and section V concluding remarks.

I. Introduction

There is a sizable literature regarding the determinants of aviation accidents—perhaps too large, considering the relatively small number of fatalities associated with flying relative to other forms of transportation.[1] Nevertheless, economists and accident analysts have learned much regarding the factors associated with aviation safety outcomes. This includes analysis of the effect of maintenance, pilot error, equipment failure, weather, and air traffic on accident rates and fatality rates (Oster and Zorn, 1989).[2] Other safety-related factors studied in the literature include the structure of flight, such as stops per trip and length of flight (Rose, 1990),[3] and shifts from one mode

[1] The Department of Transportation reported 693 fatalities associated with all forms of air travel in the US for 1999. In contrast, there were 41,611 fatalities associated with highway travel for the same year.
[2] Oster and Zorn analyze the effects of factors under the airlines' control, as well as those outside the airlines' control, in the context of deregulation.
[3] Rose finds that a longer average flight time is associated with a higher accident rate.

of transportation to another (Bylow and Savage, 1991).[4] Measures of firm financial performance have also been studied in association with airline accident and fatality rates (Globe, 1986; Rose, 1990),[5] as have measures of airline experience (Barnett and Higgins, 1989; Kanafani and Keeler, 1989; Oster and Zorn, 1989; Rose, 1990).[6]

One characteristic is shared by all of the studies above: they relied on aggregate carrier data. That is, in their studies data are at the carrier level, not at the flight or pilot level. Also, none of the studies used explicit measures of pilot experience, such as average pilot experience.[7] Two studies (Phillips and Talley, 1992; Phillips and Talley, 1996) use accident-level micro data to determine the effect of pilot characteristics on accident severity and extent of aircraft damage. In both analyses the authors find that pilot experience has a small but significant impact on severity of accident and extent of plane damage.[8] However, their approach does not model non-linearity in the returns to experience adequately. Also, their analysis is limited to commercial flights, and makes no attempt to account for selection bias in the AIDS.

---

4 Bylow and Savage discuss the decline in auto travel resulting from deregulation.

5 Here the evidence is mixed. Earlier studies, such as Globe's, found no significant correlation between financial variables and accident rates. Later studies found that lower company profit margins are associated with slightly higher accident rates.

6 Barnett and Higgins found that the increased number of new entrants after deregulation explained a significant portion of the variance in fatality rates across time. Kanafani and Keeler, and Oster and Zorn found no such relationship when looking at the effect of new entrants on accident rates, not fatality rates. Rose found a large and significant negative relationship between carrier experience and accident rates.

7 Cumulative hours flown by the carrier were used in most analysis involving experience. This is more a measure of firm experience than pilot experience. Although correlated with pilot experience, to the extent that pilots acquire experience on the job, carrier experience does not take into account initial pilot experience or different rates of pilot turnover. Furthermore, different carriers will have different average pilot tenure, and average hours flown by each pilot.

8 Phillips and Talley (1992) use an ordered probit to model the extent of aircraft damage, given that an accident has occurred. They find that every 1,000 hours of additional pilot experience is associated with a decrease of 4-6 percent of a standard deviation of the probit index. Note, however, that one cannot say which parts of the distribution shifted. Their analysis focuses on commercial accidents and incidents. Phillips and Talley (1996) look at accident severity as measured by fatalities and injuries. Using OLS they find that an increase of 1,000 hours of pilot experience is associated with a decrease of 1.4 percent of the

In addition to the Phillips and Talley references, three other studies used accident-level data. Chance and Ferris (1987), Mitchell and Maloney (1989) and Borenstein and Zimmerman (1988) looked at the effect of accidents on the equity value of airlines. In the context of this paper this is important to the extent that firms should be willing to invest heavily in training and hire experienced pilots if (a) training and experience do lower accident rates or mitigate the effects of accidents and if (b) firms bear a heavy financial burden resulting from accidents and fatalities. To the extent that markets are efficient, a drop in equity value represents the present value of the true cost to the firm—including costs related to a possible fall in consumer demand. All three studies use roughly the same data.[9] Chance and Ferris find that the one-day average equity loss is 1.18 percent. They do not find a significant effect on industry equity, only on own-firm equity. Mitchell and Maloney find a similar number, 1.19 percent. However, they do a separate analysis for accidents in which the firm was found responsible and those which were out of the firm's control. The decline in equity for the at-fault crashes was 2.2 percent, while for the not at-fault it was 1.2 percent. Borenstein and Zimmerman find a reduction of 0.94 percent, which they say amounts to roughly $4.5 million in equity loss in 1985 dollars ($6.52 million in 2000 dollars).

Since there are no equity loss studies available using data from the 1980s and 1990s, a simple tabulation of equity loss was done for accidents using equity data from four major US carriers (American, Delta, US Airways and United), from 1983 to 1998. The tabulation consisted of comparing the average stock quote before and after the

---

proportion of fatalities or injured people. Again, the population is that of commercial accidents or incidents.

9 All studies use data from the 60s to the mid 80s. Borenstein and Zimmerman's data start in 1960s, while the others start in the mid-60s.

accident. In all, these carriers had 22 fatal accidents during the period of study. Due in part to the small sample size that none of the estimates are statistically significant. Nevertheless, their magnitudes are in line with those of the studies mentioned above. The one-day equity loss due to the accident was 0.8 percent, slightly lower than the three preceding studies. If one only considers accidents in which twenty or more fatalities were involved (12 accidents), the equity loss jumps to 1.8 percent. If one considers the most serious accidents, involving forty or more fatalities, the estimate increases once more, to 3.2 percent. These percent losses, if represented as dollar losses associated with common stock of the four carriers (as of December of 2000), translate to an average loss of 6.5, 14.7 and 26 million dollars. This suggests that the accident's magnitude, as measured by loss of life, may be key in determining loss of equity. However, one must again bear in mind that these figures are not statistically significant. A more extensive sample size would likely be required to draw inferences.

The equity loss results suggest that although there are costs associated with accidents, firms are mostly insulated from these costs. Indeed, this is evidence that demand is not responsive to accidents. From Borenstein and Zimmerman's estimates, we saw that the average loss is 6.25 million in 2000 dollars, or a little over $156,000 per fatality (average number of fatalities per accident is close to 40). Given the infrequency of air accidents, an average of one every two years for major carriers, this number is quite small.

II. Model

In this paper we address three types of safety outcomes. The first relates to the incidence of pilot error in flights that resulted in an accident. The second relates to the severity of accidents. The third relates to the probability of accident occurrence. The first two outcomes are relatively uncontroversial, and the limitations associated with modeling them are well known issues, such as endogeneity and selection bias. These issues will be discussed shortly. However, in order to model accident occurrence one would need a random sample from the population of all flights.

The AIDS data set is designed to be a record of all accidents. It contains no information on non-accident flights. Since the dataset does not include observations from pilots who did not have an accident, it is difficult to compare non-accident populations with accident populations. For example, if one is to estimate average flight hours for pilots who had accidents and compare that with average hours of those who did not have an accident, observations on non-accidents would be needed also. One is then faced with the daunting task of identifying observations in the accident data set which would be representative of the population of pilots at large and use these observations as "non-accidents", properly weighted. A few candidate sub-samples emerge. The first is the sub-sample of pilots who had accidents caused by mechanical failures, or the *mechanical failure* sub-sample (MF). Since mechanical failures are determined mostly by random processes, accidents in this category are hypothesized to be a random draw from the entire pilot population. There are, nevertheless, a few problems involved with this scheme. First, it is certainly possible that careless pilots may neglect their maintenance and be prone to mechanical failures. Second, mechanical failures may be more prevalent in some makes and models of aircraft, and make and model selection may be endogenous, in the sense that risk-averse pilots may to some extent select aircraft which are inherently safer to fly. This would not hold in the case of commercial aviation, where pilots likely do not influence the company's fleet composition. Lastly, a mechanical failure does not necessarily result in an accident, so that this sub-population would be composed of pilots who had mechanical failures *and were unable to avoid an accident.* Therefore, problems associated with the sample selection of the entire NTSB data set may still be present, albeit somewhat attenuated, in this sub-sample.

In the discussion above we mentioned that pilots in the accident data set may be "worse" pilots, in that they may on average be less prudent, less experienced, have fewer ratings and certificates, etc. A second way to select a sample which may be representative of the entire population would be to select only those accidents in which the pilot was not named as a cause of the accident by the FAA investigator. We will call this sub-sample the *no pilot error* sub-sample (NPE). In other words, select only those accidents in which the pilot did nothing wrong. The problem with this scheme is exactly the opposite of that with the first scheme. Whereas we may still be disproportionately selecting "bad" pilots with the first scheme, we may well be disproportionately selecting "good" pilots with this scheme. After all, these are pilots who were able to avoid mistakes in their management of the accident. A true representation of the underlying population may well lie "between" these two sub-samples.

The selection scheme above highlights the difficulty inherent in this type of sample selection problem. It also highlights that results from this selection process need to be interpreted with caution, since we are unable to guarantee that the sub-population selected is actually a good representation of the pilot population as a whole. With this in mind, we will proceed with the analysis and present results from both selection schemes. For the accident severity and extent of damage models we will present full-sample results in addition to results obtained from the two sub-samples. For the pilot error model we will present only the full-sample results. We will also use these selection schemes to attempt to estimate accident occurrence.

The aviation environment

Before proceeding with the description of the methodology and problems in modeling the three types of safety outcome, it is instructive to describe the environment in which pilots fly.

In this study I look at three types of flights, classified by which Code of Federal Regulation (CFR) chapter they are operated under: part 121 flights (operated under chapter 121 of the CFR), part 135 flights, and part 91 flights. Part 121 flights are comprised of larger commuter and all trunk carrier flights, and will be referred to as *trunk carrier* flights. Part 135 flights are smaller commuter flights of less than thirty seats and on-demand air taxi flights; these will be referred to as *commuter* flights. Part 91 flights are private flights, mostly recreational and non-commercial flights (although some small businesses will occasionally operate under this chapter of the CFR). These flights are referred to as *general aviation* flights.[10]

---

[10] Up to now, the aviation literature has focused, almost to exclusion, on trunk carrier and commuter flights, and has largely ignored general aviation. This bias is not without rationale: commercial flights involve firms engaged in profit maximizing behavior, while general aviation flights, by and large, consist of flights by private individuals. Nevertheless, since general aviation accounts for the majority of accidents (89 percent) and fatalities (92 percent) in American aviation, we also look at pilot training and experience in this setting.

In order to operate in each of these flight environments pilots must satisfy a number of very particular training requirements, culminating in "certificates" and "ratings" awarded to the pilots. In order to operate a non-commercial flight, a pilot will need a minimum of a pilot private certificate.[11] This certificate allows for operation of small aircraft in good weather conditions. A private pilot may then train to obtain an instrument rating, which would allow the pilot to fly in bad weather conditions. These conditions are referred to as *Instrument Meteorological Conditions*, or IMC. Without an instrument rating a pilot is restricted to flying in *Visual Meteorological Conditions*, or VMC. All flights conducted in IMC must follow a set of rules called *Instrument Flight Rules*, or IFR. Flights conducted in VMC may follow either IFR or *Visual Flight Rules*, VFR. However, flights conducted under IFR are accorded greater radar and air traffic coverage than VFR flights. In fact, IFR flights receive departure-to-landing air traffic coverage, thus commercial operators will almost always fly under IFR—regardless of the weather. VFR flights are common among general aviation pilots, but very rare among commercial flights. See Table 3.13 for a description of the specific requirements that define the existence of IMC and VMC.

In order to fly commercially, pilots will need a commercial certificate. Generally, this allows pilots to fly small commuter and air taxi operations, which are conducted under part 135 of the CFR. In order to fly large commuter or trunk carriers pilots will need an air transport certificated. An air transport certificate will allow pilots to fly under part 121 of the CFR.

---

[11] Pilots may also operate under part 91 with either a recreational certificate or a student certificate. Since these pilots account for a small portion of flights, and are too heterogeneous to compare with other pilots, we are not considering them in the analysis.

The principal models in this paper will be either probit or ordered probit models. However, since we suspect that there may be significant non-linearities involved, in particular with respect to the effect of hours flown on accident probability and severity, it is useful to first analyze these relationships in a non-parametric framework. The non-parametric technique used is Cleveland's (1979) Locally Weighted Scatterplot Smoother, or LOWESS.[12] LOWESS estimates will be done for the accident probability model, using each sub-population described above as randomly selected samples of non-accidents. These results will be weighted to account for the different sampling probabilities involved. For the probability of a fatality, conditional on accident occurrence, LOWESS estimates will be obtained by using the entire sample of accidents. In addition, estimates of the effect of total hours flown on fatality rates will also be presented for the two sub-populations.

Several factors are likely to affect accident occurrence, as well as accident severity. These can be classified into four categories: pilot training characteristics, other pilot-related characteristics, environment-related characteristics, and aircraft-related characteristics. Pilot training characteristics include factors such as (i) pilot experience, measured by hours of flight, composition of those hours, pilot age and (ii) other training, such as certificates and ratings. Other pilot-related variables include (i) pilot age and gender, and (ii) behavior variables, such as filing a flight plan, wearing seat belts, and obtaining an adequate weather briefing. Environment-related characteristics include (i) weather-related variables, and (ii) the type of flight (VFR/ IFR, commercial, scheduled). Plane-related characteristics include (i) the airframe age, (ii) plane size and power, and

[12] In all estimations a bandwith of 0.8 is used. Also, the lower and upper 2.5 percent of the sample are not shown in the LOWESS graphs.

(iii) time since last maintenance. Note that although pilot age can be viewed as a proxy for experience, it is also critical in another way, since younger men and women are more likely to survive an accident, and to survive without severe injury. Since pilot fatalities account for a large portion of fatalities (particularly in general aviation, where it accounts for 52 percent of fatalities), differences in the pilot's age may be picking up mostly differences in human survivorship. To account for this, the fatality dependent variables are calculated both to include and to exclude pilot fatalities, for the general aviation case.

The probability of pilot error is modeled with a probit regression. The dependent variable is one if the FAA investigator listed pilot error as a cause of the accident and zero if he did not so name him. For accident severity and aircraft damage, three separate models are estimated. The first is a probit for accident fatality, where the dependent variable is zero if there were no fatalities or one if there were one or more fatalities. The second is an ordered probit for the severity of an accident, where the categories are: *minor or no injury, severe injury*, and *fatal injury.*[13] The ordered probit coefficients tell us by how many standard errors the injury index's distribution shifts when an explanatory variable changes by one unit. Since this is not an intuitive way to think about accident outcomes (or any outcome, for that matter), we have reported both the ordered probit coefficients and the marginal effects coefficients.[14] The third model is a Tobit regression

[13] The NTSB actually codes accident severity with four categories, *no injury, minor injury, severe injury, and fatality*. Since we see very little qualitative difference between the first two, they have been combined into a single category.

[14] What we are interested in is by how much the probability of fatality, severe injury, and minor/no injury change when an explanatory variable changes. For the first category (minor/no injury) the marginal coefficient is $\frac{dP(y=0)}{dx} = -\phi(\bar{x}\beta)\beta$, for the second it is $\frac{dP(y=1)}{dx} = \left(-\phi(\bar{x}\beta) - \phi(\mu - \bar{x}\beta)\right)\beta$, and for the last it is $\frac{dP(y=2)}{dx} = \phi(\mu - \bar{x}\beta)\beta$. Their respective standard errors are $\phi_0(\gamma_0)V(\gamma_0)'$, and $\phi_1(\gamma_1)V(\gamma_1)'$, where $\phi_0 = \phi(\beta'X)$, $\phi_1 = \phi(\mu - \beta'X)$, $\gamma_0 = I - (\beta'X)\beta X'$, $\gamma_1 = (1-\mu)I - (\beta'X)\beta X'$.

of the proportion of fatalities in a given accident, where the truncation is at zero and one. A set of models for aircraft damage will also be estimated, where categories for the ordered probit are *no damage, minor damage, severe damage, and destroyed.* Accident occurrence is also modeled with a probit regression. Here the dependent variable is a dummy equal to one if the pilot was involved in an accident, and zero otherwise. This is the only estimator that needs to be re-weighted to account for the fact that we have no non-accidents in the sample. The re-weighing procedure is described below.

Previous mention was made of the importance of both estimating models for sub-samples that are more representative of all pilots. This is so because one cannot, *a priori*, observe pilots who are likely to be in an accident. It is reasonable to assume that pilots involved in accidents are likely to have fewer hours of flight experience and be less "prudent". Furthermore, they are also less likely to be flying a commercial flight, and less likely to be flying a larger power or larger capacity plane. Parameter estimates of the above variables, based on this sub-sample, may be inconsistent, to the extent that one of these variables (probably most of them) will be correlated with the probability of having an accident. On the other hand, one must recognize that sometimes the question of interest is the probability of fatality, <u>conditional</u> on accident. For example, if one seeks to implement safety enhancements designed to reduce fatalities, the target population may well be those who are likely to have accidents. Therefore, we obtain conditional estimates, by using the full sample of observations. We also obtain estimates more closely representative of the entire pilot population, by estimating the model for our two sub-samples: NPE and MF. The equations being estimated are:

Pilot error probability

(1) $$P(error) = \Phi(\beta_1 T + \beta_2 X + \beta_3 Y + \beta_4 Z)$$

Severity outcome equations

(2) $$P_i = \Phi(\gamma_1 T + \gamma_2 X + \gamma_3 Y + \gamma_4 Z)$$

(3) $$\begin{aligned} P_1 &= 1 - \Phi(\delta_1 T + \delta_2 X + \delta_3 Y + \delta_4 Z) \\ P_2 &= 1 - \Phi(\delta_1 T + \delta_2 X + \delta_3 Y + \delta_4 Z - c_1) - P_1 \\ P_3 &= \Phi(\delta_1 T + \delta_2 X + \delta_3 Y + \delta_4 Z - c_2) \end{aligned}$$

(4) $$P_f = \eta_1 T + \eta_2 X + \eta_3 Y + \eta_4 Z + e$$

Accident probability

(5) $$P(accident) = \Phi(\beta_1 T + \beta_2 X + \beta_3 Y + \beta_4 Z)$$

Equation (1) is the probit equation for pilot error; equation (2) is the probit equation for severity of the extent of aircraft damage; equation (3) is the ordered probit equation for the same; equation (4) is the tobit equation for proportion of fatalities and proportion of accidents in which the aircraft was destroyed; the truncation is at 0 and 1. Finally, equation (5) is the probit for accident probability.

In the above equations we have: P_i denotes the probability of outcome *i*, either a fatality (in the case of the severity models) or aircraft destroyed (in the case of the extent of damage models). P_f denotes the proportion of fatalities in a given accident. *T* is the vector of pilot training characteristics, *X* is the vector of other pilot-related characteristics. *Y* is the vector of environment-related characteristics and *Z* is the vector of aircraft-related characteristics.

Parameter estimates for different types of flights (general aviation, trunk carriers and commuters) are likely to be quite different. To account for this, separate models are estimated separately for each category.

Weighing procedure

Two different probit models are estimated for accident probability. In the first we use the MF sub-sample as the sample of non-accidents. In the second we use the NPE sub-sample as the control group. The weighing procedure then consists of increasing the weight of the sample we treat as non-accidents, so that they reflect the inverse of the proportion of non-accident pilots. Weights for accidents are equal to unity since the AIDS is a record of all accidents. On the other hand, non-accidents have very large weights, typically in the high hundreds. This reflects (1) the small proportion of pilots who have accidents in any given year, and (2) the small number of observations in the non-accident control group (11 percent for the NPE and 13 percent for MF). Both the NPE sub-sample and the MF sub-sample are re-weighted.[15]

[15] In order to obtain adequate weights for these sub-samples, we must rely on the FAA's estimates of accidents per flight hour, which are tabulated by year and by CFR classification (trunk carrier, commuter and general aviation). We then convert the FAA's estimates of accidents per hour to accidents per pilot, by multiplying accidents per hour by the average number of hours flown by pilots, again, by year and CFR category. The estimate of the average number of hours per pilot is obtained from the accident/incident databases' question "hours flown last ninety days" and by expanding that value to 365 days. We now have an estimate of the average number of accidents per pilot, by CFR type and year. Since the accident/incident is a complete record, we know what the number of accidents is. If we assume that each pilot has at most one accident per year, we can then compute the number of non-accidents by subtracting the number of accidents from the number of pilots. We know how many observations are in our two sub-populations of non-accidents. Therefore we can compute their sample inclusion probability by dividing the number of observations in the data set by the estimate of the number of non-accidents in the population. This probability is then scaled to account for missing observations (observations for which some of the covariates were not recorded or were not validly coded are considered as missing). The scaling is done by multiplying the inclusion probability by the proportion of valid (non-missing) observations. Note that the sub-populations are not eliminated from the original accident sample. Rather, the entire accident sample is

III. Data

In the introduction the data was briefly described. Here we describe the data in detail. The data used are aviation accidents from the NTSB's Accident/Incident Database System for the years 1983-1998. The dataset includes pilot-related variables such as age, gender and pilot experience. The most important pilot experience measure is pilot flight time.[16] Pilots with more flight time are less likely to be involved in an accident or incident. Also, one can postulate that these pilots would also be able to better manage an accident, mitigating the extent of accident damage or severity. It is reasonable to expect the effect of flight time to be plane-specific and to expect that proportion of time spent as pilot-in-command (PIC) may have a separate impact on outcomes than total time. Thus, four variables for flight time are included: total hours of flight, proportion of time in the make and model flown, and the proportion of total hours as PIC. For general aviation equations, total hours are entered in spline form, in order to capture the dramatically different slopes associated with the first few hundred hours of flight versus higher levels of flight time.[17] In the case of the probability of pilot error, a visual inspection of the top panel of Figure 3.2 reveals a clear node at approximately 1,800 hours. For the fatality,

---

preserved, and the sub-sample is added back in with the weights as above, in effect replicating the sub-sample. Finally, the weights in the analysis are then just the inverse of the estimated sample inclusion probabilities that have just been calculated. In other words, let h_{ij} be the number of hours flown in year i and flight type j; let a_{ij} be the number of accidents, and p_{ij} the total number of pilots; let q_{ij} be the number of pilots without an accident and let s_{ij} be the number of accidents (or pilots) in our non-accident population. Then the sampling weight for the non-accident population, w_{ij}, is obtained by:

(1) $a_{ij}/p_{ij} = a_{ij}/h_{ij} \times h_{ij}/p_{ij}$

(2) $p_{ij} = a_{ij}/p_{ij} \times a_{ij}$

(3) $q_{ij} = p_{ij} - a_{ij}$ (if we assume at most one accident per pilot)

(4) $w_{ij} = q_{ij}/s_{ij}$

[16] Flight time variables available include: total time, total time as pilot in command (PIC), total time in the last 90 days and in the last 30 days. Furthermore, separate tabulations are available for all aircraft, and for the make and model flown.

severity and extent of damage equations—equations (2)-(4)—Figures 3.3 and 3.4 reveal a node at 800 hours. For the weighted accident probability probits seen in Figure 3.5, the node is clearly shown at 1,200 hours. In most of the commercial flights a visual inspection of the non-parametric results suggest that flight hours and hours squared and cubed will adequately capture non-linearities. The only possible exceptions are for fatality probability and accident probability for trunk carriers (see bottom panel of Figures 3.4 and 3.5). However, when the commercial models were estimated with a spline they did not perform as well as the polynomial specification. Therefore, we kept the polynomial specification.

Pilot training

Before proceeding with the model, it is useful to explain the training variables we are using and how they relate to the aviation industry. The main training variable in the study is total hours flown and, to a lesser extent, the proportion of those hours as pilot-in-command. As in the conventional labor literature, these variables reflect the sum total of all pilot training, as well as the effect of on-the-job training. We also have three certificate variables: private pilot certificate, commercial certificate and air transport certificate.[18] We also know if the pilot is instrument-rated.

[17] Splines are estimated subject to the restriction that the function be continuous at all points.

[18] All certificates require a combination of ground and flight school. They also require passing both a written and a practical exam. An instrument rating requires approximately 120 hours of flight training. In addition, there is a minimum number of hours in instrument conditions, a minimum number in cross-country flight, and a minimum number of landings performed in instrument conditions. A commercial certificate requires logging 250 hours, which must include a minimum amount of hours of flight instruction, a minimum number as pilot in command. It also has requirements as to time spent in various types of flight (cross country, IFR flights, etc.), as well as number of takeoffs and landing. It also requires hours in different types of planes. Air transport certificates require a minimum number of flight hours logged in addition to their training requirements, which are the most detailed and stringent.

The private pilot certificate is the lowest form of certificate and everyone in our sample must have at least this certificate. The instrument rating trains pilots how to fly in adverse weather conditions, or IMC; it is required of all commercial pilots. However, one must note that pilots who do not have the instrument rating are limited to fly in good conditions. Therefore the instrument rating has two effects. The first is the training effect mentioned above; the second is to eliminate the flight conditions restriction, so that the pilot may now fly in low/no visibility conditions. It is not possible to determine what the net effect will be. The commercial certificate is required for pilots who fly commercially, but are not involved in transporting passengers. These certificates are of limited use because they are required for entry into the different parts of the industry; however, they are the only measure of direct training available. As mentioned, every trunk carrier pilot must have an air transport certificate. Therefore, this certificate measures the lowest possible training level of trunk carrier pilots. Similarly, depending on what activity the firm is operating (transporting passengers or transporting cargo), pilots of commuter flights must have either an air transport certificate or a commercial certificate. Here, there is some variability in the type of certificate, but not much. Commuter pilots also, by and large, have instrument rating. The only instance in which we can observe the effects of the training associated with these certificates is in non-commercial flights, or general aviation flights. The drawback here is that we have no reason to suspect that training designed for the commercial arena will have the same effect on the personal and recreational arena, which is characteristic of general aviaiton.

Most training of commercial pilots' training is not represented by the training certificates in the NTSB database. In fact, training certificates are usually obtained prior to a pilot's employment, either as part of a two or four-year undergraduate program, or in an aviation course designed to award the pilot the particular training certificate. This training is mostly paid by the pilot, although airlines will often subsidize a pilot's training and offer them employment after graduation. Commercial pilots—and trunk carrier pilots in particular—undergo ongoing training throughout their careers, mostly paid in full by the firm.[19] Pilots are required to pass a "check ride" twice a year, conducted by the FAA. Furthermore, for commercial flights, every time a pilot is to fly a different make and model aircraft he is trained and certified in that aircraft. This consists of ground school, simulator flight hours, as well as real-world hours. All of this training is not represented in the certificates, although it is represented by the total flight hours and by the proportion of flight hours as pilot-in-command.

The last pilot-related class of variables is those associated with pilot behavior. Among these we include the use of seatbelts and shoulder harnesses, the filing of a flight plan,[20] and obtaining of a preflight weather briefing.[21] Note that both the weather-briefing and the seat belt and harness variables are also correlated with pilot "prudence", and in this sense may belong in the pilot error, accident severity and extent of aircraft damage equations. To account for this we include the flight plan dummy variable in all equations (indicates if a pilot filed a flight plan) to capture pilot prudence. This is only

[19] Of course we know from the training literature that whether the firm or the employee pays for training is not relevant.

[20] The FAA recommends that pilots always file a flight plan. A flight plan, even when not required, allows officials to track a flight and provides almost immediate warning of a missing plane. It is, however, of little use in preventing an accident or mitigating its immediate consequences. Therefore, the flight plan variable is used only as a proxy for pilot prudence, to the extent that prudent pilots will file more often than less prudent pilots. This discussion is not relevant for commercial aviation, since virtually all flights are conducted under IFR rules, and thus require a flight plan be filed.

[21] It should be noted that often when a fatality is involved the investigator is unable to determine if a weather briefing was obtained.

applicable to general aviation flights, as the vast majority of commercial flights file flight plans.

Environment-related variables are important factors in determining accident outcomes. The dataset includes very detailed environment-related information. The first, and perhaps most important, is weather. To control for weather, a dummy for the existence of VMC is used. This is expected to be highly correlated with pilot error, to the extent that pilots are more likely to commit mistakes in poor weather conditions. It is also certainly important for other outcomes as well as for accident rates. One expects that flights conducted in VMC will be less prone to accidents, and conditional on an accident, may be less prone to fatalities or severe injuries. The same is true, to a lesser extent, of night flying. Flights conducted during daytime enjoy greater visibility and are less susceptible to spatial disorientation. A dummy for night flying is included in all equations.

Since scheduled flights may be associated with different accident rates, a dummy variable for scheduled flight is also included in the selection equation. This is only available for the commuter and trunk carrier equations, since there are no scheduled flights conducted in general aviation.

Characteristics associated with the aircraft being flown will also affect accident rates and outcomes. In particular, larger aircraft tend to be safer than smaller aircraft. On the other hand, aircraft with more seats, *ceteris paribus*, are more likely to have at least one fatality or at least one severe injury. Thus, a continuous variable for the number of aircraft seats is included.[22] For pilot error, we postulate that larger aircraft have safeguards to reduce the incidence of pilot error. A continuous variable for airframe age, in years, is also included, as these older aircraft may be more susceptible to accidents, and may be more dangerous than newer aircraft. To capture the effect of aircraft maintenance on accident probability, a continuous variable for time since last inspection is also included in the accident equation. We must point out that this does not account for the fact that some planes will log more hours than others, and a greater time between inspection is not necessarily evidence of maintenance neglect. Lastly, pilots who fly more often are more likely to be involved in accidents, therefore the number of landings performed in the last 90 days is also included in the accident equation. Since this variable may also be considered to be a measure of recent experience, it is included in all equations. In summary, the following is a list of variables included in each estimation:

Total hours and spline (general aviation only)
Total hours flown (commercial equations only)
Total hours flown squared (commercial equations only)
Total hours flown cubed (commercial equations only)
Proportion of hours in make and model
Proportion of hours as Pilot in Command (PIC)
Gender
Age
Dummy variable for a seat belt being used (severity only)
Dummy variable for a harness being used (severity only)
Dummy variable for no flight plan filed (general aviation only)
Pilot is instrument rated (general aviation only)
Pilot has Commercial certificate (general avaition only)
Pilot has Air Transport certificate (commuter and general aviation only)
Dummy variable for VMC
Aircraft size, as measured by seats
Airframe age
Number of landings in last 90 days

[22] We could also include aircraft size, as measured by the total pounds of thrust that the aircraft is capable of. However, since these variables are highly collinear, we opted for inclusion of only the first of these two variables.

The NTSB data set contains 43,279 observations, from 1983 to 1998. However, we only use observations for which sampling weights can be obtained. Since the FAA only reports accidents per hour of flight for commuter (part 135), trunk carriers (part 121) and general aviation (part 91), only these types of records are used.[23] In order to reduce the heterogeneity in pilot characteristics, recreational, student, military and foreign pilots are excluded from the sample, as were accidents involving homebuilt airplanes. An additional concern arises because of serious underreporting of aviation incidents. Although we refer to all observations in this analysis as accidents, a small proportion of the AIDS observations are what is referred to as an *incident*. Incidents can be thought of as mild accidents.[24] This is particularly serious because the underreporting may be systematically related to pilot experience or training, although one cannot be sure in which way this correlation would go. One option to deal with this problem would be to do the analysis for aviation accidents only. Unlike incidents, aviation accidents are hard to conceal, and severe penalties are applicable for those who attempt to do so. Underreporting of aviation accidents would then seem much less of a problem than underreporting of aviation incidents. Since there are relatively few incidents in general aviation (1 percent) or commuters (12 percent) flights, little would be lost in terms of precision, and thus only accidents are used in these equations. This is not true of trunk carriers, where incidents account for the majority of events (54 percent). Furthermore, since aviation accidents are defined as events that either substantial aircraft damage or

---

[23] These flights account for 93 percent of all flights. Flights conducted under other chapters of the CFR, such as agricultural flights, or international carriers, are excluded.

[24] The NTSB defines an accident as "an occurrence associated with the operation of an aircraft which takes place between the time any person boards the aircraft with the intention of flight and all such persons have disembarked, and in which any person suffers death or serious injury, or in which the aircraft receives

severe or fatal injury, the dependent variables in the severity and damage extent equations would be incidentally truncated. Therefore, for trunk carriers separate estimates are done for accidents and incidents and only for accidents. It should be noted that for incidents pilots only report total flight hours, and do not report hours as pilot in command or hours in make model. The main trunk carrier estimates, done with incidents, will not have estimates for the proportion of time as pilot in command or the proportion of time in make and model. The final dataset includes 40,921 observations. However, when only observations with no missing values are used, the sample size falls to:

	General Aviation	Commuter	Trunk Carrier
Pilot error/ Fatality/Severity models	12,851	1,083	671
Extent of damage models	15,649	1,253	751

IV. Results

Summary statistics are presented in Table 3.1. We see that characteristics vary significantly between the three aviation categories. First notice the difference in the accident rates between the three categories. In this sense, commuters and general aviation flights are very similar, with an accident rate of 18 ½ and 19 ½ accidents per 1,000 pilots per year, respectively. Trunk carriers have a much lower accident rate of 1.7 accidents per 1,000 pilots. One must note that these accident rates are <u>per pilot</u> and not <u>per hour</u> or per takeoff/landing. Since commercial pilots fly more often than do general aviation pilots, their accident rates <u>per hour</u> will be substantially lower. Among general

substantial damage." An incident is defined as "an occurrence other than an accident, associated with the aviation pilots we see that half have instrument training, 35 percent hold commercial certificates and 10 percent hold air transport certificates. Among commuter pilots, 64 percent have commercial certificates and 57 percent have air transport certificates. All of them hold either of the two certificates, a requirement for operation under part 135. All are instrument rated. All of trunk pilots are instrument rated and hold Air Transport certificates, which are required for operation under part 121. Commercial pilots have more experience than general aviation pilots do, with an average of almost 5,800 hours for commuter pilots and almost 13,000 for trunk carriers versus an average of 2,600 hours for their general aviation counterparts. However, the proportion of hours in make and model is lower for commercial pilots (0.21 for commuters, 0.27 for trunk carriers) than for general aviation pilots (0.32). This is to be expected, since general aviation plane turnover is much lower than commercial plane turnover. The percentage of hours as PIC is higher for commuters (86 percent) than general aviation (79 percent), but it is much lower for trunk carriers (59 percent).

A cursory inspection of the fatality and damage means in Table 3.1 shows that these are much worse for the general aviation and the commuter population than they are for trunk carriers. Trunk carriers have a very low fatality rate: only four percent of trunk carrier accidents result in fatalities. Compare this to 20 percent for both general aviation and commuters. Other safety outcomes mirror this result.

When considering accident rates, one must, nevertheless, keep in mind that these are accident rates per pilot, not per hour. Therefore, *ex ante*, one could expect these rates to be higher for commercial pilots—both commuter and trunk carriers—given that they have far more flight hours than do general aviation pilots. This is not the case.

---

operation of an aircraft, which affects or could affect the safety of operations."

Non-parametric results

Before reviewing the discrete and OLS models, it is instructive to look at non-parametric results. Non-parametric results are important, for they capture the non-linear functional relationship between safety outcomes and explanatory variables. In this case we are interested in the non-linear relationship between pilot experience, as measured by total flight hours, and pilot error, fatality rates, and accident rates. For fatality rates we present both full-sample results and results for the sub-samples, which throughout this paper have hypothesized to be more representative of all pilots—not just pilots with accidents. As mentioned previously, the non-parametric technique used is LOWESS. The main advantage of LOWESS is that it is robust to outliers and "far out" response variables.[25]

In Figure 3.1 we can see the estimated distribution of flight hours, for both general aviation and commercial flights. We should note that all non-parametric results are obtained by using a LOWESS bandwidth of 0.8. Also, since there is enough heterogeneity between the three types of flights being studied, separate estimation was conducted for each. One first sees that average hours are higher for trunk carriers than for commuter, which in turn are higher than general aviation pilots. The distribution of the general aviation pilots also seems to be much more concentrated than the other two categories (at about 800-1000 hours), with the distribution of trunk carriers the most being spread out. Commuter hours peak at 3,500, but fall slowly thereafter.

[25] Median smoothing techniques are also resistant to outliers. Results using median smoothing produced similar results.

In Figure 3.2 we see the first bivariate non-parametric result. The top panel shows the non-parametric relationship between pilot error probability and thousands of flight hours for general aviation. The middle panel shows the results for commuter aircraft and the bottom panel for trunk carriers. For general aviation there is a monotonic relationship, with most of the gains in pilot error coming within the first two thousand flight hours. The returns decrease markedly after the first 1,800 hours. For commuters we also see considerable gains concentrated around the first few thousand hours of flight. However, for commuters the gains not only level out, but actually reverse themselves at higher levels of flight hours. The trunk carrier panel, again, shows marked reductions in pilot error probability associated with the first few thousand hours of flight. In this case the decrease is not as dramatic as in, say, the general aviation case, and is not concentrated along the first couple thousand hours. The gains persist linearly until 6-7 thousand hours of flight, then level off (or even increase slightly).

Figures 3.3 and 3.4 present the results for the probability of a fatality and the probability of the aircraft being destroyed. In both figures estimates are presented separately for three categories: (1) all observations, (2) only observations from the MF sub-sample, and (3) only observations from the NPE sub-sample. For general aviation we see that both the slopes and the levels are smaller as you progress from (1) to (3). This is consistent with the hypothesis that the MF strategy may still over-sample "bad" pilots, but that the NPE strategy may in fact over-sample "good" pilots. However, what is most remarkable is that flight hours tend to (at least initially) increase both the probability of fatality and the probability of the aircraft being destroyed. This initial effect is seen until roughly 1,200 flight hours. There can be several explanations for this.

On

mo

exp

eit

sug

spl

inf

inc

sar

ext

pro

see

sar

inc

3.4

Fig

ide

sin

ger

rou

rate

One is that as rookie pilots gain experience they may become overconfident and thus more prone to be involved in a serious accident. Another is simply that as a pilot gains experience, he may better able to eliminate minor accidents than major accidents. In either case one would observe the initial increase in fatality rates. The top panels also suggest that for the multivariate analysis, pilot experience would be well modeled by a spline, with a node at 1,200 flight hours.

For commercial flights one observes a myriad of patterns, none of which are very informative. For trunk carriers fatality rates increase, then decrease slightly, and then increase again. For commuters we see a slight increase for both the MF and NPE sub-samples, but a slight decrease for the entire sample. For trunk carriers and—to a lesser extent commuters—there is a marked non-linearity. However, when one considers the probability of having the aircraft destroyed, the patterns change. For trunk carriers we see a monotonic decline in probability. The decline is more pronounced for the MF sub-sample. For commuters both sub-samples show an initial decline followed by an increase, and a monotonic decline for the full sample estimate. Again, Figures 3.3 and 3.4 suggest a cubic polynomial for the commercial flights.

Accident probability results are presented in the last non-parametric Figure, Figure 3.5. Two estimates are presented in each panel, one using the MF population to identify non-accidents and the other using the NPE population. Both techniques deliver similar results, with the exception of commuter pilots, for which results are mixed. For general aviation we see a marked decrease in estimated accident rates, corresponding to roughly the 1,200 hours of flight. After that there is little or no reduction in accident rates (NPE), and a slight increase in accident probability (MF). This result motivated modeling accident probabilities for general aviation with a spline function, with a node at 1,200 hours. For commercial pilots we likewise see a large initial reduction in accident rates, albeit less dramatic. For trunk carriers the effect seems to level off at about 6,000 hours. However, for the MF specification, there is a muted decrease, followed by a slight increase and another decrease. For commuter flights there is little discernable trend for the MF specification, although the NPE specification shows a sharp decline, followed by a dramatic increase. Based on these results both the trunk carrier and the commuter equations will be modeled with a cubic polynomial.

Multivariate results

The most obvious effect of training and experience on safety is its effect on pilot error. One could indeed argue that other outcomes are in fact a function of the reduction in pilot error. Therefore, the first result discussed is the pilot error probit results of Table 3.2. The flight hours results strongly mirror the non-parametric results we saw previously. For general aviation, for every thousand flight hours we see a corresponding reduction of 6.2 percentage points in the probability of pilot error. This number decreases to 0.2 percentage points after the first 1,800 hours. In addition to the magnitude, the results are very significant, at any conventional level. The other training variables tell a similar story. We see that it is not only the total flight experience which is driving the changes in pilot error, but also the proportion of those hours spent as pilot-in-command, and the proportion spent in the make and model concerned. For the proportion in make and model, we see that a 0.3 increase in the proportion of time in make and model (say, from 0.5 to 0.8) is associated with a one percentage point decrease in the probability of pilot error. For the proportion of time as pilot-in-command the magnitude is twice as big—and in both cases the results are extremely significant (at the 1 percent level).

General aviation presents an interesting medium in which to evaluate certificates and ratings. This is because there is very little variability in these measures outside general aviation: commercial flights require certain certificates and in all cases require an instrument rating. The first pattern to observe is that all certificates and training variables are associated with lower incidences of pilot error. These results are significant for both the instrument rating and the air transport certificate. For the instrument rating, we find that instrument-rated pilots are 1.7 percent less likely to be at fault in an accident. For the air transport certificate the magnitude is somewhat larger: pilots with this certificate have a probability of pilot error 2.2 percentage points smaller than those without it. The last training-related variable is age. Although older pilots are in fact more likely to be at fault in an accident, the result is modest: one percentage point less more likely for every 10 years.

The experience and training results for commercial aviation are mixed. For commuters we see no significant effect of increased pilot experience of pilot error probability. Furthermore, we also find no evidence that the makeup of these hours affects pilot error either. The only training variable that is significant for commuters is the air transport certificate. Pilots with this certificate are almost eight percentage points less likely to be at fault in an accident. This is a bit puzzling, since we do find significant experience results for trunk carriers, all consistent with what one would expect. It is certainly possible that experience and training affects pilot error through a censoring mechanism. For example, more experienced pilots may be able to eliminate less severe accidents—where pilot error is not at play—then these observations would be censored in our data. This is, however, at odds with the results from the two other flight groups, and we have no a priori reason to believe that censoring would occur only for commuter pilots.

Table 3.2 is the only table in which results for trunk carriers are separated for accidents and incidents and for accidents alone. In all other estimations we present the results for accidents + incidents together. This is in part because the small number of accidents does not allow us to draw significant inferences. However, in the case of pilot error we see that the results for both accidents + incidents and for accidents alone are quite similar. They are more significant for accidents + incidents, no doubt because of the larger number of observations, but are the same sign in both cases. The only difference is that in the case of accidents alone the results are larger in magnitude—roughly twice as large across the board. In both cases we see large returns to experience. In order to better visualize these effects, the predicted probability of pilot error is graphed for different levels of flight hours in Figure 3.6. In the bottom panel, which corresponds to trunk carriers, we can see that the first four thousand flight hours cuts the probability of pilot error in half—from 60 percent to 30 percent. This is encouraging and speaks to the potentially large effects of experience in reducing pilot error.

It is unclear exactly how filing a flight plan could affect pilot error. Although a flight plan requires some planning, and thus may reduce the workload on the pilot in-flight, one must be mindful of the endogeneity of this variable. Pilots who file a flight plan may in fact be more cautious pilots, and therefore the variable may be picking up prudence. In any case, filing a flight plan certainly reduces the incidence of pilot error: general aviation pilots who file a plan are almost five percentage points less likely to be at fault in an accident.

Lastly, although weather is certainly not the focus of this paper, we cannot ignore the enormous magnitude of the VMC variable. Pilots who fly in VMC are much less likely to err. This result is extremely significant, and is the only result that is true across all flight categories: general aviation and commercial flights. The magnitude of this effect varies from 19 percentage points to 31 percentage points—in all cases a remarkable and sobering result.

A second aspect of training and aviation safety is the role of training in mitigating the results of an aviation accident. These results refer to both aircraft damage and to personal injury. It is reasonable to expect that some facets of pilot training may better prepare a pilot to deal with an accident, and in so doing reduce the severity of the accident. In the non-parametric section we saw that in some cases experience tends to increase the probability of a fatal accident. Nevertheless, one must always keep in mind that, for our purposes, there are three categories in both the severity of injury and the extent of damage analysis. This means that an increase in one category may shift probability into either of the other two categories. We deal with this, to a certain extent, by using an ordered probit and reporting the marginal effects for each category. Since there are only three categories, interpretation of the ordered probit is not as complicated as it is in the case of multiple categories. If two adjacent parameter values are positive (negative) that implies that the net effect of a marginal change is to shift probability mass out of (into) the remaining category. Lastly, we estimate the proportion of fatalities in each accident. As can be expected, this variable is truncated at 0 and 1, since the most severe accidents will likely kill all passengers and the milder accidents will not kill anyone. Therefore a tobit estimator is used for the proportion killed.

Tables 3.3 and 3.4 present the general aviation probit and ordered probit results for accident fatality and severity, respectively. In both tables flight time is modeled with a spline. The spline node is at 800 hours, so that the slope coefficient is allowed to differ, before and after 800 hours. Although the flight time estimates, by and large, confirm the non-parametric results (a sharp increase, followed by a modest decrease), they are not significant. In fact, the only significant result is an initial (up to 800 hours) very significant negative effect of hours on the probability of a fatality in the NPE sub-sample. This effect is a 0.6 percentage point decrease in the probability of an accident for each hundred hours up to 800. The only training results that are consistent in sign across all specifications are the hours makeup variables: proportion of time as pilot-in-command and the proportion in make and model. For the proportion as pilot in command we see that a 0.1 increase in the proportion is associated with a half percentage point decrease in the probability of a fatality, for both the full sample and the MF sub-sample. Lastly, we see that, as in the pilot error equation, having an air transport certificate is significant, and is associated with a four percentage point reduction in the probability of a fatality. However, this last result is not robust to the sample type, and is only present in the full sample.

An inspection of the ordered probit results largely validates the probit results, with one important exception.[26] Hours beyond 800 have a negative and significant affect on fatality probability, when the analysis is done using the entire sample. Here we see that an increase in flight hours beyond 800 tends to shift probability mass out of both the "fatality" and the "severe injury". In other words, it appears general aviation pilots, conditional on accident occurrence, are able to transform some fatal and severe accidents into those with only minor injuries or no injuries. With respect to both the proportion of time as pilot-in-command and the air transport certificate, again, the shifting is from fatality and severe injury to the more innocuous "minor/no injury" category. Since the ordered probit coefficients for both the "fatality" and the "sever injury" categories are comparable, the severity outcomes may actually be best modeled by a probit of fatality or severe injury. This is done and the results are reported in Table 3.9. The results are somewhat different in that in the fatality/severe injury probit, the returns to experience are much larger and significant. On the other hand the air transport certificate is not as significant. Otherwise, the results are comparable.

The last interesting result in Table 3.3 is pilot gender. Male pilots are four percentage points more likely to be in accidents with fatalities. It is possible that this figure is picking up survivability differences between the male and female pilot. If women do in fact survive accident trauma more than men do, the parameter estimate may be picking up this differential. To deal with this, the fatality model is estimated for non-pilot fatalities only. This way, the only link between the pilot and the fatality would be through passengers, eliminating the survivability issue. However, for this second model

---

[26] Ordered probit results are only reported for the "all observations" category, and not for the NPE and MF sub-samples.

the results are identical—a 4.2 percentage point increase in fatality rates associated with a male pilot. This suggests that the mechanism has little to do with survivability, and more to do with other differences in male and female characteristics.

The extent of damage table (Table 3.4) echoes the previous results, if attenuated. Here we see that, again, the proportion as pilot in command reduce the probability of an aircraft being totally destroyed. We also see a negative coefficient on the air transport certificate variable, and—again—a positive result for the male dummy. The results are only significant in the full sample. Nevertheless, they are similar across all three samples (full sample, MF, NPE), and the magnitudes are similar to those in Table 3.3.

At first, the result of the instrument training variable may be somewhat of a paradox. We see that pilots with that rating, even controlling for weather and other conditions, seem to be associated with higher probabilities of aircraft being destroyed. Although this same result was not statistically significant for the fatality equations, it was of the same sign and similar magnitude. One would expect the opposite: pilots who receive this training would be better able to manage an accident and mitigate the nefarious effects of the accident. While this may be true if we randomly assigned instrument-trained and non-instrument-trained pilots into accident situations, this is certainly not the case here. An IFR rating not only indicates that the pilot received the training, it also allows the pilot to legally fly into hazardous conditions which his non-rated counterpart would not legally be able to fly in. Many of these hazards may not be properly picked up by the weather and VMC control variables. Therefore, we may in fact be observing the effect of the removal of a safety restriction on safety outcomes, rather than the effect of pilot training. If this last hypothesis is correct, then the IFR results seen in Tables 3.3 and 3.4 are perfectly reasonable.

When interpreting the other training variables—commercial and air transport certificates—one must be weary that these variables are extremely endogenous. Since these certificates are required by certain aviation occupations, they are highly correlated with pilot occupation in the airline business.[27] And in so being, they will be correlated with a myriad of other characteristics that professional pilots have and others don't. To the extent that these characteristics are, at best, imperfectly measured in the variables available to us, they will be heavily biased due to an omitted variable problem. That much said, there is no surprise that their signs are always negative (i.e. those who have the certificate have lower fatality probabilities and aircraft destroyed probabilities), of large magnitude, if of varying degrees significant. In both cases, the presence of a certificate shifts probability mass away from the most severe outcome, and increases probability mass in the severe injury and minor injury categories.

The commercial equation results (Tables 3.5-3.8) have fewer significant coefficients than the general aviation equations. Part of this is due to the fact that there are fewer commercial observations. However, when present, the results tend to be as expected. Flight hours turns out to have, generally, the correct sign. That is, more hours are associated with less severe outcomes. In commuter equations this effect is insignificant. A test for the joint significance of all hours variables shows that they are only significant for the trunk carriers and only, and only for the the NPE sub-sample. The exact shape of the hours curve depends on the sample, but tends to follow the non-

[27] For example, pilots are required to carry an air transport certificate in order to operate under parts 121 of the CFR.

parametric results. The proportion of time as pilot in command and the proportion of time in make and model have a negative impact on severity outcomes for commuters, but neither is statistically significant—in both the severity and the extent of damage equations.

We have omitted both the commercial certificate and the instrument rating variable from the commercial equations. We included the air transport certificate in the commuter equation. The coefficient, as expected, is negative (i.e. indicating a reduction in severity and extent of damage outcomes) but in all specifications is not significant.

The last set of accident severity results presented is those in Table 3.10. Table 3.10 does not model the probability of fatality (or another outcome category). Rather, it models the proportion of people who were killed with a tobit. The estimates are similar to the probit estimates. Results for the full sample are reported. For general aviation, again we see large albeit statistically insignificant flight hours results. These are in line with both the probit and the ordered probit severity results. Also, again we see a large and significant effect of proportion of hours as pilot in command, with magnitudes in line with out probit estimates. The same is true for the proportion of hours in make and model. In the commercial flights equations we see an absence of any flight hours effects, a result similar to the full sample probit results. In fact the only significant results are those for VMC and night flying for the commuter equation—an effect ubiquitous throughout all of our analysis results. For trunk carriers the only significant result is a negative coefficient on male dummy. This is at odds with the positive coefficient found in most of the general aviation equations, and suggests that gender plays very different roles in outcomes for the two types of flights.

Tables 3.11 and 3.12 present the probit results for the accident probability. We left these tables for last as the identification strategy used to obtain a sample of "non-accidents" is contentious. However, with this in mind, we will discuss the results.

For general aviation one can see that all the experience-related variables are significant at conventional confidence levels. One can also observe that the effects are in the expected directions. Total hours of experience reduce the chance of an accident by 0.6 percentage points (MF) and 1.0 percentage points (NPE) for every thousand hours of flight time up to 1,200. Thereafter the reduction is much smaller, at 0.2 percentage points per thousand flight hours for both NPE and MF. Controlling for total hours, a 0.1 increase in the proportion of time as pilot in command reduces the probability of an accident by 0.1 percentage points (MF) and 0.11 percentage points (NPE). The same is not true for proportion of time in make and model; here we see an increase of 0.2 percentage points associated with an increase of 0.1 in the proportion in make and model for the MF strategy. For NPE there is no significant effect. Age also has an effect on accident probability, but not in the expected direction. Older pilots are associated with a higher accident rate, when one includes flight time and other training variables. This is, however, consistent with the hypothesis that almost all training is already captured by other variables, and that age is simply measuring the effect of aging on accident probability. If this is the case, then the direction of the effect is perfectly understandable. In any case, the magnitude is modest, with one additional year being associated with an increase in accident probability of 0.01 percentage points. Also of note is the fact that the explicit training variables are not significant. For the bottom part of the panel we see that the remainder of the covariates estimates are of the expected sign. Greater frequency of flight, measured by landings in the last 90 days is associated with an increased accident probability, as expected. Also as expected, flying in good weather, as measured by the VMC variable is associated with a remarkable reduction in accident probability (a 1.5 percentage point reduction). Filing a flight plan, one of our "prudence" variables is also associated with a 0.65 percentage point decrease in accident probability.

The results for commercial flights are not as strong. We see that hours tend to reduce accident probability. However, this effect is only statistically significant for commuters, and even so only in the NPE model. As with general aviation, the age, landings, and weather variables are of the expected sign. However, their magnitudes are smaller. Surprisingly, night flying only increases the accident probability of general aviation pilots, and even so only in the NPE identification scheme. Before moving to the other outcome equations, one should note that the estimates for trunk carriers in Table 3.11 include both accidents and incidents. Separate estimates done for just accidents are presented in Table 3.12. Although most of the observations are lost, we see a remarkable increase in the parameter estimates under the MF scheme. Now the effect of hours is very significant. Furthermore, the magnitude of the effects increased markedly. This effect is only present in the MF scheme.

The accident probit results are very significant, both statistically and with respect to the magnitude of the effect. It suggests that more experienced and better trained pilots are able to substantially decrease their chances of having an accident. This effect is substantially larger for general aviation pilots than for commercial pilots. We also see that the proportion of time as pilot in command is a significant training variable in determining accident rates in general aviation. This suggests that skills acquired while pilot in command are not only crucial to aircraft safety, but apparently not easily substitutable by other pilot training.

Taken as a whole, the results of the empirical models estimated point to potentially large effects of training on aviation safety outcomes. These results are not limited to the general aviation sample, but pervade every flight type. In all outcome models the results are stronger for general aviation, but we still see flight hour effects present for trunk carriers. This is particularly true in the pilot error and in the accident probits. There we see large experience effects in both flight categories. In the fatality/severity and destroyed/extent of damage equations the results are varied, but in almost all cases point to better safety outcomes associated with more experience pilots, pilot with more time in command, and, to a lesser extent, pilots who have more certificates. In the pilot error and in the fatality or severe injury probit we see largest experience effects, effects that clearly show that more experienced pilots are better able to avoid mistakes and better able to mitigate the effects of accidents. These effects are larger for general aviation, but again, are present for trunk carriers and are of large magnitudes. The composition of these hours was also found to be important, as pilots with more time as pilot in command and in the make and model of aircraft being flown were at lower risk for pilot error, and were—again—better able to mitigate the effects of an accident, or to avoid accidents altogether. This result was mostly concentrated in the general aviation models, although there was sporadic evidence of their effect on commuters also.

We also studied the effect of certificates on safety outcomes, and found general aviation to be a possible test case for these certificates, since we were able to compare safety outcomes of pilots with certificates with safety outcomes of pilots without them. We saw training effects associated with certificates in the expected direction, albeit not always statistically significant at conventional levels. Possession of an air transport certificate, in particular, is associated with lower pilot error probabilities, as well as lower fatality probabilities. The only training variable which did not behave in the way expected is the instrument-rating variable in the general aviation equation which, for reasons discussed above, may measure the relaxation of a flight restriction more than it does the effects of training.

Before concluding one should point out that when dealing with processes that vary over time, it is always possible that some of the dependent variables may be cointegrated with the explanatory variable. In particular, one must be wary of trending processes, that are integrated I(1). If both the dependent variable and the one or more of the explanatory variables are I(1), then the two series may be cointegrated, and one can attribute causality to the explanatory variable when in fact other factors may be affecting both variables. In a time-series setting one could test for unit roots using test such as those developed by Dickey and Fuller (1981) in order to rule out cointegration. Since our data consist of several cross-sections, and not a time-series per se, what can be done is to include a trend in the various outcome models and interact this trend with the explanatory variables. In other words, we estimate

$$(6) \qquad DV = f(\alpha TREND + \sum_i \beta_i EV_i + \sum_i \eta_i EV_i xTREND),$$

where DV is the dependent variable and EV are the explanatory variables.

If the net trend effect is significant and the interactions terms are also significant, then the two series may be cointegrated and we may be improperly attributing causality.

We perform this procedure on all outcomes. However, the net trend term is only significant for the pilot error equation. And for the pilot error equation, the only significant interaction is for the proportion of time as pilot in command. These results, which are presented in the appendix table 3.14, suggest that the safety outcome findings of this paper are not driven by cointegration between the safety outcomes and the explanatory variables.

V. Conclusion

The foregoing pages have been an attempt at modeling the effect of pilot training in accident safety outcomes. We modeled several different safety outcomes, and when significant results were found, they always highlighted the importance of experience and training in reducing pilot error, preventing accidents, and ultimately in preserving lives. Most of the results are concentrated in the general aviation category, a category which accounts for the vast majority of accidents and fatalities, but which has more often than not been ignored in the aviation safety literature. This is not to say that significant results were not found for commercial aviation, and in particular for trunk carriers, the aviation category that accounts for most aviation travel. In fact, the role of experience in reducing the incidence of pilot error is just as strong for trunk carriers as it is for general aviation. Furthermore, both the parametric and the non-parametric results point to significant non-linearities in these returns, with most of the gains coming from the first thousand hours of flight (general aviation), or the first three or four thousand hours of flight (commercial flights). This suggests that policy measures should focus on young and inexperienced pilots, for this is where the largest opportunity to save lives exists.

REFERENCES

Barnett, Arnold and Mary Higgins, "Airline Safety: The Last Decade." *Management Science* 29 (November 1983): 1225-36.

Borenstein, Severin and Martin Zimmerman, "Market incentives for safe commercial airline operation," *American Economic Review* 79 (December 1988): 913-35.

Bylow, Lance F. and Ian Savage, "The Effect of Airline Deregulation on Automobile Fatalities." *Accident Analysis and Prevention* 23 (May 1991): 443-52.

Chance, Don and Stephen P. Ferris, "The effect of aviation disasters on the air transport industry: a financial market perspective," *Journal of Transport Economics and Policy* 21 (May 1987): 151-65

Code of Federal Regulations, 1998, online version. National Archives and Records Administration's Office of the Federal Register.

Heckman, J., "Sample Selection Bias as Specification Error." *Econometrica* 47 (1979): 153-161.

Mitchell, Mark L. and Michael T. Maloney, "Crisis in the cockpit? The role of market forces in promoting air travel safety," *Journal of Law and Economics* 32 (October 1989): 329-55.

Oster, Clinton V. and C. K. Zorn, "Airline Deregulation: Is it Still Safe to Fly?"in Moses, Leon and Ian Savage, eds. *Transportation Safety in an Age of Deregulation.* Oxford University Press (1989).

Phillips, R. A. and W. K. Talley, "Determinants of the Severity of Airline Accidents: a Passanger Perspective." *International Journal of Transport Economics* 23 (June 1996): 141-156.

Phillips, R. A. and W. K. Talley, "Airline Safety Investments and Operating Conditions." *Southern Economic Journal* 59 (October 1992): 157-164.

Rose, Nancy L. "Profitability and Product qaulity: Economic Determinants of Airline Performance," *Journal of Political Economy* 98 (October 1990): 944-64.

Rose, Nancy L. "Fear of Flying? Economic Analysis of Airline Safety." *Journal of Economic Perspectives* 6 (Spring 1992): 75-94.

Zellner, A. " An Efficient Method of Estimating Seemingly Unrelated Regressions and Tests for Aggeregation Bias", *Journal of the American Statistical Association* 57: 348-368.

APPENDIX 1

Tables and figures for chapter 1

Table 1.1: Summary statistics, 1988 PNAD and 1987/88 POF

PNAD		POF	
Mother's schooling			
Less than elementary	0.210	Less than elementary	0.119
Elementary	0.263	Elementary	0.242
Some middle school	0.133	Middle	0.352
Middle	0.106	High school	0.156
Some high school	0.047	More than high	0.131
High school	0.148		
More than high school	0.093		
Father's schooling			
Less than elementary	0.218	Less than elementary	0.102
Elementary	0.270	Elementary	0.230
Some middle school	0.115	Middle	0.331
Middle	0.102	High school	0.146
Some high school	0.040	More than high	0.191
High school	0.137		
More than high school	0.118		
Number of children			
Boy, ages 9 and under	0.633	Boy, ages 0-5	0.324
Girl, ages 9 and under	0.637	Girl, ages 0-5	0.345
Boy, ages 10 to 11	0.108	Boy, ages 6-12	0.319
Girl, ages 10 to 11	0.108	Girl, ages 6-12	0.344
Boy, ages 12 to 13	0.096	Boy, ages 13-15	0.165
Girl, ages 12 to 13	0.094	Girl, ages 13-15	0.170
Boy, ages 14 to 15	0.089		
Girl, ages 14 to 15	0.084		
Household size	4.55		4.44
Mother's age	34.93		39.15
Father's age	38.62		42.79
Proportion of mothers employed	0.41		...
Proportion of fathers employed	0.93		...
Mother's hours	36.45		...
Father's hours	48.43		...

Monthly household consumption, in 1000s of 1987 \$Cz.'			3,402.38
			percent of total
(1) Food expenditures	...	722.28	(21.23%)
(2) Health care expenditures	...	210.83	(6.20%)
(3) Education expenditures	...	109.71	(3.22%)
(4) Tobacco expenditures	...	42.24	(1.24%)
(5) Adult clothing expenditures	...	240.38	(7.07%)
(6) Beer expenditures	...	7.55	(0.22%)
(7) Adult goods expenditures equals (4)+(5)+(6)	...	290.17	(8.53%)

Table 1.2: Men and women's earnings and employment rates, 1988

Age group	Average monthly earnings[1] Men	Average monthly earnings[1] Women	Employment rate Men	Employment rate Women
		All schooling levels		
16-24	476.27	300.12	0.72	0.42
25-34	1149.07	666.21	0.93	0.55
35-44	1699.25	847.82	0.95	0.58
45-54	1598.51	661.20	0.89	0.47
55-64	1140.86	624.86	0.74	0.25
65+	765.71	284.24	0.38	0.08
		Less than completed primary		
16-24	255.34	118.52	0.78	0.40
25-34	475.55	223.11	0.93	0.43
35-44	631.33	249.33	0.94	0.47
45-54	684.90	287.13	0.88	0.40
55-64	663.23	304.18	0.71	0.23
65+	777.05	282.99	0.35	0.07

[1] From 1988 PNAD. Values deflated to thousands of 1987 Cruzados.

Table 1.3: Outlay equivilancy ratios for selected goods[1]

Age Category	Tobacco	Adult clothing	Beer	Adult goods
Boys, 5 and under	-0.7010	-0.4660	-0.8220	-0.5040
	(0.087)	(0.054)	(0.207)	(0.058)
Boys, 6-12	-0.6750	-0.5420	-0.6140	-0.5600
	(0.065)	(0.044)	(0.165)	(0.055)
Boys, 15-15	-0.9140	-0.3890	-0.8140	-0.4630
	(0.092)	(0.064)	(0.207)	(0.087)
Girls, 5 and under	-0.7920	-0.4240	-0.6090	-0.4730
	(0.085)	(0.055)	(0.234)	(0.058)
Girls, 6-12	-0.6200	-0.4510	-0.2880	-0.4670
	(0.067)	(0.053)	(0.207)	(0.057)
Girls, 13-15	-0.8050	-0.3540	-1.1420	-0.4300
	(0.105)	(0.069)	(0.185)	(0.084)
Boy-girl, 5 and under	0.0910	-0.0420	-0.2120	-0.0300
	(0.092)	(0.059)	(0.295)	(0.078)
Boy-girl, 6-12	-0.0560	-0.0910	-0.3260	-0.0930
	(0.093)	(0.075)	(0.276)	(0.085)
Boy-girl, 13-15	-0.1100	-0.0350	0.3280	-0.0340
	(0.140)	(0.092)	(0.272)	(0.126)

Wald Tests [2]

	(1)		**(2)**	
Boys, 5 and under	6.080	*(0.048)*	2.976	*(0.226)*
Boys, 6-12	1.476	*(0.478)*	0.132	*(0.936)*
Boys, 13-15	9.917	*(0.007)*	1.912	*(0.384)*
Girls, 5 and under	10.035	*(0.007)*	0.795	*(0.672)*
Girls, 6-12	3.563	*(0.168)*	0.640	*(0.726)*
Girls, 13-15	12.170	*(0.002)*	6.962	*(0.031)*

[1] Results are weighted. Standard errors in parenthesis.

[2] (1) is the hypothesis that all adult goods *OERs* are equal, (2) is the hypothesis that Tobacco and Beer are equal. Statistics are distributed Chi-squared with three degrees of freedom for (1) and two degrees of freedom for (2).

Table 1.4: OLS regressions for the effect of child gender on expenditure shares, by expenditure type[1]

	Consumption	Health	Education	Adult good
Mother's education				
Completed primary	0.0900**	0.0012	0.0031	0.0032
	(2.50)	(0.34)	(1.58)	(1.00)
Completed middle	0.1377**	0.0025	0.0069**	0.0019
	(3.85)	(0.74)	(3.31)	(0.59)
Completed high school	0.2598**	0.0000	0.0137**	-0.0012
	(4.35)	(0.00)	(4.75)	(-0.28)
College +	0.4450**	0.0020	0.0217**	0.0008
	(6.86)	(0.32)	(5.11)	(0.14)
Mother's age/10	-0.0026	-0.0002	0.0011**	-0.0010
	(-0.40)	(-0.33)	(3.74)	(-1.53)
Mother's age/10 squared	0.0000	0.0000	0**	0.0000
	(0.47)	(0.51)	(-4.31)	(0.87)
Father's education				
Completed primary	0.1106**	-0.0028	0.0013	0.0022
	(2.86)	(-0.89)	(0.91)	(0.67)
Completed middle	0.1959**	0.0003	0.0076**	-0.0016
	(5.91)	(0.09)	(4.36)	(-0.52)
Completed high school	0.3406**	0.0000	0.0105**	-0.0055
	(6.06)	(0.00)	(4.95)	(-1.35)
College +	0.5605**	0.0027	0.017**	-0.0098**
	(6.79)	(0.50)	(5.38)	(-2.08)
Father's age/10	0.0342**	-0.0017**	0.0008**	-0.0027**
	(4.95)	(-2.28)	(2.58)	(-4.38)
Father's age/10 squared	-0.0003**	0**	0**	0**
	(-4.29)	(2.69)	(-2.40)	(3.33)
Male household	0.1337	0.0038	0.0035	-0.0216**
	(1.39)	(0.42)	(1.00)	(-2.08)
Log of household size	0.4712**	-0.0156**	0.0161**	0.0459**
	(7.67)	(-3.70)	(6.44)	(10.29)
Log of household expenditures	-0.5714**	...	...	...
	(-15.12)	...	...	...
Log of household consumption	...	0.0110**	0.0008	0.0012
	...	(5.41)	(0.64)	(0.69)

Table 1.4 (cont.)

	Consumption	Health	Education	Adult good
Number of children				
Boys 5 and under	-0.0891**	0.0053**	-0.0054**	-0.0102**
	(-3.82)	(2.77)	(-5.04)	(-4.64)
Girls 5 and under	-0.1078**	0.0071**	-0.0035**	-0.0117**
	(-5.30)	(3.06)	(-3.12)	(-5.50)
Boys, 6-9	-0.0818**	0.0015	0.0008	-0.0112**
	(-3.25)	(0.69)	(0.55)	(-4.47)
Girls, 6-9	-0.1026**	0.0003	0.0035**	-0.011**
	(-4.00)	(0.12)	(2.77)	(-4.03)
Boys, 10-11	-0.0720*	-0.0023	-0.0001	-0.0121**
	(-1.90)	(-0.87)	(-0.05)	(-3.65)
Girls, 10-11	-0.0350	-0.0010	0.0020	-0.007*
	(-0.94)	(-0.34)	(1.00)	(-1.89)
Boys, 12-13	-0.0706**	-0.0003	0.0021	-0.0117**
	(-2.04)	(-0.09)	(1.12)	(-3.53)
Girls, 12-13	-0.0458	0.0035	0.0027	-0.0105**
	(-0.99)	(1.15)	(1.36)	(-3.19)
Boys, 14-15	-0.0727**	-0.0023	0.0040	-0.0013
	(-2.22)	(-0.78)	(1.35)	(-0.35)
Girls, 14-15	-0.0964**	0.0014	-0.0008	-0.0006
	(-3.19)	(0.48)	(-0.48)	(-0.18)
Difference in boy-girl coefficients				
Ages 5 and under	0.0187	-0.0018	-0.0019	0.0014
	(0.64)	*(0.45)*	*(2.13)*	*(0.32)*
Ages 6 to 9	0.0208	0.0012	-0.0027	-0.0002
	(0.43)	*(0.19)*	*(2.20)*	*(0.00)*
Ages 10 to 11	-0.0370	-0.0013	-0.0021	-0.0051
	(0.54)	*(0.12)*	*(0.61)*	*(1.16)*
Ages 12 to 13	-0.0248	-0.0038	-0.0006	-0.0012
	(0.20)	*(0.85)*	*(0.05)*	*(0.70)*
Ages 14 to 15	0.0237	-0.0037	0.0049	-0.0007
	(0.41)	*(0.90)*	*(2.20)*	*(0.20)*
Joint significance of boy-girl coefficients *(F- statisitc)*	*(0.60)*	*(0.44)*	*(1.66)*	*(0.32)*

[1] Results are weighted. t-statistics in parenthesis. F-statistics reported for difference in boy-girl coefficients.* indicates 5% significance level; ** indicates 1%.

Table 1.5: OLS regressions for the effect of child gender on expenditure shares, mothers with low education[1]

	Consumption	Health	Education	Adult good
Number of children				
Boys 5 and under	-0.0987**	-0.0007	-0.0033**	-0.0067*
	(-2.60)	(-0.27)	(-2.26)	(-1.78)
Girls 5 and under	-0.0976**	-0.0028	-0.0051**	-0.0097**
	(-3.28)	(-1.00)	(-4.14)	(-2.69)
Boys, 6-9	-0.0442	0.0003	0.0010	-0.0111**
	(-1.27)	(0.12)	(0.37)	(-2.37)
Girls, 6-9	-0.139**	0.0066**	0.0025	-0.0159**
	(-3.43)	(1.98)	(1.46)	(-3.19)
Boys, 10-11	-0.0742	-0.0028	0.0015	-0.0124**
	(-0.96)	(-0.74)	(0.61)	(-2.26)
Girls, 10-11	-0.0534	-0.0019	0.0049	-0.0056
	(-1.00)	(-0.43)	(1.40)	(-0.87)
Boys, 12-13	-0.091*	-0.0045	0.0030	-0.0087
	(-1.84)	(-1.21)	(1.04)	(-1.58)
Girls, 12-13	-0.0875	0.0041	-0.0007	-0.0129**
	(-1.49)	(0.91)	(-0.28)	(-2.52)
Boys, 14-15	-0.0193	-0.0022	0.0053	-0.0002
	(-0.35)	(-0.47)	(0.98)	(-0.03)
Girls, 14-15	-0.0530	-0.0042	-0.0002	-0.0036
	(-1.21)	(-0.94)	(-0.09)	(-0.65)
Difference in boy-girl coefficients				
Ages 5 and under	-0.0011	0.0187	0.0187	0.0187
	(0.01)	*(0.41)*	*(0.98)*	*(0.39)*
Ages 6 to 9	0.0948*	-0.0062	-0.0015	0.0048
	(2.78)	*(2.49)*	*(0.21)*	*(0.46)*
Ages 10 to 11	-0.0207	-0.0009	-0.0034	-0.0068
	(0.04)	*(0.03)*	*(0.65)*	*(0.75)*
Ages 12 to 13	-0.0035	-0.0086	0.0037	0.0042
	(0.00)	*(2.15)*	*(0.98)*	*(0.30)*
Ages 14 to 15	0.0337	0.0021	0.0055	0.2100
	(0.32)	*(0.13)*	*(1.02)*	*(0.41)*
Joint significance of boy-girl coefficients *(F- statisitc)*	*(1.12)*	*(1.00)*	*(0.77)*	*(0.42)*

[1] Low schooling defined as less than primary completed. Results are weighte, t-statistics in parenthesis. F-statistics reported for difference in coefficients. * indicates 5% significance level; ** indicates 1%.

Table 1.6: Probit of mother's employment, 1988 PNAD and full sample of PNADs[1]

	Specification is *Model 1*			
	1988		Full Sample	
Mother's years of schooling				
One	0.0099	(0.71)	0.0219**	(6.08)
Two	-0.0267**	(-2.35)	0.0180**	(6.26)
Three	-0.0201*	(-1.92)	0.0306**	(11.58)
Four	0.0011	(0.13)	0.0465**	(21.20)
Five	0.0043	(0.92)	0.0583**	(16.71)
Six	0.0060	(0.36)	0.0534**	(14.06)
Seven	0.0136	(0.83)	0.0653**	(17.30)
Eight	0.0103	(0.89)	0.0702**	(24.70)
Nine	0.0452*	(1.93)	0.1011**	(18.78)
Ten	0.0742**	(3.66)	0.1227**	(25.38)
Eleven	0.1926**	(17.98)	0.2305**	(89.56)
Twelve	0.2451**	(6.01)	0.2991**	(33.01)
Thirteen	0.3100**	(8.83)	0.3063**	(39.17)
Fourteen	0.3078**	(11.35)	0.3756**	(61.23)
Fifteen	0.4005**	(23.75)	0.4225**	(111.13)
Sixteen	0.4319**	(14.30)	0.4585**	(76.43)
Seventeen or more	0.5196**	(5.36)	0.4961**	(26.71)
Mother's age/10	0.04149**	(20.88)	0.0443**	(96.07)
Mother's age/10 squared	-0.0006**	(-21.72)	-0.0006**	(-98.31)
Number of adults (i.e.16 years and older)	-0.0009	(-0.38)	-0.0055**	(-9.09)
Child status variables (number of children)				
Boys, nine and under	-0.0246**	(-7.65)	-0.0368**	(-46.99)
Girls, nine and under	-0.0210**	(-6.39)	-0.0351**	(-44.28)
Boys, ten or eleven years old	0.0074	(0.86)	-0.0009	(-0.44)
Girls, ten or eleven years old	0.0214**	(2.47)	0.0113**	(5.28)
Boys, twelve or thirteen years old	0.0067	(0.78)	0.0074**	(3.55)
Girls, twelve or thirteen years old	0.0360**	(4.13)	0.0262**	(12.41)
Boys, fourteen or fifteen years old	0.0033	(0.38)	-0.0029	(-1.38)
Girls, fourteen or fifteen years old	0.0390**	(4.36)	0.0281**	(13.06)
Child dummy (1 if family has a child)	0.0086	(0.87)	0.0153**	(6.57)

[1] Results are weighted. Marginal effects are reported, z-statistics in parenthesis.
* indicates 5% significance level; ** indicates 1%.

Table 1.7: Probit of mother's employment using full sample of PNADs, Models 1-4[1]

	Model 1		*Model 2*[2] (1) + prop. at work	
Number of adults	-0.0050**	(-9.09)	-0.0060**	(-9.81)
Child status variables (number of children)				
Boys, nine and under	-0.0368**	(-46.99)	-0.0394**	(-50.03)
Girls, nine and under	-0.0351**	(-44.28)	-0.0377**	(-47.32)
Boys, ten or eleven years old	-0.0009	(-0.44)	0.0002	(0.11)
Girls, ten or eleven years old	0.0113**	(5.28)	0.0221**	(10.37)
Boys, twelve or thirteen years old	0.0074**	(3.55)	-0.0047**	(-2.25)
Girls, twelve or thirteen years old	0.0262**	(12.41)	0.0293**	(13.89)
Boys, fourteen or fifteen years old	-0.0029	(-1.38)	-0.0364**	(-16.62)
Girls, fourteen or fifteen years old	0.0281**	(13.06)	0.0191**	(8.85)
Child dummy (1 if family has a child)	0.0153**	(6.57)	-0.0069**	(-2.92)
Proportion of children at work	...	...	0.2006**	(68.04)
Proportion of children doing housework	...	...	...	...
Difference in boy-girl coefficients[3]				
Ages 9 and under	-0.0017	*(2.45)*	-0.0017	*(2.47)*
Ages 10 to 11	-0.0122**	*(23.61)*	-0.0219**	*(76.11)*
Ages 12 to 13	-0.0188**	*(53.64)*	-0.0341**	*(175.58)*
Ages 14 to 15	-0.0310**	*(136.87)*	-0.0555**	*(426.59)*
Joint significance test of all difference in boy-girl coefficients (*Chi-squared*)	*(217.90)*		*(673.34)*	

Table 1.7 (cont.)

	Model 3 [2] (1) + prop. housework		*Model 4* (1) + both proportions	
Number of adults	-0.0044	(-7.22)	-0.0048	(-7.81)
Child status variables (number of children)				
Boys, nine and under	-0.0358**	(-45.64)	-0.0384**	(-48.68)
Girls, nine and under	-0.0341**	(-42.99)	-0.0367**	(-46.02)
Boys, ten or eleven years old	0.0075**	(3.54)	0.0094**	(4.42)
Girls, ten or eleven years old	0.0123**	(5.77)	0.0235**	(10.99)
Boys, twelve or thirteen years old	0.0141**	(6.70)	0.0022	(1.03)
Girls, twelve or thirteen years old	0.0238**	(11.24)	0.0268**	(12.65)
Boys, fourteen or fifteen years old	0.003	(1.53)	-0.0305**	(-13.87)
Girls, fourteen or fifteen years old	0.0234**	(10.88)	0.0139**	(6.42)
Child dummy (1 if family has a child)	-0.011**	(-4.66)	-0.0363**	(-14.58)
Proportion of children at work	...	...	0.2055**	(69.60)
Proportion of children doing housework	0.0775**	(33.07)	0.0841**	(35.83)
Difference in boy-girl coefficients[3]				
Ages 9 and under	-0.0016**	*(2.38)*	-0.0017**	*(2.40)*
Ages 10 to 11	-0.0048**	*(3.63)*	-0.0140**	*(31.20)*
Ages 12 to 13	-0.0096**	*(14.11)*	-0.0245**	*(90.34)*
Ages 14 to 15	-0.0202**	*(57.21)*	-0.0443**	*(269.14)*
Joint significance test of all difference in boy-girl coefficients (*Chi-squared*)	*(76.69)*		*(383.91)*	

[1] Regional dummies included but not reported. Results are weighted. Marginal effects are reported, z-statistics in parenthesis. Sample is married households with household heads 18 to 64 years old. Mother's schooling, household size and mother's age and age squared variables were included in the estimation but are also not reported. * indicates 5% significance level; ** indicates 1%.

[2] Proportion worked variable is defined as the proportion of children who reported working in the reference week, and is coded 0 for households without children. Alternate prop. worked variables used: (1) children who worked at least 10 hours per week and (2) proportion of reference person's sons and daughters who worked (as opposed to children in the household). Results were similar. Proportion working at home is defined as proportion who reported "helping with household chores" as main activity for the reference week.

[3] Chi-squared for the hypothesis that boy effects equal girl effects are reported in parenthesis and italics.

Table 1.8: Probit of mother's employment, high versus low child wage areas, difference-in-differences[1]

	Model 5	
Child status variables (number of children)		
Boys, nine and under	-0.0411**	(-31.39)
Girls, nine and under	-0.0391**	(-29.54)
Boys, ten or eleven years old	-0.0068**	(-2.10)
Girls, ten or eleven years old	0.0023	(0.71)
Boys, twelve or thirteen years old	0.0020	(0.62)
Girls, twelve or thirteen years old	0.0243**	(7.33)
Boys, fourteen or fifteen years old	-0.0030	(-0.89)
Girls, fourteen or fifteen years old	0.0279**	(8.28)
Child dummy (1 if family has a child)	0.0186**	(7.93)
Number of adults	-0.0051**	(-8.53)
High child wage dummy and interactions		
High child wage	-0.0204**	(-10.48)
High wage x boys, nine and under	0.0105**	(6.67)
High wage x girls, nine and under	0.01**	(6.32)
High wage x boys, ten or eleven years old	0.0101**	(2.68)
High wage x girls, ten or eleven years old	0.015**	(3.92)
High wage x boys, twelve or thirteen years old	0.0101**	(2.60)
High wage x girls, twelve or thirteen years old	0.0040	(1.02)
High wage x goys, fourteen or fifteen years old	0.0012	(0.31)
High wage x girls, fourteen or fifteen years old	0.0017	(0.44)
Difference in boy-girl marginal effects[2]		
Ages 9 and under	-0.0019	*(2.61)*
Ages 10 to 11	-0.0116**	*(19.72)*
Ages 12 to 13	-0.0193**	*(51.96)*
Ages 14 to 15	-0.0311**	*(128.06)*
High vs. low child wage area difference-in-differences[3]		
Ages 9 and under	0.0004	*(0.04)*
Ages 10 to 11	-0.0049	*(0.87)*
Ages 12 to 13	0.0061	*(1.29)*
Ages 14 to 15	-0.0005	*(1.01)*
Joint significance test of all difference-in-differences *(Chi-squared)*	*(2.21)*	

[1] See note 1 in Table 1.7.

[2] (dP/dboy - dP/dgirl). Chi-squared for difference in coefficients are reported in italics.

[3] High (low) wage cells are defined by the top (bottom) 50th centile of *year* x *region* x *urban* categories.

Table 1.9: Probit of mother's employment, households with and without childcare availability, difference-in-differences[1]

	Model 6	
Child status variables (number of children)		
Boys, nine and under	-0.0462**	(-52.34)
Girls, nine and under	-0.0444**	(-49.71)
Boys, ten or eleven years old	-0.0018	(-0.74)
Girls, ten or eleven years old	0.0124**	(5.01)
Boys, twelve or thirteen years old	0.0072**	(2.90)
Girls, twelve or thirteen years old	0.0309**	(12.23)
Boys, fourteen or fifteen years old	-0.0054**	(-2.03)
Girls, fourteen or fifteen years old	0.0361**	(13.44)
Child dummy (1 if family has a child)	0.0156**	(6.60)
Number of adults	-0.0168**	(-21.49)
Childcare availability dummy and interactions		
Childcare available	0.0156**	(6.13)
Childcare x boys, nine and under	0.0388**	(21.61)
Childcare x girls, nine and under	0.0388**	(21.31)
Childcare x boys, ten or eleven years old	-0.0036	(-0.88)
Childcare x girls, ten or eleven years old	-0.0109**	(-2.65)
Childcare x boys, twelve or thirteen years old	-0.0091**	(-2.29)
Childcare x girls, twelve or thirteen years old	-0.023**	(-5.75)
Childcare x goys, fourteen or fifteen years old	-0.0036	(-0.92)
Childcare x girls, fourteen or fifteen years old	-0.0286**	(-7.20)
Difference in boy-girl coefficients[2]		
Ages 9 and under	-0.0018	*(2.37)*
Ages 10 to 11	-0.0116**	*(20.95)*
Ages 12 to 13	-0.0187**	*(53.20)*
Ages 14 to 15	-0.0326**	*(148.20)*
Childcare availability difference-in-differences[3]		
Ages 9 and under	0.0000	*(0.00)*
Ages 10 to 11	0.0073	*(1.75)*
Ages 12 to 13	0.0139	*(6.67)*
Ages 14 to 15	0.0250	*(21.59)*
Joint significance test of all difference-in-differences *(Chi-squared)*	*(30.22)*	

[1-2] See note 1 in Table 1.8.

[3] Childcare availability is defined as the presence of a female relative, 16 years or older.

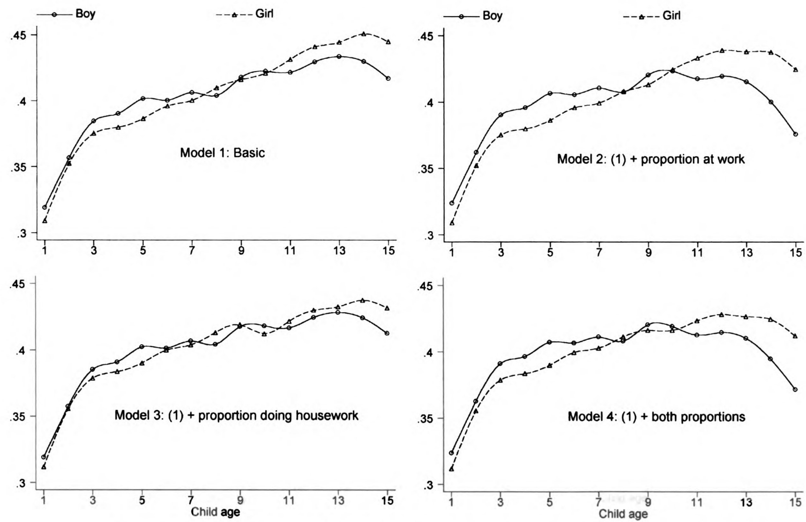

Figure 1.1: Wife's predicted employment for different child age and gender categories, Models 1 - 4

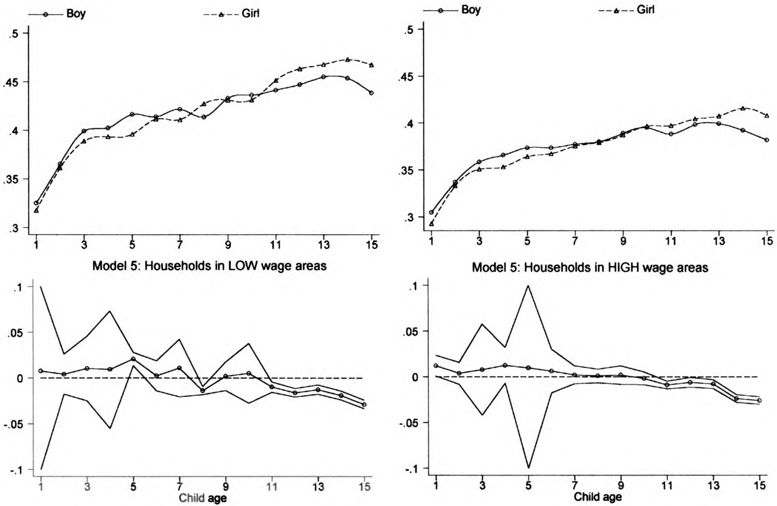

Figure 1.2: Wife's predicted employment for households in high and low wage regions (top) and 95 percent confidence interval for boy-girl difference in predicted employment (bottom)

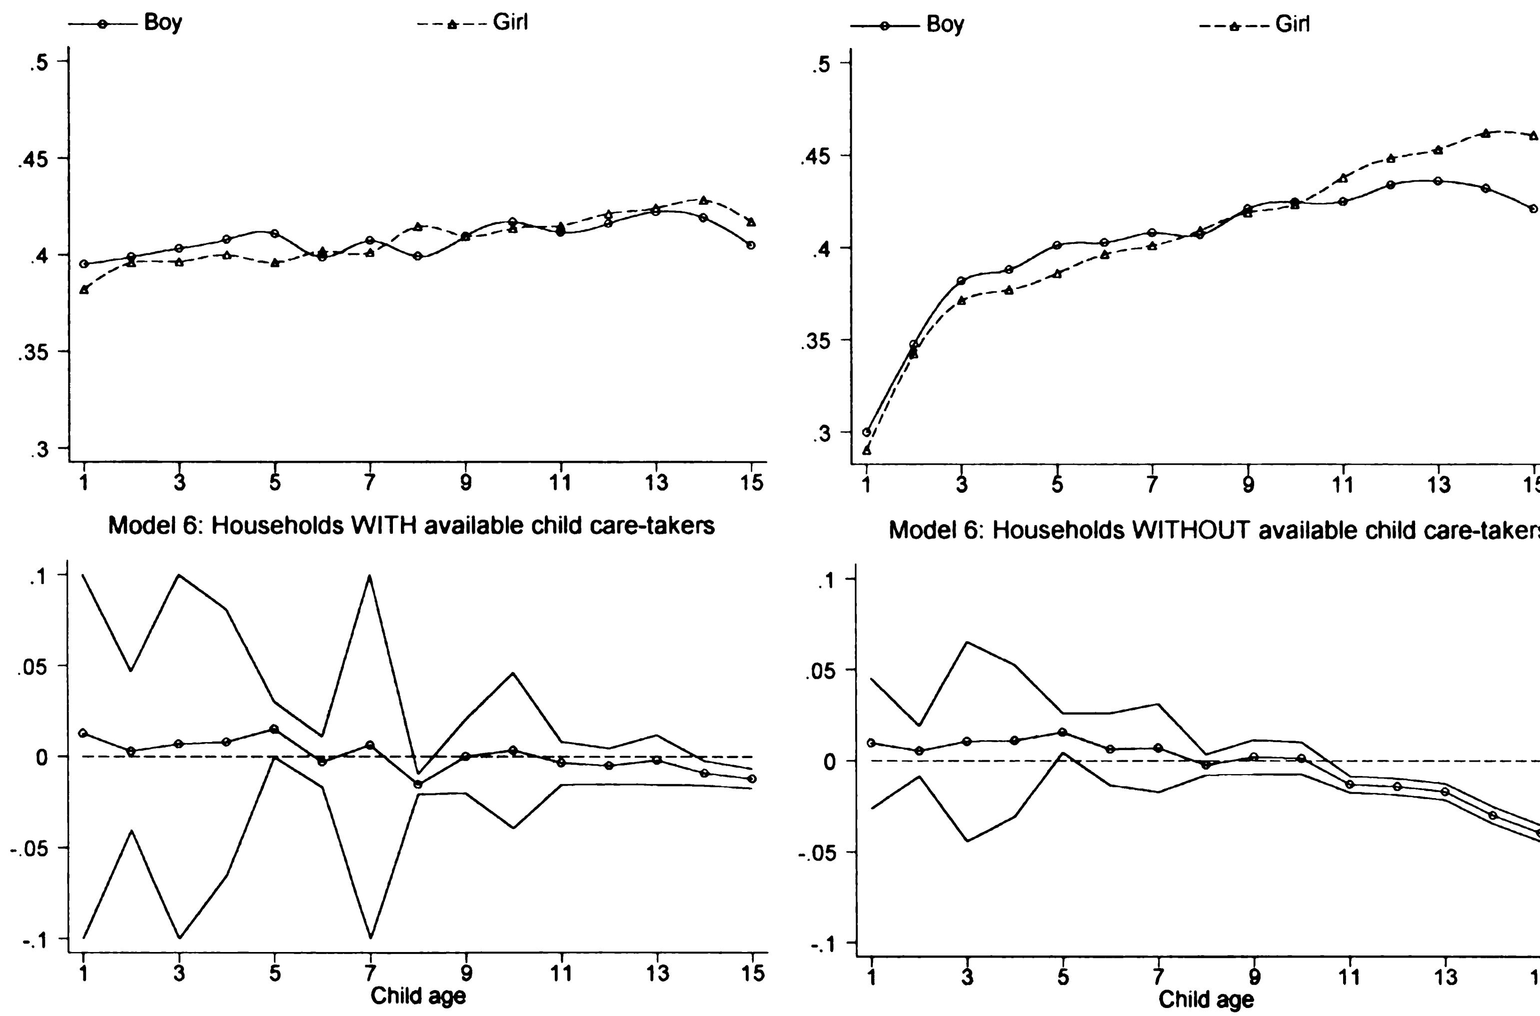

Figure 1.3: Wife's predicted employment probability for households with and without availability of child care-takers (top) and 95 percent confidence interval for boy-girl difference in predicted employment (bottom)

APPENDIX 2

Tables and figures for chapter 2

Table 2.1: Summary statistics, women ages 15 through 70, three years[1]

	Married women					
	1981		1990		1999	
	mean	obs	mean	obs	mean	obs
Age	37.706	77,403	38.393	53,478	39.277	58,738
Years education	3.835	77,403	4.912	53,478	6.136	58,738
Proportion rural	0.276	77,403	0.257	53,478	0.194	58,738
LFP	0.233	77,403	0.334	53,478	0.431	58,738
Proportion at work	0.227	77,403	0.328	53,478	0.380	58,738
Proportion who are formal wage workers	0.456	18,894	0.457	18,247	0.451	22,900
Proportion who are informal wage workers	0.184	18,894	0.181	18,247	0.259	22,900
Proportion who are self-employed	0.360	18,894	0.361	18,247	0.289	22,900
Household size	5.035	77,403	4.566	53,478	4.145	58,738
No daughters 0-4	0.338	77,403	0.243	53,478	0.177	58,738
No daughters 5-9	0.313	77,403	0.273	53,478	0.193	58,738
No. sons 0-4	0.344	77,403	0.251	53,478	0.186	58,738
No. sons 5-9	0.314	77,403	0.280	53,478	0.201	58,738

Table 2.1 (cont.)

	Single women					
	1981		1990		1999	
	mean	obs	mean	obs	mean	obs
Age	29.634	70,506	30.921	45,598	31.583	56,153
Years education	5.019	70,506	5.818	45,598	7.139	56,153
Proportion rural	0.212	70,506	0.189	45,598	0.139	56,153
LFP	0.459	70,506	0.516	45,598	0.543	56,153
Proportion at work	0.428	70,506	0.489	45,598	0.444	56,153
Proportion of worker who are formal wage workers	0.525	30,609	0.517	22,295	0.487	25,072
Proportion of workers who are informal wage workers	0.313	30,609	0.313	22,295	0.350	25,072
Proportion of workers who are self-employed	0.162	30,609	0.170	22,295	0.162	25,072
Household size	5.877	70,506	5.226	45,598	4.684	56,153

1 Results are weighted. Statistics computed from the 1981, 1990, and 1999 PNADs.

Table 2.2: Work probit, married women[1]

	Basic model (**M1**)		+ education controls (**M2**)		+ husband's controls (**M3**)		+ fertility/demographic controls (**M4**)	
Region (Central-west omitted)								
Minas Gerais	-0.0021	(-1.02)	0.0036*	(1.74)	0.0050**	(2.33)	0.0085**	(3.90)
North	0.0213**	(8.31)	0.0352**	(13.58)	0.0340**	(12.90)	0.0418**	(15.70)
Northeast	0.0378**	(21.25)	0.0507**	(28.08)	0.0486**	(24.01)	0.0571**	(28.00)
Rio	0.0315**	(14.07)	0.0205**	(8.95)	0.0290**	(11.95)	0.0272**	(11.19)
South	0.0176**	(9.20)	0.0195**	(10.04)	0.0214**	(10.66)	0.0205**	(10.22)
Sao Paulo	0.0016	(0.83)	0.0016	(0.82)	0.0047**	(2.28)	0.0057**	(2.79)
Rural dummy	-0.1660**	(-116.48)	-0.1121**	(-74.87)	-0.1181**	(-71.65)	-0.1117**	(-67.62)
Age	0.0443**	(7.78)	0.0025	(0.43)	-0.0016	(-0.28)	0.0865**	(14.60)
Age squared	-0.0002	(-0.94)	0.0012**	(5.29)	0.0014**	(6.29)	-0.0019**	(-8.40)
Age cubed/10000	-0.0077**	(-2.08)	-0.0280**	(-7.54)	-0.0313**	(-8.42)	0.0193**	(5.08)
Age quartic/10000	0.0001**	(2.40)	0.0002**	(7.26)	0.0002**	(8.05)	-0.0001**	(-3.98)

[1] Marginal effects reported. Z-statistics in parenthesis

Table 2.2 (cont.)

	Basic model (M1)		+ education controls (M2)		+ husband's controls (M3)		+ fertility/demographic controls (M4)	
Years of completed education								
one	..	..	0.0024	(0.77)	0.0055*	(1.74)	0.0032	(1.01)
two	..	..	-0.0001	(-0.06)	0.0050**	(1.97)	0.0008	(0.33)
three	..	..	0.0081**	(3.53)	0.0157**	(6.82)	0.0103**	(4.45)
four	..	..	0.0239**	(12.87)	0.0389**	(20.07)	0.0309**	(15.87)
five	..	..	0.0344**	(10.76)	0.0560**	(17.00)	0.0481**	(14.57)
six	..	..	0.0380**	(10.92)	0.0644**	(17.94)	0.0556**	(15.45)
seven	..	..	0.0522**	(15.17)	0.0828**	(23.13)	0.0742**	(20.71)
eight	..	..	0.0598**	(23.95)	0.1010**	(36.60)	0.0919**	(33.19)
nine	..	..	0.0878**	(17.60)	0.1327**	(25.71)	0.1225**	(23.75)
ten	..	..	0.1060**	(24.07)	0.1573**	(34.09)	0.1482**	(32.14)
eleven	..	..	0.2038**	(92.03)	0.2718**	(100.58)	0.2635**	(97.05)
twelve	..	..	0.2783**	(33.49)	0.3707**	(43.22)	0.3601**	(42.07)
thirteen	..	..	0.2827**	(39.50)	0.3808**	(51.30)	0.3712**	(50.11)
fourteen	..	..	0.3626**	(63.87)	0.4612**	(76.54)	0.4554**	(75.55)
fifteen	..	..	0.4078**	(115.95)	0.5271**	(125.11)	0.5235**	(123.91)
sixteen	..	..	0.4575**	(78.43)	0.6054**	(92.61)	0.6004**	(91.96)
seventeen +	..	..	0.5184**	(27.94)	0.6864**	(35.54)	0.6834**	(35.41)
Husband's variables[2]								
Wage ($R/hr, predicted)	..	..	.. ..	..	0.0077**	(4.73)	0.0044**	(2.67)
Wage squared/100	..	..	.. ..	..	-0.0006**	(-10.33)	-0.0005**	(-7.66)
Schooling (completed years)	..	..	.. ..	..	-0.0096**	(-15.06)	-0.0094**	(-14.65)
Age	..	..	.. ..	..	-0.0029**	(-25.77)	-0.0033**	(-29.35)
Household size	..	..	..	..	..	..	0.0047**	(12.97)

[2] Husband's wage is predicted in a first stage (see Table 2.13), based on husband schooling and age--standard errors adjusted.
Sons and daughters are defined relative to the household head, and may miss children from other unions or children no longer at home.

Table 2.2 (cont.)

	Basic model **(M1)**		+ education controls **(M2)**		+ husband's controls **(M3)**		+ fertility/demographic controls **(M4)**	
Child status								
Daughters, 0-4, 1	..	..	..	..	..	..	-0.0793**	(-49.94)
Daughters, 0-4, 2	..	..	..	..	..	..	-0.1330**	(-39.10)
Daughters, 0-4, 3+	..	..	..	..	..	..	-0.1704**	(-16.96)
Sons, 0-4, 1	..	..	..	..	..	..	-0.0774**	(-48.97)
Sons, 0-4, 2	..	..	..	..	..	..	-0.1365**	(-40.75)
Sons, 0-4, 3	..	..	..	..	..	..	-0.1764**	(-18.45)
Daughters, 5-9, 1	..	..	..	..	..	..	-0.0185**	(-12.15)
Daughters, 5-9, 2	..	..	..	..	..	..	-0.0314**	(-10.27)
Daughters, 5-9, 3+	..	..	..	..	..	..	-0.0210**	(-2.52)
Sons, 5-9, 1	..	..	..	..	..	..	-0.0244**	(-16.15)
Sons, 5-9, 2	..	..	..	..	..	..	-0.0349**	(-11.57)
Sons, 5-9, 3	..	..	..	..	..	..	-0.0369**	(-4.51)
GDP growth	0.0006**	(4.05)	0.0007**	(4.65)	0.0007**	(4.63)	0.0007**	(4.64)
Selected birthyear cohort dummies (effects relative to 1921 cohort)								
Born 1925	0.0045	(0.38)	-0.0008	(-0.07)	0.0019	(0.16)	0.0027	(0.23)
Born 1930	0.0179*	(1.71)	0.0065	(0.62)	0.0092	(0.88)	0.0119	(1.14)
Born 1935	0.0315**	(3.08)	0.0101	(0.97)	0.0134	(1.30)	0.0175*	(1.70)
Born 1940	0.0665**	(6.64)	0.0352**	(3.48)	0.0396**	(3.93)	0.0457**	(4.55)
Born 1945	0.1061**	(10.58)	0.0605**	(5.98)	0.0668**	(6.62)	0.0711**	(7.05)
Born 1950	0.1498**	(15.04)	0.0893**	(8.88)	0.0949**	(9.45)	0.0952**	(9.50)
Born 1955	0.1953**	(19.57)	0.1193**	(11.84)	0.1252**	(12.44)	0.1170**	(11.65)
Born 1960	0.2301**	(22.88)	0.1438**	(14.15)	0.1497**	(14.73)	0.1328**	(13.09)
Born 1965	0.2703**	(26.32)	0.1758**	(16.94)	0.1821**	(17.55)	0.1585**	(15.28)
Born 1970	0.3047**	(28.44)	0.2096**	(19.35)	0.2171**	(20.04)	0.1889**	(17.45)

Table 2.3: Work probit, single women[1]

	Basic model (S1)		+ education controls (S2)		+ household size (S3)	
Region (Central-west omitted)						
Minas Gerais	-0.0199**	(-7.41)	-0.0191**	(-7.08)	-0.0143**	(-5.27)
North	-0.0749**	(-23.94)	-0.0709**	(-22.52)	-0.0584**	(-18.38)
Northeast	-0.0680**	(-29.07)	-0.0594**	(-25.23)	-0.0535**	(-22.62)
Rio	-0.0316**	(-10.85)	-0.0466**	(-15.72)	-0.0447**	(-14.99)
South	-0.0087**	(-3.20)	-0.0145**	(-5.34)	-0.0179**	(-6.55)
Sao Paulo	0.0422**	(15.81)	0.0348**	(12.91)	0.0355**	(13.10)
Rural dummy	-0.1519**	(-74.08)	-0.1195**	(-56.10)	-0.1162**	(-54.41)
Age	0.1810**	(42.99)	0.1424**	(33.37)	0.1651**	(38.26)
Age squared	-0.0052**	(-29.09)	-0.0040**	(-21.73)	-0.0049**	(-26.71)
Age cubed/10000	0.0648**	(20.29)	0.0464**	(14.39)	0.0623**	(19.10)
Age quartic/10000	-0.0003**	(-17.11)	-0.0002**	(-11.40)	-0.0003**	(-16.00)
Years of completed education						
one	..	..	0.0364**	(8.14)	0.0364**	(8.14)
two	..	..	0.0441**	(12.31)	0.0470**	(13.08)
three	..	..	0.0494**	(15.37)	0.0531**	(16.49)
four	..	..	0.0600**	(22.90)	0.0649**	(24.70)
five	..	..	0.0591**	(15.95)	0.0657**	(17.64)
six	..	..	0.0408**	(10.38)	0.0467**	(11.83)
seven	..	..	0.0371**	(9.65)	0.0422**	(10.92)
eight	..	..	0.0744**	(23.03)	0.0792**	(24.34)
nine	..	..	0.0285**	(6.43)	0.0341**	(7.64)
ten	..	..	0.0516**	(11.91)	0.0560**	(12.83)
eleven	..	..	0.2004**	(66.60)	0.2038**	(66.97)
twelve	..	..	0.1238**	(17.44)	0.1224**	(16.91)
thirteen	..	..	0.1472**	(20.47)	0.1446**	(19.98)
fourteen	..	..	0.2155**	(31.40)	0.2128**	(30.88)
fifteen	..	..	0.3391**	(67.79)	0.3349**	(66.57)
sixteen	..	..	0.3447**	(42.50)	0.3382**	(41.49)
seventeen +	..	..	0.4645**	(19.63)	0.4528**	(19.12)
Household size	..	..	..	..	-0.0164**	(-53.60)

Table 2.3 (cont.)

Selected birthyear cohort dummies (effects relative to 1921 cohort)

Born 1925	-0.0118	(-1.11)	0.0063	(0.62)	0.0067	(0.63)
Born 1930	-0.0174*	(-1.90)	-0.0044	(-0.52)	-0.0026	(-0.30)
Born 1935	-0.0010	(-0.10)	0.0047	(0.56)	0.0065	(0.74)
Born 1940	0.0298**	(3.24)	0.0284**	(3.43)	0.0293**	(3.36)
Born 1945	0.0480**	(5.00)	0.0290**	(3.36)	0.0293**	(3.21)
Born 1950	0.0482**	(5.08)	0.0206**	(2.44)	0.0177**	(1.97)
Born 1955	0.0493**	(5.26)	0.0069	(0.84)	0.0030	(0.34)
Born 1960	0.0299**	(3.35)	-0.0250**	(-3.26)	-0.0291**	(-3.49)
Born 1965	0.0280**	(3.20)	-0.0321**	(-4.33)	-0.0382**	(-4.71)
Born 1970	0.0322**	(3.57)	-0.0325**	(-4.22)	-0.0404**	(-4.81)

[1] Marginal effects reported. Z-statistics in parenthesis. Change in GDP also included but not reported.

Table 2.4: Age, period and cohort effects for model M4 probit[1]

	No period effects		No cohort effects		Both period and cohort effects	
Age effects						
Age	0.0865**	(14.60)	0.0756**	(13.18)	0.0872**	(14.40)
Age squared	-0.0019**	(-8.40)	0.0016**	(-7.35)	-0.0020**	(-8.48)
Age cubed/10000	0.0193**	(5.08)	0.0134**	(3.67)	0.0196**	(5.11)
Age quartic/10000	-0.0001**	(-3.98)	-0.0001**	(-2.30)	-0.0001**	(-3.98)
Selected birthyear cohort dummies (effects relative to 1921 cohort)						
Born 1925	0.0027	(0.23)	...	...	-0.0041	(-0.32)
Born 1930	0.0119	(1.14)	...	...	-0.0005	(-0.03)
Born 1935	0.0175*	(1.70)	...	...	-0.0002	(-0.01)
Born 1940	0.0457**	(4.55)	...	...	0.0231	(1.02)
Born 1945	0.0711**	(7.05)	...	...	0.0434	(1.59)
Born 1950	0.0952**	(9.50)	...	...	0.0631*	(1.96)
Born 1955	0.1170**	(11.65)	...	...	0.0807**	(2.17)
Born 1960	0.1328**	(13.09)	...	...	0.0924**	(2.19)
Born 1965	0.1585**	(15.28)	...	...	0.1143**	(2.42)
Born 1970	0.1889**	(17.45)	...	...	0.1434**	(2.74)
Period dummies (effects relative to 1996-1999)						
Year 1981	...	...	-0.0964**	(39.82)	-0.0368**	(-2.13)
Year 1982	...	...	-0.0712**	(30.85)	-0.0150	(-0.92)
Year 1983	...	...	-0.0579**	(25.32)	-0.0051	(-0.33)
Year 1984	...	...	-0.0597**	(26.51)	-0.0103	(-0.73)
Year 1985	...	...	-0.0489**	(22.16)	-0.0029	(-0.22)
Year 1986	...	...	0.0357**	(12.94)	0.0067	(0.55)
Year 1987	...	...	-0.0239**	(-8.83)	0.0150	(1.34)
Year 1988	...	...	-0.0239**	(-8.82)	0.0116	(1.14)
Year 1989	...	...	-0.0189**	(-7.01)	0.0129	(1.40)
Year 1990	...	...	-0.0133**	(-5.03)	0.0149*	(1.81)
Year 1992	...	...	-0.0158**	(-6.03)	0.0049	(0.77)
Year 1993	...	...	-0.0117**	(-4.53)	0.0052	(0.98)
Year 1995	...	...	0.0065**	(2.57)	0.0159**	(4.37)

[1] Marginal effects reported. Z-statistics in parenthesis

Table 2.5: Change in employment rates of married women, based on probit models M1-M4[1]

	Secular change: 1990-1981			
Actual change	0.094	0.094	0.094	0.094
Predicted change in employment	0.049	0.050	0.050	0.053
Percentage explained by				
Change in age (and region)	2.0	-4.1	5.3	-1.4
Change in education	...	34.9	46.9	43.5
Change in husband's wage	...	...	-21.2	-22.6
Change in household composition	...	...	...	25.6
Change in cohort composition	98.0	69.2	69.0	54.8
	Secular change: 1999-1990			
Actual change	0.066	0.066	0.066	0.066
Predicted change in employment	0.090	0.094	0.093	0.092
Percentage explained by				
Change in age (and region)	48.9	37.8	48.3	36.3
Change in education	...	26.0	35.4	35.0
Change in husband's wage	...	...	-20.6	-23.5
Change in household composition	...	...	...	20.2
Change in cohort composition	51.1	36.2	36.9	32.2
	Cohort change: 1952-1942			
Actual change	0.058	0.058	0.058	0.058
Predicted change in employment	0.047	0.055	0.056	0.056
Percentage explained by				
Change in age (and region)	88.7	39.0	-6.8	51.0
Change in education	...	53.7	73.4	68.6
Change in husband's wage	...	...	26.5	32.5
Change in household composition	...	...	...	-57.9
Change in cohort composition	11.3	7.3	6.9	5.8
	Cohort change: 1962-1952[1]			
Actual change	-0.034	-0.034	-0.034	-0.034
Predicted change in employment	-0.035	-0.038	-0.036	-0.034
Percentage explained by				
Change in age (and region)	157.0	147.0	201.2	162.0
Change in education	...	-9.6	-17.9	-17.1
Change in husband's wage	...	...	-45.0	-61.0
Change in household composition	...	...	...	51.5
Change in year composition	-57.0	-37.4	-38.3	-35.4

[1] See text for description of methodology. Since there was a fall in employment, a positive value correspon to a percentage decrease in employment, a negative value correspons to an increase.

Table 2.6: Sector choice model, multinomial logit[1]

	Wage workers				Self-employed	
	Formal		Informal			
Region (Central-west omitted)						
Minas Gerais	1.0037	(0.26)	0.9991	(-0.05)	1.1165**	(7.08)
North	1.2808**	(14.91)	0.8930**	(-4.61)	1.3720**	(17.06)
Northeast	1.1985**	(13.73)	0.9034**	(-5.48)	1.7259**	(38.42)
Rio	1.0099	(0.63)	1.2438**	(10.42)	1.2235**	(11.86)
South	1.2459**	(17.11)	1.0345*	(1.94)	1.0008	(0.05)
Sao Paulo	1.0449**	(3.31)	1.1734**	(9.24)	0.8963**	(-7.24)
Rural dummy	0.5044**	(-52.04)	0.6337**	(-32.00)	0.6049**	(-44.31)
Age	2.2685**	(16.34)	1.3643**	(6.04)	1.6147**	(10.24)
Age squared	0.9744**	(-12.88)	0.9943**	(-2.80)	0.9884**	(-6.57)
Age cubed/10000	1.4641**	(11.03)	1.0563	(1.59)	1.1412**	(4.58)
Age quartic/10000	0.9977**	(-10.86)	0.9997	(-1.54)	0.9993**	(-4.13)
Years of completed education						
one	1.3938**	(11.23)	0.8410**	(-7.08)	1.0464**	(2.15)
two	1.5267**	(18.07)	0.8141**	(-10.61)	0.9883	(-0.69)
three	1.7520**	(26.55)	0.7768**	(-14.03)	1.0478**	(2.97)
four	2.2921**	(45.33)	0.7019**	(-22.52)	1.1559**	(11.22)
five	2.6696**	(38.13)	0.7536**	(-11.14)	1.1696**	(6.71)
six	2.7799**	(37.23)	0.7162**	(-11.53)	1.2913**	(10.33)
seven	3.2993**	(42.86)	0.6992**	(-11.19)	1.3818**	(13.26)
eight	4.0405**	(63.44)	0.6258**	(-19.10)	1.3856**	(17.58)
nine	5.0241**	(48.18)	0.6849**	(-7.99)	1.5202**	(11.67)
ten	6.0047**	(59.73)	0.6566**	(-9.33)	1.6323**	(15.38)
eleven	13.3965**	(122.74)	0.8784**	(-5.20)	1.6533**	(26.69)
twelve	21.8596**	(67.33)	1.3844**	(3.74)	1.9899**	(10.18)
thirteen	23.3251**	(78.03)	1.4247**	(4.59)	2.0990**	(13.12)
fourteen	38.1036**	(106.21)	1.6820**	(8.14)	2.2011**	(16.13)
fifteen	54.8751**	(147.47)	2.2276**	(18.79)	3.2666**	(39.32)
sixteen	69.2131**	(111.33)	4.0047**	(22.80)	7.5660**	(49.15)
seventeen +	142.8921**	(46.96)	4.2344**	(7.47)	6.5122**	(14.17)
Husband's variables						
Wage ($R/hr, predicted)	1.0601**	(5.75)	0.9604**	(-2.45)	1.0380**	(3.26)
Wage squared	0.9959**	(-11.54)	1.0004	(0.66)	0.9984**	(-3.80)
Schooling (completed ye	0.9394**	(-14.90)	0.9346**	(-11.05)	0.9767**	(-5.34)
Age	0.9736**	(-32.65)	0.9818**	(-16.29)	0.9942**	(-7.87)
Household size	1.0265**	(9.48)	1.0326**	(10.21)	1.0119**	(4.99)

Table 2.6 (cont.)

	Wage workers				Self-employed	
	Formal		Informal			
Child status						
Daughters, 0-4, 1	0.6263**	(-44.21)	0.6754**	(-29.09)	0.7567**	(-24.38)
Daughters, 0-4, 2	0.4140**	(-34.99)	0.5129**	(-22.71)	0.6655**	(-16.70)
Daughters, 0-4, 3+	0.2756**	(-13.85)	0.4070**	(-10.32)	0.6220**	(-6.86)
Sons, 0-4, 1	0.6276**	(-44.23)	0.6681**	(-29.95)	0.7823**	(-21.67)
Sons, 0-4, 2	0.4049**	(-37.12)	0.5235**	(-22.43)	0.6404**	(-18.20)
Sons, 0-4, 3	0.3046**	(-14.52)	0.3841**	(-11.21)	0.5810**	(-7.83)
Daughters, 5-9, 1	0.7893**	(-23.28)	0.9914	(-0.70)	1.0081	(0.78)
Daughters, 5-9, 2	0.665**	(-18.32)	0.9698	(-1.27)	0.9919	(-0.40)
Daughters, 5-9, 3+	0.6545**	(-6.03)	1.0754	(1.16)	1.0281	(0.50)
Sons, 5-9, 1	0.7756**	(-25.14)	0.9456**	(-4.51)	0.9848	(-1.48)
Sons, 5-9, 2	0.6504**	(-19.55)	0.9716	(-1.21)	0.9705	(-1.48)
Sons, 5-9, 3	0.6173**	(-7.16)	0.9139	(-1.38)	1.0026	(0.05)
GDP growth	1.0033**	(3.05)	1.0083**	(5.87)	1.0012	(1.10)
Selected birthyear cohort dummies (effects relative to 1921 cohort)						
Born 1925	1.3077*	(1.81)	1.0896	(0.65)	0.9808	(-0.25)
Born 1930	1.4177**	(2.63)	1.1728	(1.36)	1.0347	(0.50)
Born 1935	1.3577**	(2.35)	1.3956**	(2.91)	1.0426	(0.62)
Born 1940	1.6475**	(3.89)	1.9438**	(6.00)	1.0530	(0.78)
Born 1945	1.7977**	(4.59)	2.3891**	(7.84)	1.1579**	(2.22)
Born 1950	1.8285**	(4.73)	3.2666**	(10.73)	1.2624**	(3.54)
Born 1955	1.8502**	(4.82)	4.6697**	(13.97)	1.3024**	(3.99)
Born 1960	1.7670**	(4.45)	6.7019**	(17.20)	1.3288**	(4.23)
Born 1965	1.7923**	(4.53)	9.8187**	(20.49)	1.4111**	(4.97)
Born 1970	1.9561**	(5.14)	13.8253**	(23.14)	1.4446**	(4.94)

[1] Relative risk ratios reported, z-statistics in parenthesis, not at work base category.

Table 2.7: Sector choice model, multinomial logit, both year and cohort effects[1]

	Wage workers				Self-employed	
	Formal		Informal			
Selected birthyear cohort dummies (effects relative to 1921 cohort)						
Born 1925	1.2474	(1.46)	0.9993	(0.00)	0.9946	(-0.03)
Born 1930	1.298*	(1.76)	0.9820	(-0.13)	1.0758	(1.20)
Born 1935	1.2025	(1.13)	1.0666	(0.40)	1.1139	(1.10)
Born 1940	1.419*	(1.89)	1.3549	(1.60)	1.1609	(1.60)
Born 1945	1.507*	(1.95)	1.5169*	(1.87)	1.3165	(1.20)
Born 1950	1.4977*	(1.70)	1.8887**	(2.47)	1.4856*	(1.74)
Born 1955	1.4833	(1.48)	2.455**	(3.06)	1.5923*	(1.88)
Born 1960	1.3858	(1.10)	3.203**	(3.53)	1.69*	(1.69)
Born 1965	1.3795	(0.98)	4.2628**	(3.94)	1.8767*	(2.20)
Born 1970	1.5029	(1.13)	5.4634**	(4.19)	2.0394**	(2.89)
Year effects (effects relative to 1996-1999 period)						
Year 1981	0.8552	(-1.38)	0.659**	(-3.16)	0.9778	(-1.30)
Year 1982	0.9226	(-0.76)	0.7193**	(-2.66)	1.1345	(0.70)
Year 1983	0.9159	(-0.88)	0.9141	(-0.78)	1.1099	(0.90)
Year 1984	0.9287	(-0.80)	0.816*	(-1.88)	1.1023	(0.99)
Year 1985	0.9840	(-0.19)	0.8251*	(-1.92)	1.1309	(0.98)
Year 1986	1.0412	(0.51)	0.8105**	(-2.24)	1.2193**	(1.99)
Year 1987	1.0944	(1.23)	0.8443**	(-1.97)	1.245**	(2.30)
Year 1988	1.0805	(1.16)	0.7878**	(-3.04)	1.2598**	(2.10)
Year 1989	1.0911	(1.45)	0.8576**	(-2.17)	1.2044**	(2.55)
Year 1990	1.0868	(1.55)	0.781**	(-3.90)	1.3011**	(2.54)
Year 1992	1.0434	(1.04)	0.9538	(-0.99)	1.0436	(1.56)
Year 1993	1.0231	(0.66)	1.0207	(0.51)	1.0236	(1.35)
Year 1995	1.064**	(2.62)	1.0221	(0.79)	1.1414**	(2.90)

[1] Child status variables, household size, mother's schooling variables, region and rural dummies and husband's controls are also included but are not reported .

Table 2.8: Effect of policy change, informal sector probit, by control group[1]

	Control group			
	All women, ages 40-55		Single women, ages 20-35	
Region (Central-west omitted)				
Minas Gerais	-0.0537**	(-7.15)	-0.0436**	(-5.65)
North	-0.0594**	(-6.60)	-0.0592**	(-6.40)
Northeast	0.0803**	(12.15)	0.0744**	(10.89)
Rio	-0.1052**	(-12.72)	-0.0909**	(-10.91)
South	-0.2741**	(-37.35)	-0.2594**	(-34.31)
Sao Paulo	-0.3250**	(-44.52)	-0.3067**	(-40.84)
Rural dummy	0.2565**	(39.22)	0.2542**	(37.55)
Years of completed education				
one	-0.1320**	(-9.57)	-0.2499**	(-11.21)
two	-0.1876**	(-17.58)	-0.2922**	(-16.08)
three	-0.2730**	(-29.42)	-0.4112**	(-25.35)
four	-0.4777**	(-66.07)	-0.6359**	(-45.75)
five	-0.5074**	(-48.14)	-0.7072**	(-42.89)
six	-0.5200**	(-86.67)	-0.7566**	(-43.23)
seven	-0.6953**	(-70.16)	-0.8640**	(-50.56)
eight	-0.8880**	(-102.43)	-1.0596**	(-71.59)
nine	-0.9845**	(-71.86)	-1.1162**	(-57.54)
ten	-1.0958**	(-85.88)	-1.2333**	(-67.43)
eleven	-1.4152**	(-187.69)	-1.5742**	(-116.18)
twelve	-1.5035**	(-69.87)	-1.6396**	(-62.46)
thirteen	-1.4648**	(-74.36)	-1.5794**	(-65.37)
fourteen	-1.5378**	(-103.63)	-1.6740**	(-75.75)
fifteen	-1.5710**	(-145.46)	-1.6739**	(-99.22)
sixteen	-1.1790**	(-78.81)	-1.2653**	(-59.77)
seventeen +	-1.5340**	(-33.43)	-1.6556**	(-23.00)
Household size	0.0077**	(8.89)	0.0141**	(11.64)
GDP growth	0.0091**	(15.91)	0.0117**	(15.57)
Difference-in-differences				
Control dummy	-0.0021	(-0.32)	0.2569**	(34.48)
1990-1999 dummy	0.1643**	(32.21)	0.2585**	(34.56)
Control x 1990-1999	-0.0536**	(-5.98)	-0.0416**	(-3.84)

[1] Marginal effects reported. Z-statistics in parenthesis. Sample is of working women. Treatment group is all working women, ages 20-35 for left panel and working married women, ages 20-35 for right panel.

Table 2.9: Effect of policy change, informal sector multinomial logit[1]

	Not at work		Informal wage workers		Self-employed	
	Control: women ages 40-55. Treatment: women ages 20-35					
Region (Central-west omitted)						
Minas Gerais	0.9292**	(-7.51)	0.8938**	(-8.28)	0.9927	(-0.72)
North	0.9704**	(-3.17)	0.7254**	(-15.18)	1.2164**	(10.59)
Northeast	1.0524**	(2.07)	0.8947**	(-6.20)	1.6305**	(34.74)
Rio	0.9375**	(-6.39)	0.8296**	(-8.94)	1.0600**	(2.88)
South	0.7745**	(-23.44)	0.6065**	(-32.84)	0.7250**	(-19.61)
Sao Paulo	0.7228**	(-27.87)	0.6219**	(-30.03)	0.6504**	(-26.10)
Rural dummy	2.4013**	(88.03)	1.2630**	(19.05)	1.6585**	(37.37)
Years of completed education						
one	0.7358**	(-14.87)	0.6937**	(-15.01)	0.7835**	(-9.78)
two	0.6856**	(-22.50)	0.6445**	(-22.67)	0.6862**	(-17.78)
three	0.5981**	(-34.12)	0.5481**	(-33.77)	0.6088**	(-25.61)
four	0.4517**	(-60.42)	0.3582**	(-64.84)	0.5034**	(-42.50)
five	0.4249**	(-53.03)	0.3708**	(-49.18)	0.3903**	(-38.06)
six	0.4164**	(-53.10)	0.3139**	(-51.90)	0.4142**	(-35.10)
seven	0.3576**	(-63.64)	0.2454**	(-60.88)	0.3559**	(-41.01)
eight	0.2739**	(-90.17)	0.1594**	(-92.79)	0.3137**	(-59.49)
nine	0.2626**	(-73.12)	0.1586**	(-64.58)	0.2231**	(-46.41)
ten	0.2101**	(-86.30)	0.1165**	(-76.64)	0.2246**	(-52.36)
eleven	0.0986**	(-65.63)	0.0622**	(-47.28)	0.1466**	(-8.45)
twelve	0.0978**	(-84.19)	0.0567**	(-58.15)	0.1200**	(-42.08)
thirteen	0.0889**	(-90.94)	0.0592**	(-62.24)	0.1286**	(-45.79)
fourteen	0.0597**	(-15.70)	0.0482**	(-74.85)	0.1071**	(-57.96)
fifteen	0.0398**	(-71.44)	0.0369**	(-12.38)	0.1200**	(-89.65)
sixteen	0.0391**	(-13.35)	0.0590**	(-71.42)	0.2581**	(-46.15)
seventeen +	0.0217**	(-41.63)	0.0348**	(-24.77)	0.1234**	(-22.98)

Table 2.9 (cont.)

	Not at work		Informal wage workers		Self-employed	
Difference-in-differences						
	Control: women ages 40-55. Treatment: women ages 20-35					
Control dummy	1.0026**	(11.35)	0.4938**	(-40.62)	2.0548**	(57.46)
1990-1999 dummy	0.9143**	(-7.66)	1.2530**	(26.99)	1.3607*	(1.64)
Control x 1990-1999	1.0087**	(8.12)	0.7203**	(-2.10)	1.2181**	(17.59)
	Control: married women ages 20-35. Treatment: single women ages 20-35					
Control dummy	4.6417**	(51.73)	0.8414**	(-11.89)	3.0735**	(70.70)
1990-1999 dummy	1.2848**	(13.68)	1.5600**	(36.67)	1.3991**	(7.58)
Control x 1990-1999	0.6764**	(-26.07)	0.7971**	(-6.33)	1.1377**	(10.02)

[1] Relative risk ratios reported, *z*-statistics in parenthesis. Formal work is base category. For top panel control group is women ages 40-55. Results for married 20-35 control group were similar. Household size and GDP growth included but not reported.

Table 2.10: Change in annual occupational and industry segregation, measured by Dissimilarity Index[1]

Year	81	85	90	95	99
Two-digit occupations					
Index (D)	60.56	58.31	56.82	57.81	55.89
Change due to change in cell size (dc)		0.10	-0.92	-0.32	-0.70
Change due to cell composition (ds)		-2.35	-2.82	-2.43	-3.97
Two-digit industries					
Index (D)	54.42	52.43	51.92	52.28	50.89
Change due to change in cell size (dc)		-0.41	-1.02	-0.39	-0.65
Change due to cell composition (ds)		-1.58	-1.47	-1.74	-2.88
Three-digit occupations					
Index (D)	67.06	65.26	63.43	63.42	61.65
Change due to change in cell size (dc)		0.22	-0.67	-0.19	-0.49
Change due to cell composition (ds)		-2.03	-2.96	-3.45	-4.92
Three-digit industries					
Index (D)	58.59	57.25	56.65	57.06	55.99
Change due to change in cell size (dc)		-0.03	-0.58	0.44	-0.39
Change due to cell composition (ds)		-1.31	-1.36	-1.96	-2.21

[1] 1981 is the base year. See text equations (9) and (10) for definition of index.

Table 2.11: Changes in the degree of feminization of occupations, three years[1]

Feminization decile	Proportion of women in each category					Difference in average age of women vs. men			Difference in average schooling of women vs. men		
	1981	1990		1999		1981	1990	1999	1981	1990	1999
		Observed	Expected	Observed	Expected						
1	0.001	0.006	0.002	0.008	0.002	-0.03	-0.83	0.94	-0.07	1.39	2.15
2	0.016	0.019	0.020	0.025	0.023	-4.38	-4.23	-0.64	0.32	1.06	1.73
3	0.061	0.097	0.074	0.109	0.086	-3.84	-3.63	-1.69	0.80	1.67	1.79
4	0.103	0.107	0.124	0.105	0.144	0.36	0.73	1.78	-0.14	-0.31	0.00
5	0.120	0.209	0.145	0.271	0.167	-1.79	-2.93	-2.05	0.37	0.95	0.98
6	0.181	0.253	0.215	0.282	0.245	-2.26	-0.69	0.57	0.89	1.36	1.46
7	0.333	0.421	0.382	0.460	0.422	-2.95	-2.72	-1.77	0.66	0.61	0.78
8	0.434	0.505	0.487	0.516	0.529	-2.72	-1.92	-0.80	1.02	0.44	1.05
9	0.683	0.701	0.728	0.722	0.760	1.21	2.75	3.35	-0.13	-0.40	-0.41
10	0.914	0.885	0.929	0.874	0.939	-3.95	-2.93	-0.90	-2.05	-1.49	-1.96
Overall	0.292	0.347		0.377		-2.06	-1.60	-0.86	1.30	1.18	1.25

[1] In all three years deciles are based on the 1981 orderings. That is, in each decile the occupations are the same across years. "Expected" proportions are those one would expect had women increased their participation in each decile in proportion to their share of employment in that decile. "Observed" proportions are those actually observed.

Table 2.12: Probit regression for the probability of an occupation being highly seggregated[1]

							All effects
Change in occupation's share of employment							
Current period	-4.196**	...	...	...	...	...	-1.297
	(-2.31)	...	...	...	...	...	(-0.34)
Lagged one year	...	-2.617	...	...	...	...	0.852
	...	(-1.54)	...	...	...	...	(0.25)
Lagged two years	...	...	-2.312**	...	...	...	0.520
	...	...	(-1.99)	...	...	...	(0.21)
Lagged three years	...	...	...	-2.25**	...	...	1.160
	...	...	...	(-2.43)	...	...	(0.49)
Lagged four years	...	...	...	...	...	-2.140**	-3.304*
	...	...	...	...	...	(-2.97)	(-1.83)
Test of joint significance (Chi-squared)	...	...	...	...	...	...	12.42**
Change in proportion young in occupation							
Current period	-12.640	...	...	...	...	...	-14.428
	(-0.59)	...	...	...	...	...	(-0.56)
Lagged one year	...	-11.204	...	...	...	...	-0.187
	...	(-0.52)	...	...	...	...	(-0.10)
Lagged two years	...	...	-7.021	...	...	...	1.078
	...	...	(-0.43)	...	...	...	(0.04)
Lagged three years	...	...	...	-6.160	...	...	-0.885
	...	...	...	(-0.44)	...	...	(-0.03)
Lagged four years	...	...	...	...	-6.453	-4.185	-1.840
	...	...	...	...	(-0.72)	(-0.35)	(-0.09)
Test of joint significance (Chi-squared)	...	...	...	...	...	...	0.62
Women's share of total employment	-0.092	-0.089	-0.119	-0.057	-0.065	-0.080	-0.065
	(-1.45)	(-1.05)	(-1.07)	(-1.04)	(-1.20)	(-1.01)	(-0.99)
Occupation is highly segregated in 1981	0.593**	0.595**	0.589**	0.592**	0.595**	0.597**	0.596'**
	(12.62)	(11.86)	(12.26)	(11.55)	(14.94)	(10.28)	(9.60)
Log likelihood	-293.42	-274.10	-255.75	-237.16	-218.60	-218.51	-217.65

1 Highly segregated defined as 95 percent female or male. Marginal effects reported. Z-statistics in parenthesis. Regression weighted by occupation size and is based on values for two digit occupations, from 1981-1999. Highly segregated dummy is defined as one if occupation is 95 percent male or female and zero otherwise. * indicates significant at 10% level; ** at 5%.

Table 2.13: Husband's wage first stage estimates, for married (M) models[1]

Region (Central-west omitted)		
Minas Gerais	-0.4683**	(-13.61)
North	-0.3775**	(-10.44)
Northeast	-0.7187**	(-24.17)
Rio	-0.5957**	(-11.99)
South	-0.3384**	(-10.65)
Sao Paulo	0.3959**	(11.47)
Rural dummy	-0.4805**	(-44.33)
Age	0.1983**	(47.68)
Age squared	-0.0016**	(-33.08)
Years of completed education		
one	0.4335**	(25.40)
two	0.5648**	(39.13)
three	0.8411**	(50.57)
four	1.4182**	(103.24)
five	1.7284**	(66.00)
six	2.0881**	(64.28)
seven	2.3201**	(74.35)
eight	2.9343**	(114.21)
nine	3.0414**	(58.58)
ten	3.7798**	(67.72)
eleven	5.2904**	(154.83)
twelve	7.5071**	(49.43)
thirteen	7.9891**	(63.87)
fourteen	9.0205**	(65.48)
fifteen	12.6260**	(90.93)
sixteen	16.3736**	(132.43)
seventeen +	19.2229**	(39.66)

1 *t*-statistics in parenthesis.

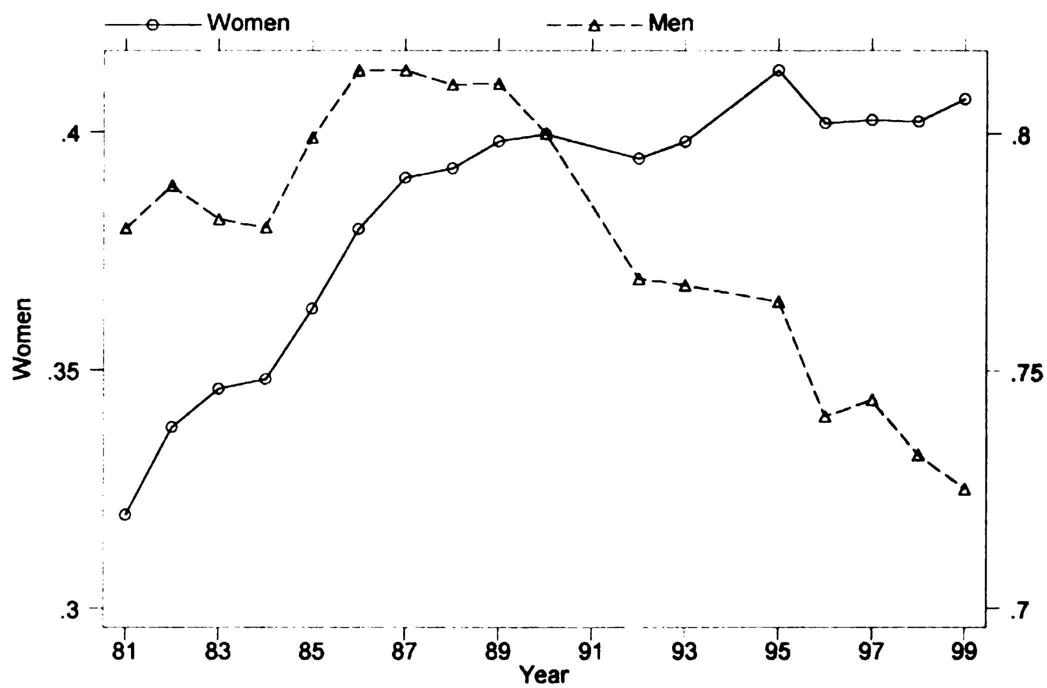

Figure 2.1: Employment rates for men and women, secular trend

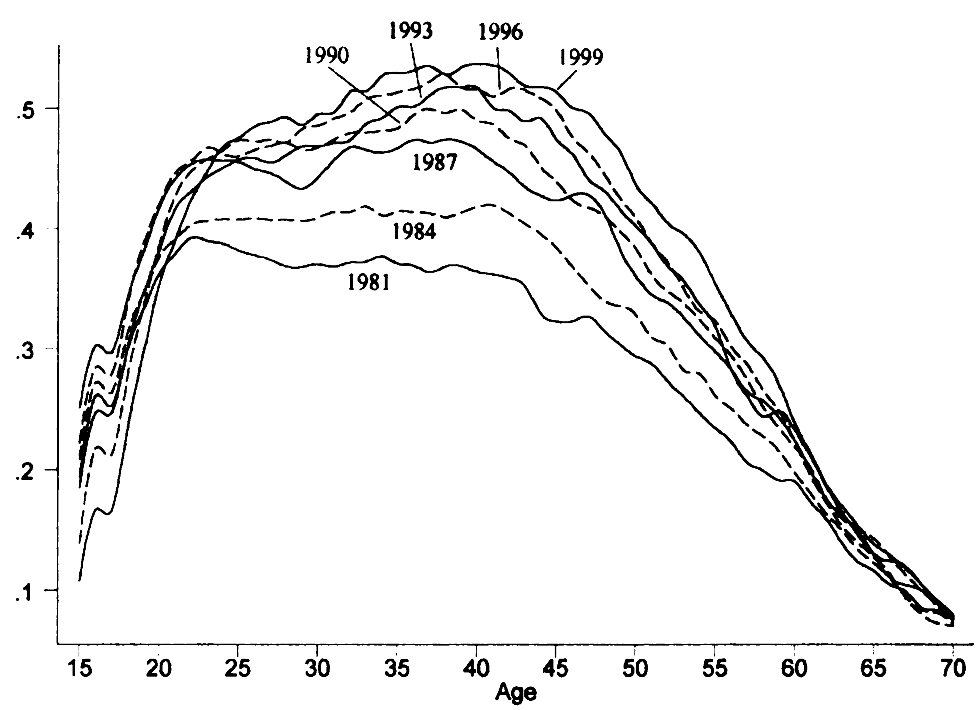

Figure 2.2: Female employment rates, select years, three year moving average

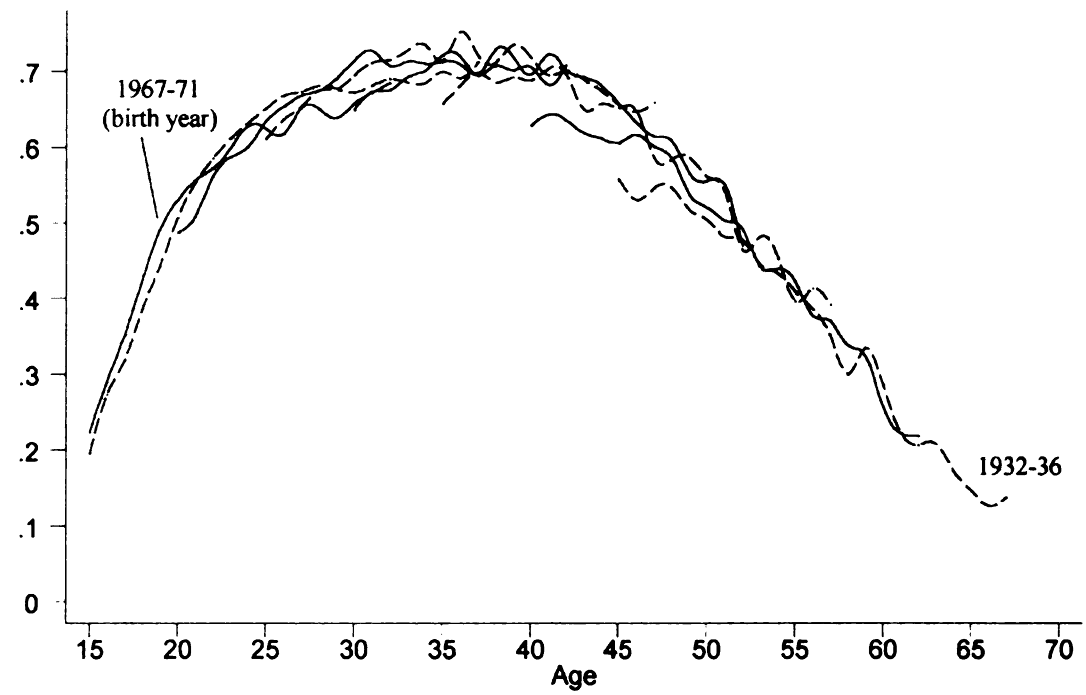

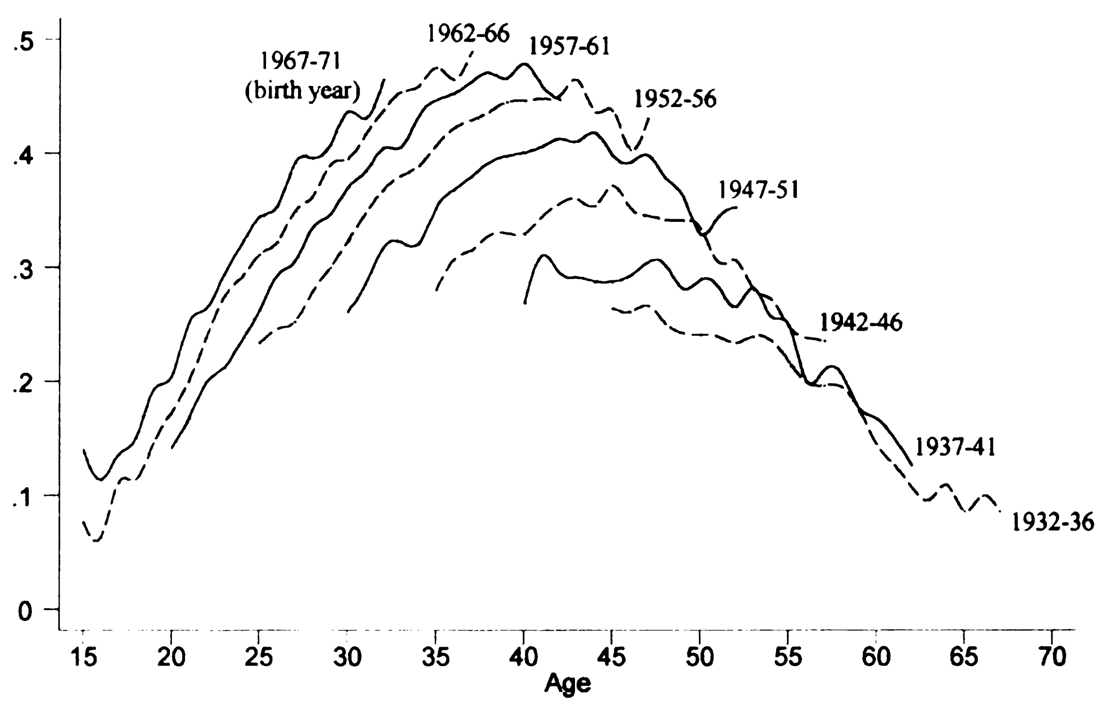

Figure 2.3: Employment rates for single (top) and married (bottom) women, five year cohort groups

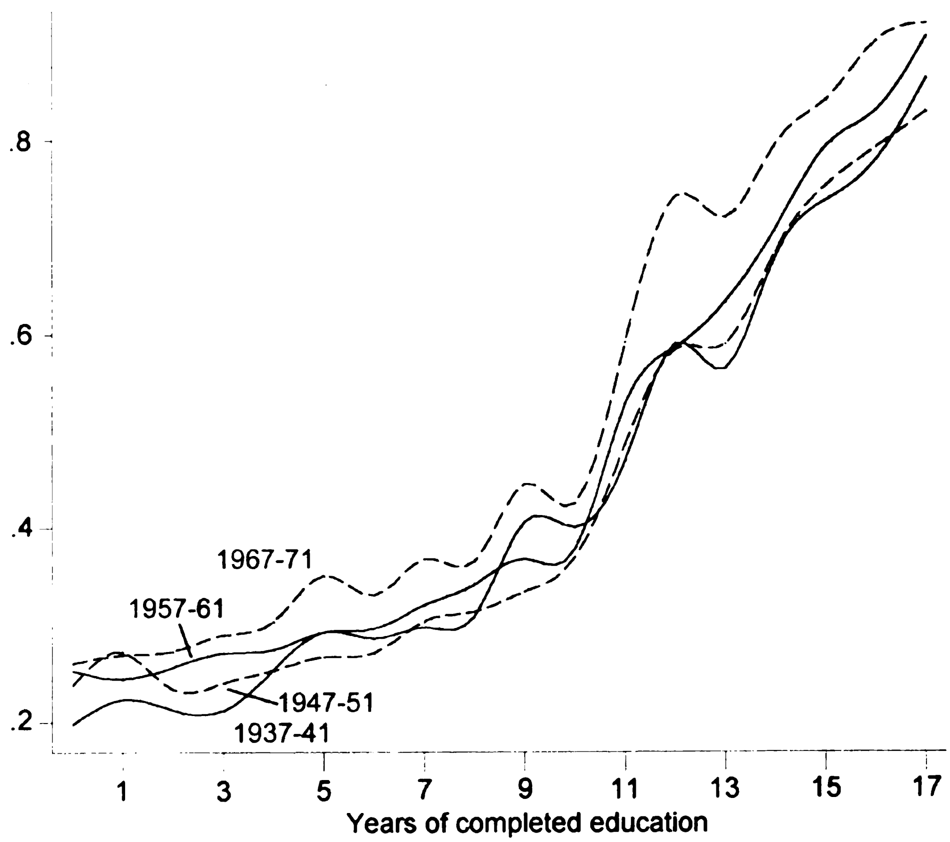

Figure 2.4: Predicted employment rates versus years of education, married women, select cohort groups

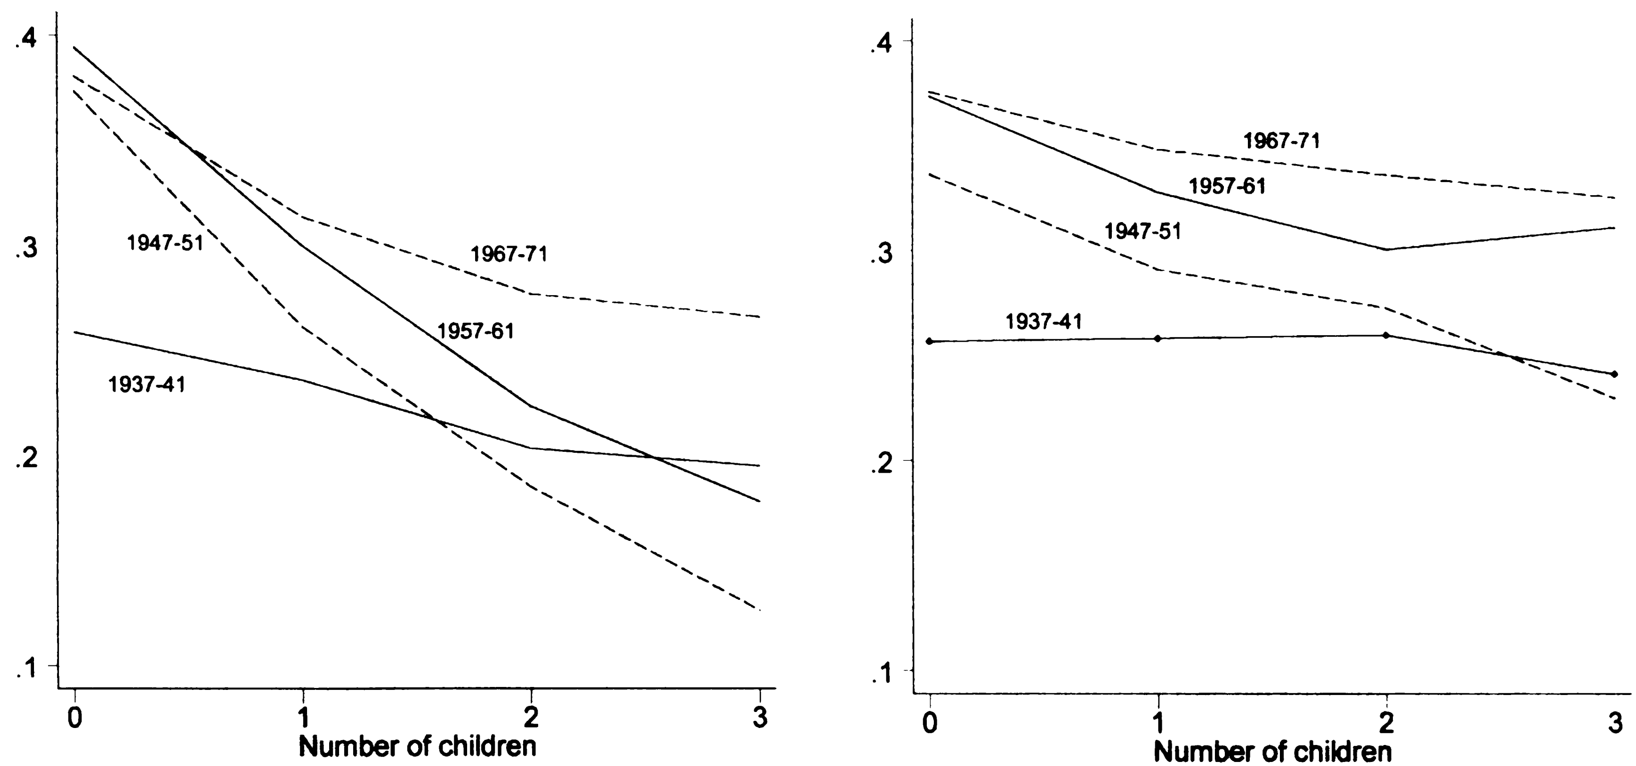

Figure 2.5: Predicted employment rates versus number of boys ages 4 and under (left) and number of boys ages 5-9 (right), select cohort groups

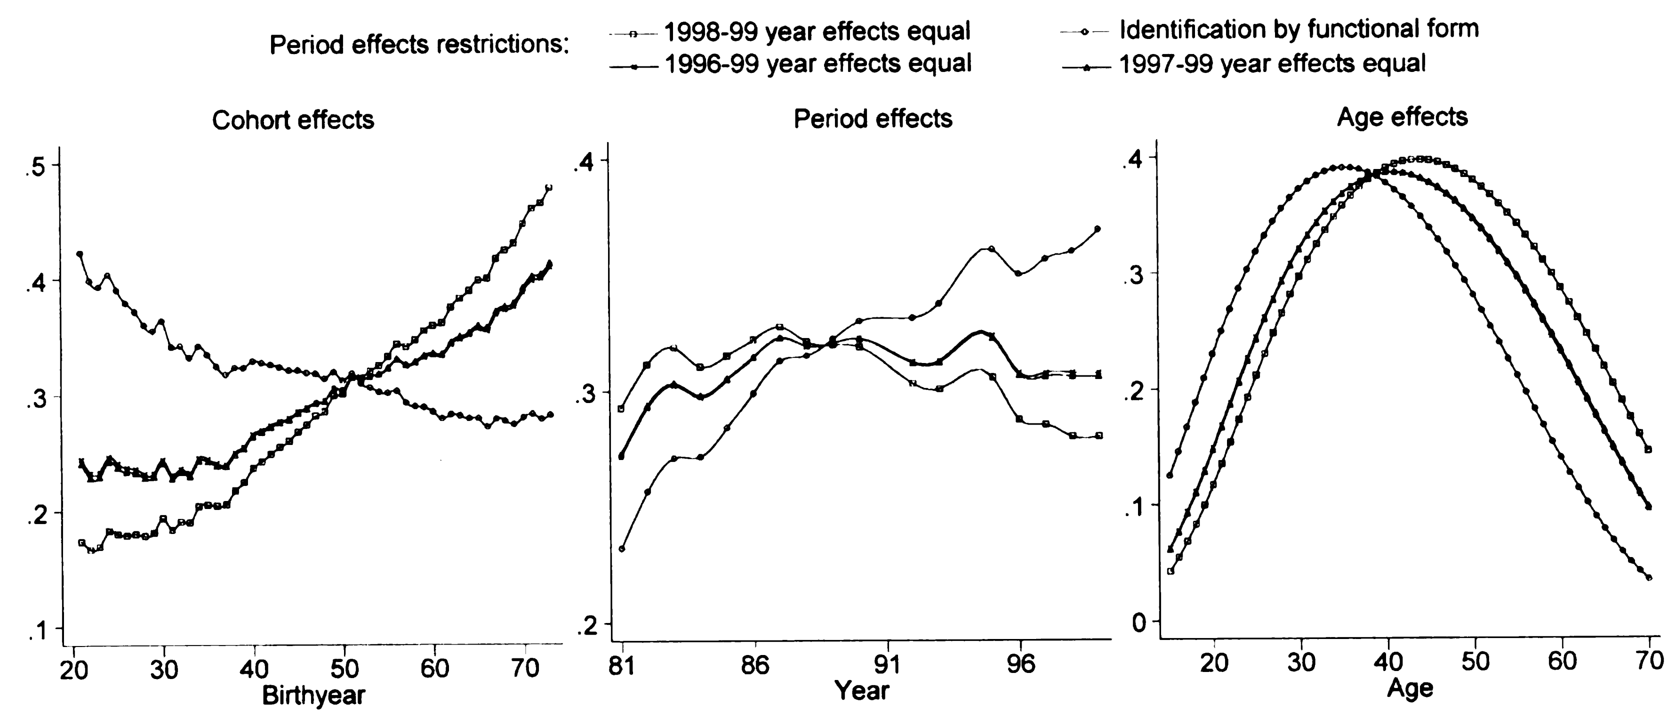

Figure 2.6: Predicted cohort, period and age employment effects for married women, based on model M4, four identification restrictions

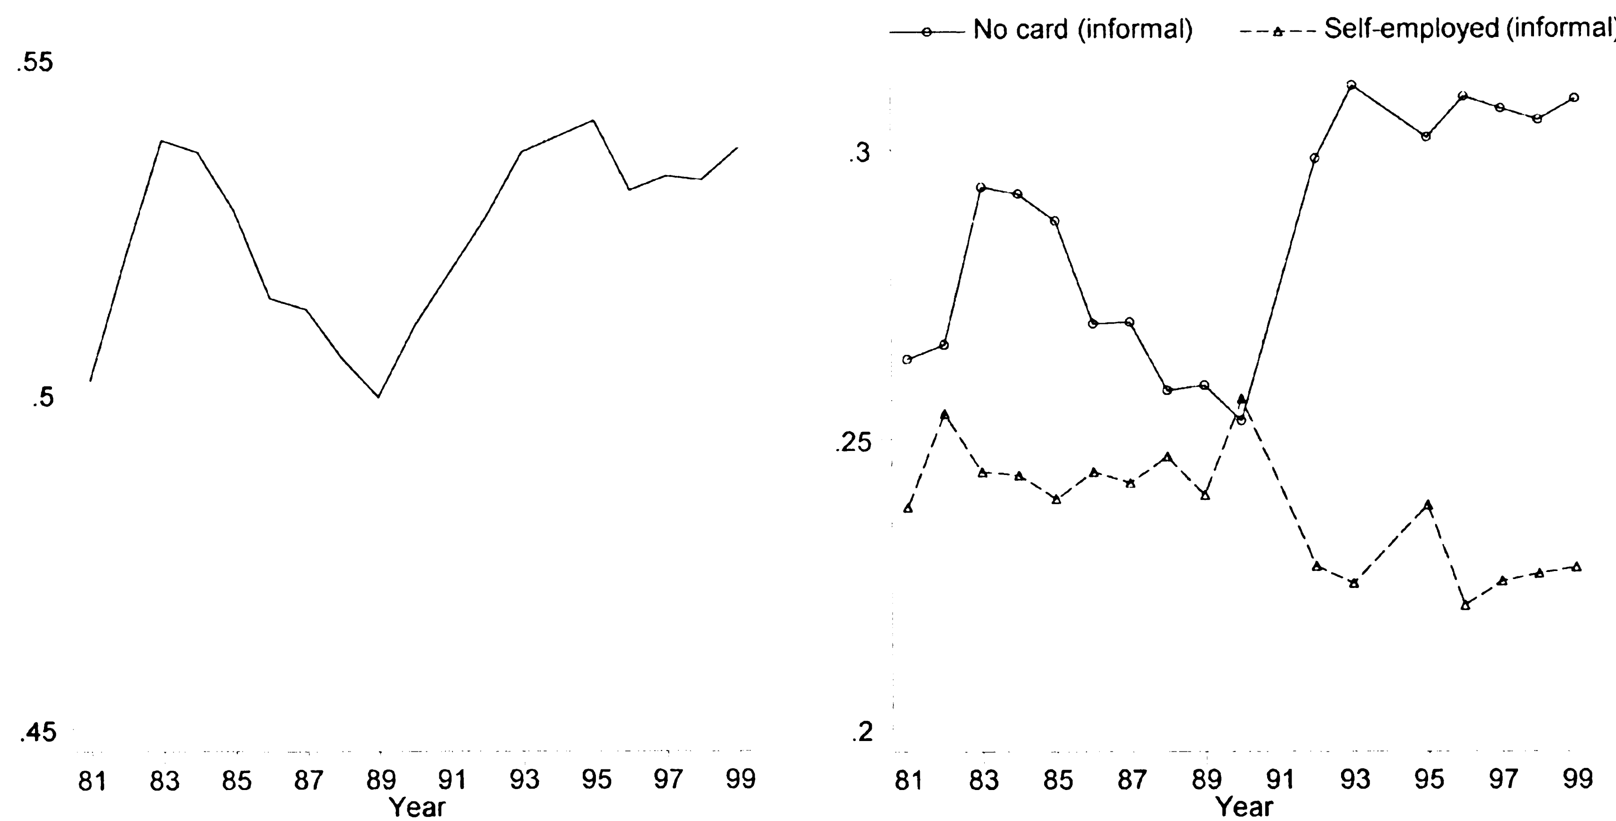

Figure 2.7: Proportion of working women employed informally (left), and by type of informal employment (right)

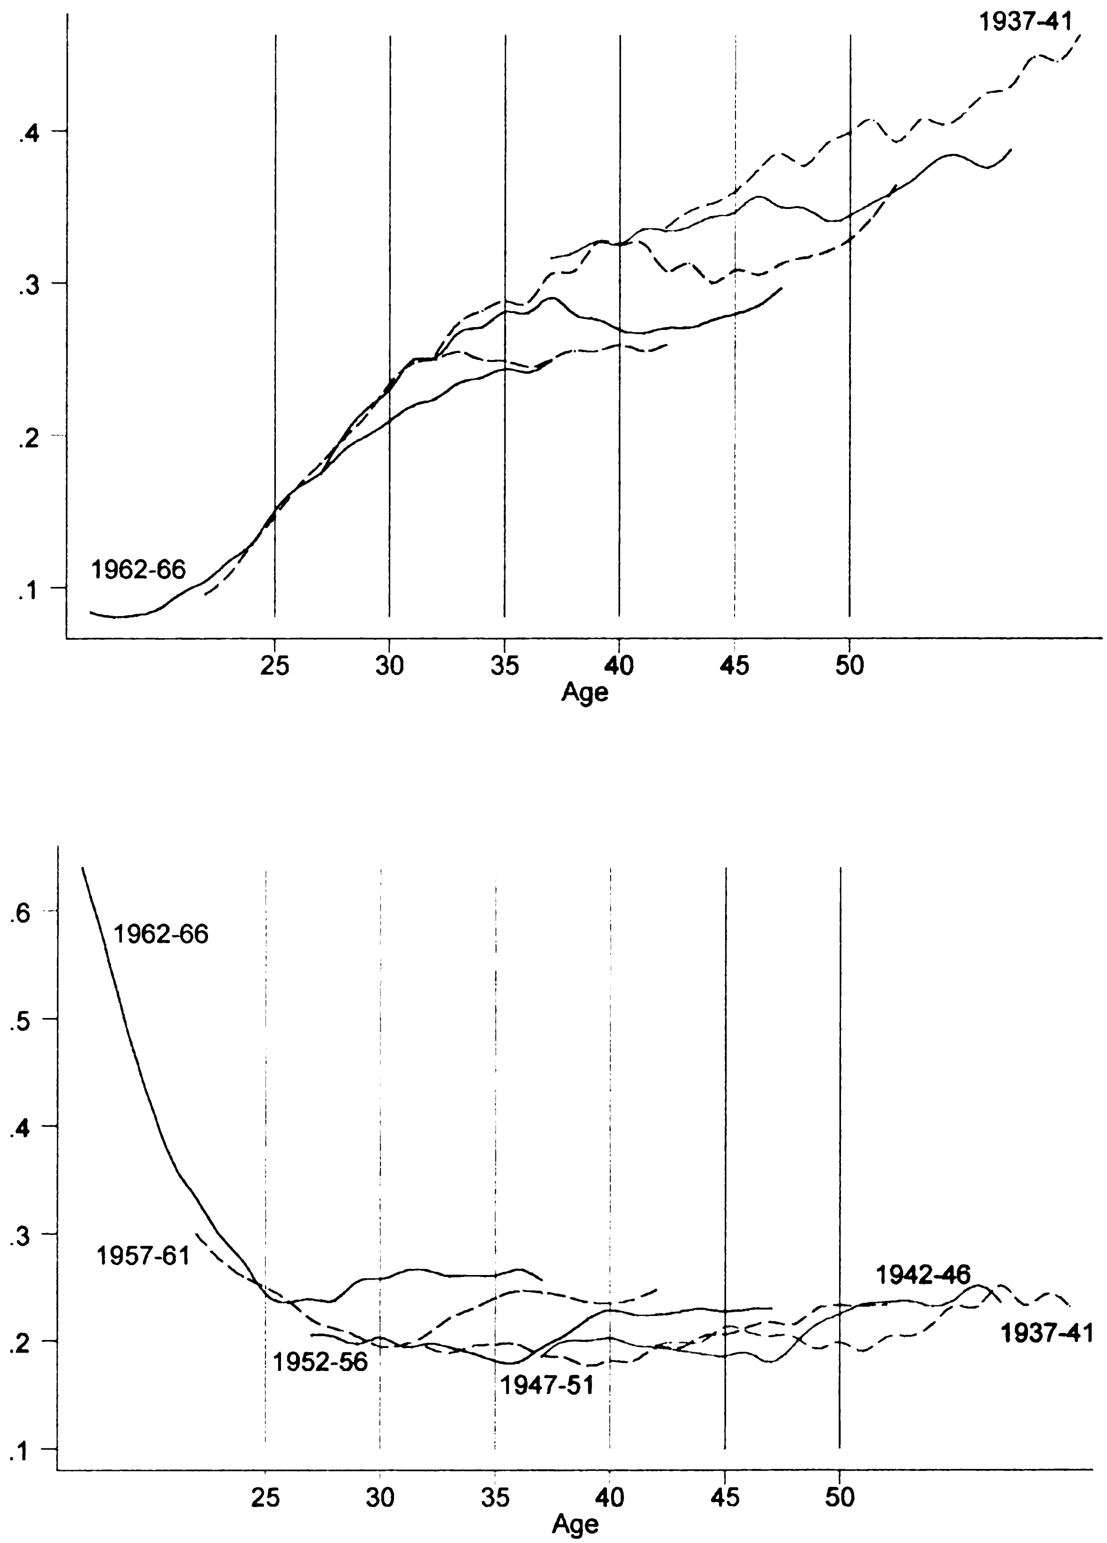

Figure 2.8: Proportion of working women self-employed (top) and proportion employed as informal wage workers (bottom)

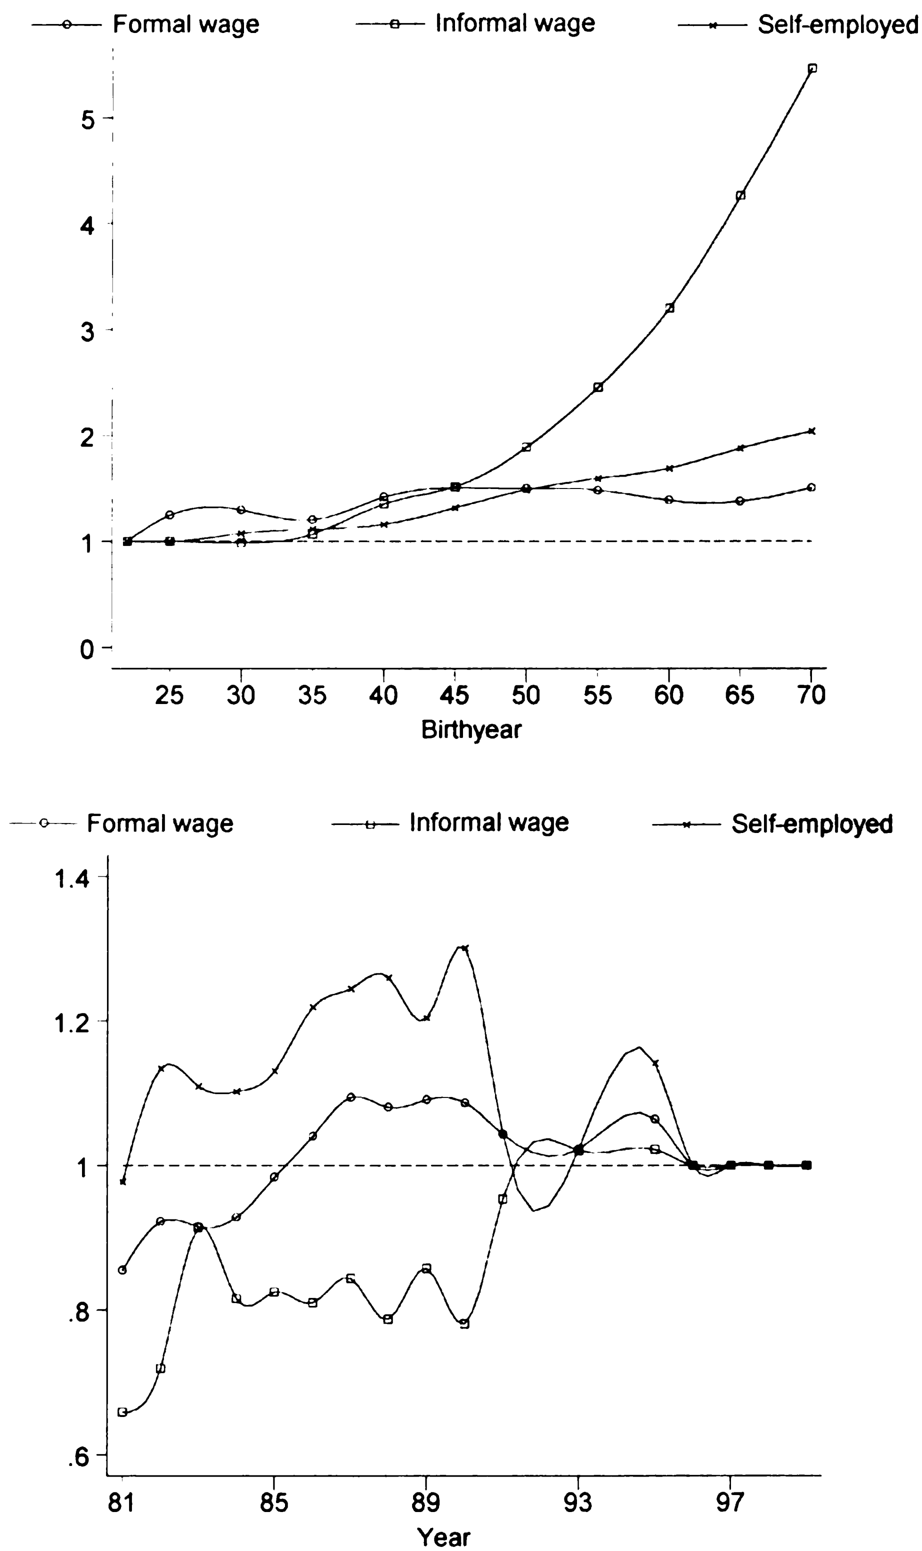

Figure 2.9: Relative risk ratios for period (bottom) and cohort (top) dummy coefficients from sector choice model

APPENDIX 3

Tables and figures for chapter 3

Table 3.1: Summary Statistics

	General Aviation (Part 91)	Commuters (Part 135)	Trunk Carriers (Part 121)
Accident rate[1]			
Mechanical failure	0.019	0.019	0.002
No pilot error	0.019	0.019	0.002
Most serious injury ("none" omitted)			
Fatal	0.203	0.205	0.049
Serious	0.103	0.088	0.196
Minor	0.239	0.103	0.126
Damage ("none" omitted)			
Destroyed	0.270	0.245	0.039
Substantial	0.709	0.587	0.149
Some	0.014	0.117	0.383
Proportion of all on board who suffered:			
Fatal injuries	0.184	0.181	0.022
Serious injuries	0.089	0.062	0.010
Pilot's flight hours (thousands)			
Total	2.6	5.8	12.9
Proportion in make/model	0.321	0.209	0.275
Proportion as pilot in command	0.793	0.863	0.580
Other training			
Pilot's age	44.0	37.0	46.6
Instrument rated	0.496	1.000	1.000
Commercial certificate	0.353	0.636	0.241
Air transport certificate	0.104	0.573	1.000
Pilot is male	0.967	0.968	0.993
VMC	0.900	0.789	0.840
Hours since last inspection	48.5	56.3	224.3
Thrust (thousands of pounds)	394	922	48154
Number of seats	4.0	7.8	149.3
Airframe (hours of flight)	2523	6431	18744
No seat belt used	0.008	0.006	0.004
No harness used	0.187	0.095	0.056
Scheduled Flight	0.000	0.232	0.937
No Flight Plan Filed	0.888	0.459	0.000

[1] Accident rate, accidents per pilot, is weighted.

Table 3.2: Probit of pilot error, full sample

	General Aviation (Part 91)	Commuters (Part 135)	Trunk Carriers (Part 121) All obs.	Trunk Carriers (Part 121) No incidents
Pilot flight hours (thousands)				
Spline, first 1,800 hours	-0.0622**	...	...	...
	(-8.28)	...	...	...
Spline, hours beyond 1,800	-0.0021**	...	...	...
	(7.92)	...	...	...
Total hours (thousands)	...	0.0176	-0.1008**	-0.1744*
	...	(0.82)	(-2.71)	(-1.74)
Total hours squared/100	...	-0.0023	0.008**	0.013*
	...	(-1.10)	(2.96)	(1.74)
Total hours cubic/100	...	0.0001	-0.0002**	-0.0003*
	...	(1.22)	(-2.92)	(-1.68)
Joint significance test (chi-squared)	...	1.79	10.85*	3.07
Proportion in make/model	-0.033**	0.0213	...	0.1050
	(-2.89)	(0.30)	...	(0.55)
Proportion as pilot in command	-0.0663**	0.0509	...	0.1089
-	-3.7200	0.6100	...	0.5000

Table 3.2 (cont.)

	General Aviation (Part 91)	Commuters (Part 135)	Trunk Carriers (Part 121) All obs.	Trunk Carriers (Part 121) No incidents
Male pilot	-0.0150	-0.0203	0.0337	
	-0.8500	-0.2700	0.1600	
Pilot's age	0.0011**	0.0046**	(0.00)	(0.00)
	(3.94)	(2.59)	(0.60)	(0.58)
Landings	(0.01)	0.0558**	0.2115**	0.2127*
	(1.56)	(3.63)	(2.64)	(1.68)
Flight plan filed	-0.0461**	...	...	
	(-4.01)	...	...	
Certificates/ratings				
Instrument rated	-0.0171*	...	...	
	(-1.87)	...	...	
Commercial certificate	-0.0060	...	...	
	(-0.69)	...	...	
Air transport certificate	-0.0214*	-0.0687*	...	
	(-1.73)	(-1.95)	...	
Visual meteorological conditions	-0.1912**	-0.2335**	-0.2596**	-0.2853**
	(-10.95)	(-5.89)	(-5.46)	(-2.71)
Night flying	0.0109	-0.0941**	0.1105**	0.0138
	(1.03)	(-2.94)	(2.77)	(0.16)
N. Observations	12851	1083	671	166
Log of Likelihood	-5894.23	-608.22	77.4981.99	-96.61

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported. z-statistics in parethesis.

Table 3.3: Accident severity probit and ordered probit, General Aviation

	Fatality probit (dP/dx reported)			Ordered probit for severity, all observations	Marinal effects (dP/dx)		
	All obs.	MF	NPE	Beta	Fatality	Severe injury	Minor injury
Pilot's flight hours							
Spline, first 800 hours	-0.0181	-0.0132	-0.0640**	-0.0808	-0.0148	-0.0090	0.0230
	(-0.86)	(-0.58)	(-2.46)	(-1.15)	(-1.11)	(-1.11)	(-1.11)
Spline, hours beyond 800	-0.0009	-0.0018	0.0010	-0.0077**	-0.0014**	-0.0009**	0.0022**
	(-1.20)	(-1.50)	(0.96)	(-2.10)	(-2.10)	(-2.09)	(2.10)
Proportion in Make/Model	-0.0096	-.0271*	-0.0076	-0.0450	-0.0080	-0.0050	0.0131
	(-1.02)	(-1.87)	(-0.49)	(-1.05)	(-1.05)	(-1.05)	(1.05)
Proportion as pilot in command	-0.0557**	-0.0484**	-0.0182	-0.1247*	-0.0225*	-0.0138*	0.0363*
	(-3.80)	(-2.32)	(-0.75)	(-1.83)	(-1.83)	(-1.82)	(1.83)
Other training							
Pilot's age	0.0000	0.0001	0.0001	0.0011	0.0002	0.0001	-0.0003
	(0.40)	(0.36)	(0.36)	(1.05)	(1.05)	(1.05)	(-3.28)
Instrument rated	0.0085	0.0059	0.0109	0.0599	0.0108	0.0067	-0.0174
	(1.04)	(0.14)	(0.93)	(1.62)	(1.62)	(1.62)	(-1.62)
Commercial certificate	-0.0039	-0.0066	0.0006	-0.0278	-0.0050	-0.0031	0.0081
	(-0.51)	(-0.63)	(0.05)	(-0.80)	(-0.81)	(-0.80)	(0.80)
Air Transport certificate	-0.0388**	-0.0127	-0.0125	-0.1277**	-0.0216	-0.0141**	0.0357**
	(-3.37)	(-0.82)	(-0.80)	(-2.49)	(-2.66)	(-2.50)	(2.60)
Male pilot	0.0418**	0.0008	0.0196	.1288*	0.0215**	0.0142*	-0.0357**
	(2.63)	(0.04)	(0.69)	(1.92)	(2.08)	(1.94)	(-2.02)
N. Observations	12851	3639	2100	12851			
Log of Likelihood	-4184.52	-819.81	-360.81	-8300.63			

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported for probit. z-statistics in parethesis.

Table 3.4: Extent of damage probit and ordered probit, General Aviation

	Fatality probit (dP/dx reported)			Ordered probit for severity, all observations	Marinal effects (dP/dx)		
	All obs.	MF	NPE	Beta	Fatality	Severe injury	Minor injury
Pilot's flight hours							
Spline, first 800 hours	-0.0181	0.0242	-0.0307	-0.1069*	-0.0290*	0.0252*	0.0044*
	(-0.91)	(0.72)	(-0.73)	(-1.76)	(1.77)	(1.77)	(1.76)
Spline, hours beyond 800	-0.0007	-0.0014	0.0011	-0.0066**	-0.0018	0.0016**	0.0003**
	(-0.78)	(-0.84)	(0.62)	(-2.22)	(-2.22)	(2.22)	(2.21)
Proportion in make/model	0.0038	-0.0203	0.0005	-0.0287	-0.0080	0.0068	0.0012
	(0.32)	(-0.96)	(0.02)	(0.77)	(-0.77)	(0.77)	(0.77)
Proportion as pilot in command	-0.0300*	-0.0350	-0.0123	-0.0709	-0.0198	0.0168	0.0030
	(1.65)	(-1.08)	(-0.30)	(-1.20)	(-1.20)	(1.20)	(1.20)
Other training							
Pilot's age	-0.0004	-0.0003	-0.0010	0.0002	0.0001	-0.0001	-0.0001
	(-1.24)	(-0.63)	(-1.56)	(0.24)	(0.24)	(-0.24)	(-0.24)
Instrument rated	0.0216**	0.0351**	0.040*	0.0800**	0.0223**	-0.0189**	-0.0033**
	(2.21)	(2.07)	(1.89)	(2.50)	(2.50)	(-2.50)	(-2.48)
Commercial certificate	-0.0137	-0.0414**	-0.0418**	-0.0600**	-0.0166**	0.0141**	0.0026*
	(-1.48)	(-2.67)	(-2.21)	(-1.99)	(-2.00)	(2.01)	(1.94)
Air Transport certificate	-0.0402**	-0.0335	-0.0310	-0.2091**	-0.0542**	0.0437*	0.0105**
	(-2.94)	(-1.46)	(-1.16)	(-4.33)	(-5.12)	(5.45)	(3.97)
Male pilot	0.0304*	0.0329	0.0536	0.1059*	0.0283*	-0.0234**	-0.0050
	(1.88)	(0.95)	(1.18)	(1.83)	(1.91)	(-1.98)	(-1.65)
N. Observations	15649	4388	2500	15649			
Log of Likelihood	-7462.99	-2001.51	-1029.61	-8936.78			

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported for probit. z-statistics in parethesis.

Table 3.5: Accident severity probit and ordered probit, Commuters

	Fatality probit (dP/dx reported)			Ordered probit of severity, all observations	Marginal effects (dP/dx)		
	All obs.	MF	NPE	Beta	Fatality	Severe injury	Minor injury
Pilot's flight hours							
Total hours	-0.0047	-0.0015	0.0106	0.0135	0.0023	0.0014	-0.0038
	(-0.36)	(-0.10)	(0.99)	(0.21)	(0.21)	(0.21)	(-0.21)
Hours squared/100	0.0000	0.0007	-0.0014	-0.0021	-0.0004	-0.0002	0.0006
	(0.01)	(0.38)	(-1.15)	(-0.33)	(-0.33)	(-0.33)	(0.33)
Hours cubed/100	0.0000	-0.0001	0.0000	0.0000	0.0000	0.0000	0.0000
	(-0.11)	(-0.60)	(0.98)	(0.02)	(0.02)	(0.02)	(-0.02)
Test of joint significance (Chi-squared)	4.89	2.02	3.43				
Certificates							
Proportion in make/model	-0.0318	-0.0775**	-0.0071	0.0002	0.0001	0.0000	-0.0001
	(-0.68)	(-1.98)	(-0.32)	(0.01)	(0.01)	(0.01)	(-0.01)
Proportion as pilot in comma	-0.0460	-0.0405	-0.0002	-0.2469	-0.0387	-0.0240	0.0627
	(-0.91)	(-1.40)	(-0.01)	(-1.28)	(-0.89)	(-0.88)	(0.89)
Other training							
Pilot's age	0.0028**	0.0008	0.0014**	0.0096*	0.0017*	0.0010*	-0.0027*
	(2.57)	(1.13)	(2.59)	(1.86)	(1.85)	(1.83)	(-1.85)
Air Transport certificate	-0.0048	0.0004	0.0217	-0.0797	-0.0139	-0.0085	0.0224
	(-0.22)	(0.03)	(1.60)	(-0.75)	(-0.74)	(-0.75)	(0.75)
Male pilot	0.0290	...	...	-0.0947	-0.0173	-0.0100	0.0275
	(0.52)	...	...	(-0.41)	(-0.39)	(-0.41)	(0.40)
N. Observations	1083	363	308	1083			
Log of Likelihood	-353.28	-59.70	-42.09	-678.22			

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported for probit. z-statistics in parenthesis.

Table 3.6: Extent of damage probit and ordered probit, Commuters

	A/C destroyed probit (dP/dx reported)			Ordered probit of extent of damage, all observations	Marginal effects (dP/dx)		
	All obs.	MF	NPE	Beta	Destroyed	Sig. damage	Minor damage
Pilot's flight hours							
Total hours	-0.0051	0.0170	-0.0213	-0.0080	-0.0021	0.0005	0.0016
	(-0.31)	(0.54)	(-0.94)	(-0.24)	(-0.24)	(0.24)	(0.24)
Hours squared/100	0.0005	-0.0003	0.0003	-0.0003	-0.0001	0.0001	0.0001
	(0.32)	(-0.07)	(0.96)	(-0.14)	(-0.14)	(0.14)	(0.14)
Hours cubed/100	-0.0001	-0.0001	-0.0001	0.0000	0.0000	0.0000	0.0000
	(-0.59)	(-0.32)	(-0.99)	(0.16)	(0.16)	(-0.16)	(-0.16)
Test of joint significance (Chi-squared)	3.75	3.45	1.33				
Certificates							
Proportion in make/model	-0.0086	-0.0366	-0.0642	-0.0872	-0.0260	0.0050	0.0175
	(-0.15)	(-0.47)	(-1.00)	(-0.52)	(-0.52)	(0.52)	(0.52)
Proportion as pilot in comma	-0.0439	-0.0004	0.0320	-0.0402	-0.0104	0.0023	0.0081
	(-0.74)	(-0.04)	(0.48)	(-0.22)	(-0.22)	(0.22)	(0.22)
Other training							
Pilot's age	0.0012	-0.0002	0.0023*	0.0064	0.0017	-0.0004	-0.0013
	(0.88)	(-0.08)	(1.91)	(1.55)	(1.55)	(-1.42)	(-1.54)
Air Transport certificate	-0.0030	-0.0449	-0.0112	-0.0534	-0.0139	0.0032	0.0107
	(-0.11)	(-1.25)	(-0.41)	(-0.65)	(-0.65)	(0.62)	(0.66)
Male pilot	-0.0043	...	...	0.0265	0.0075	-0.0015	-0.0060
	(-0.07)	...	...	(0.16)	(0.16)	(-0.18)	(-0.16)
N. Observations	1253	432	481	1253			
Log of Likelihood	-566.44	-160.06	-149.66	-1004.12			

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported for probit. z-statistics in parenthesis.

Table 3.7: Accident severity probit and ordered probit, Trunk Carriers

	Fatality probit (dP/dx reported)			Ordered probit of severity, all observations	Marginal effects (dP/dx)		
	All obs.	MF	NPE	Beta	Fatality	Severe injury	Minor injury
Pilot's flight hours							
Total hours	-0.0020	-0.0016	0.0172	-0.1153	-0.0105	-0.0240	0.0344
	(-0.20)	(-0.15)	(1.22)	(-1.16)	(-1.15)	(-1.16)	(1.16)
Hours squared/100	0.0001	-0.0001	-0.0011	0.0058	0.0005	0.0012	-0.0017
	(0.06)	(-0.05)	(-1.19)	(0.82)	(0.64)	(0.82)	(0.21)
Hours cubed/100	0.0000	0.0000	0.0000	-0.0001	0.0000	-0.0001	0.0000
	(0.21)	(0.29)	(1.30)	(-0.62)	(-0.61)	(-0.62)	(0.62)
Test of joint significance (Chi-squared)	5.14	2.61	7.22*				
Pilot's age	-0.0002	-0.0006	-0.0006	0.0157*	0.0015*	0.0034*	-0.0050*
	(-0.19)	(-0.53)	(0.46)	(1.79)	(1.76)	(1.78)	(-1.80)
Male pilot	-0.0767	-0.0688**	-0.0711**	-0.2468	-0.0276	-0.0526	0.0801
	(-1.49)	(-2.31)	(-2.20)	(-0.43)	(-0.36)	(-0.43)	(0.40)
N. Observations	671	295	467	671			
Log of Likelihood	-117.92	-33.31	-58.41	-429.94			

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported for probit. z-statistics in parenthesis.

Table 3.8: Extent of damage probit and ordered probit, Trunk Carriers

	A/C destroyed probit (dP/dx reported)			Ordered probit of extent of damage, all obs.	Marginal effects (dP/dx)		
	All obs.	MF	NPE	Beta	Fatality	Severe injury	Minor injury
Pilot's flight hours							
Total hours	0.0009	-0.0096	0.0055	0.0770	.0053.	0.0161	-0.0214
	(0.11)	(-0.73)	(0.71)	(0.78)	(0.77)	(0.78)	(-0.78)
Hours squared/100	-0.0002	0.0003	-0.0463	-0.0054	-0.0004	-0.0011	0.0015
	(-0.19)	(0.34)	(-0.87)	(-0.79)	(0.78)	(-0.79)	(0.79)
Hours cubed/100	0.0000	0.0000	0.0001	0.0001	0.0000	0.0000	-0.0001
	(0.49)	(0.05)	(1.08)	(1.00)	(0.99)	(1.00)	(-1.00)
Test of joint significance (Chi-squared)	8.27**	5.90	0.73				
Pilot's age	-0.0005	0.0007	-0.0004	-0.0050	-0.0040	-0.0011	0.0014
	(-0.59)	(0.51)	(-0.58)	(-0.56)	(-0.56)	(-0.56)	(0.56)
Male pilot	-0.0478	-0.0766*	-.0377**	-0.1029	-0.0078	-0.0220	0.0298
	(-1.32)	(-1.70)	(-2.03)	(-0.18)	(-0.16)	(-0.17)	(0.17)
N. Observations	751	315	502	751			
Log of Likelihood	-440.50	-51.67	-35.65	-440.50			

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported for probit. z-statistics in parenthesis.

Table 3.9: Probit of fatality or severe injury

	General Aviation (Part 91)	Commuters (Part 135)	Trunk Carriers (Part 121)
Pilot flight hours (thousands)			
Spline, first 1,800 hours	-0.0024**	...	...
	(-2.21)	...	...
Spline, hours beyond 1,800	-0.0177	...	...
	(-0.83)	...	...
Total hours (thousands)	...	0.0108	-0.0424
	...	(0.64)	(-1.37)
Total hours squared/100	...	-0.0011	0.0023
	...	(-0.68)	(1.03)
Total hours cubic/100	...	0.0000	0.0000
	...	(0.25)	(-0.89)
Joint significance test (chi-squared)		5.40	5.13
Proportion in make/model	-0.0116	0.0209	...
	(-0.90)	(0.34)	...
Proportion as pilot in command	-0.0186	-0.0510	...
	(-0.90)	(-0.71)	...
Male pilot	0.0279	-0.0503	-0.0007
	(1.39)	(-0.75)	(0.00)
Pilot's age	(0.000)	(0.002)	0.0059**
	(1.35)	(1.44)	(2.07)
Landings	(-0.000)	(0.009)	(0.032)
	(-0.05)	(0.83)	(0.63)
Flight plan filed	-0.0257*	...	...
	(-1.87)	...	...
Certificates/ratings		...	...
Instrument rated	0.0177	...	...
	(1.60)	...	...
Commercial certificate	-0.0083	...	...
	(-0.79)	...	...
Air transport certificate	-0.029*	-0.0344	...
	(-1.89)	(-1.14)	...
Visual meteorological conditions	-0.3061**	-0.2106**	-0.0779*
	(-21.66)	(-7.25)	(-1.83)
Night flying	0.0778**	0.0663**	0.0024
	(6.92)	(2.35)	(0.07)
N. Observations	12851	1083	671
Log of Likelihood	-6392.91	-515.76	-348.91

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported. z-statistics in parethesis.

Table 3.10: Tobit model for proportion killed, all observations

	General Aviation (Part 91)	Commuters (Part 135)	Trunk Carriers (Part 121)
Pilot flight hours (thousands)			
Spline, first 1,800 hours	-0.0349	...	...
	(-1.300)	...	...
Spline, hours beyond 1,800	-0.39960	...	...
	(-0.770)	...	...
Total hours (thousands)	...	-0.1178	-0.2474
	...	(-0.320)	(-0.600)
Total hours squared/100	...	0.0008	0.0147
	...	(0.020)	(0.500)
Total hours cubic/100	...	-0.0001	-0.0002
	...	(-0.130)	(-0.310)
Joint significance test (Chi-squared)	...	1.32	0.58
Proportion in make/model	-0.2381	-0.8552	...
	(-0.760)	(-0.660)	...
Proportion as pilot in command	-1.7902**	-1.4572	...
	(-3.600)	(-1.050)	...
Male pilot	1.1549**	0.77960	-2.7472*
	2.2000	0.5000	-1.7000
Pilot's age	(-0.000)	0.0789**	(-0.014)
	-0.0100	2.4600	-0.3300
Landings	(0.097)	(0.167)	(-0.026)
	0.6100	0.7700	-0.0400
Flight plan filed	(-0.394)	...	...
	-1.2200	...	...
Certificates/ratings		...	...
Instrument rated	0.3429	...	...
	(1.250)	...	...
Commercial certificate	-0.2600	...	...
	(-1.010)	...	...
Air transport certificate	-1.3317**	0.1437	...
	(-3.410)	(0.240)	...
Visual meteorological conditions	-6.8281**	-3.7468**	-1.0204*
	(-14.960)	(-4.660)	(-1.690)
Night flying	1.7594**	1.4208**	0.5008
	(6.540)	(2.420)	(0.960)
N. Observations	12851	1083	671
Log of Likelihood	-4836.35	-411.52	-102.15

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported. z-statistics in parethesis.

Table 3.11: Probit of accident, by selection criteria

	Mechanical failure			No pilot error		
	G. Aviation	Commuters	Trunk Carriers	G. Aviation	Commuters	Trunk Carriers
Pilot flight hours (thousands)						
Spline, first 1,200 hours	-0.0064**	...	...	-0.0101**	...	...
	(-6.010)	...	...	(-8.050)	...	...
Spline, hours beyond 1,200	-0.0002**	...	...	-0.0002**	...	...
	(-3.020)	...	...	(-2.750)	...	...
Total hours (thousands)	...	-0.0009	-0.0001	...	0.0016	-0.0001
	...	(0.790)	(-0.710)	...	(0.980)	(-0.390)
Total hours squared/100	...	-0.0002	0.0007	...	-0.0275*	0.0000
	...	(-0.020)	(0.500)	...	(-1.880)	(0.560)
Total hours cubic/100	...	0.0001	0.0000	...	0.0097**	0.0000
	...	(0.530)	(-0.460)	...	(2.500)	(0.580)
Joint significance test (Chi-squared)		5.40	2.37		11.97**	1.46
Proportion in make/model	0.0023**	0.0000	...	-0.0009	0.0001	...
	(2.280)	(0.000)	...	(-0.730)	(0.020)	...
Proportion as pilot in command	-0.0096**	0.0001	...	-0.0112**	-0.0011	...
	(-6.470)	(0.030)	...	(-6.020)	(-0.020)	...

Table 3.11 (cont.)

	Mechanical failure			No pilot error		
	G. Aviation	Commuters	Trunk Carriers	G. Aviation	Commuters	Trunk Carriers
Other training						
Pilot's age	0.0001**	0.0001	0.0001**	0.0001**	0.0001	0.0000
	(3.810)	(1.100)	(2.120)	(2.910)	(1.000)	(0.213)
Instrument rated	0.0006	...	...	-0.0007	...	...
	(0.810)	...	...	(-0.710)	...	...
Commercial certificate	0.0007	...	...	0.0001	...	...
	(0.850)	...	...	(0.150)	...	...
Air transport certificate	0.0018	0.0021	...	-0.0010	-0.0038	...
	(1.630)	(0.870)	...	(-0.810)	(-1.470)	...
Male	-0.0021	0.0011	0.0003	-0.0023	0.0068	0.0002
	(-1.320)	(0.170)	(0.360)	(-1.090)	(0.790)	(0.190)
Landings, last 90 days (thousands)	0.0012**	0.003**	0.0002*	0.0005	0.0044	0.0015**
	(2.980)	(3.010)	(1.810)	(0.890)	(3.630)	(2.130)
Scheduled flight	...	...	...	...	0.0020	...
	...	...	...	...	(0.520)	...
Visual meteorological conditions	-0.0151**	-0.0158**	-0.011**	-0.0239**	-0.0219**	-0.0006**
	(-7.370)	(-4.520)	(-3.080)	(-8.430)	(-5.640)	(-1.980)
Night flying	-0.0005	0.0004	0.0005*	0.0024*	-0.0047	0.0002
	(-0.530)	(0.700)	(1.870)	(1.880)	(-1.680)	(0.710)
Flight plan filed	-0.0065**	...	...	-0.0046**	...	...
	(-5.230)	...	...	(-2.990)	...	...
N. Observations	28190	2157	1130	25824	2076	1341
Log of Likelihood	-2474.88	-195.65	-11.09	-2236.75	-191.34	-15.36

*** indicates significant at the 10% level; ** at the 5% level. dF/dx reported. z-statistics in parenthesis.**

Table 3.12: Probit of accident, Trunk carriers, incidents not used

	Mechanical failure	No pilot error
Pilot flight hours (thousands)		
Total hours (thousands)	-0.0027**	-0.0005
	(-2.450)	(-0.880)
Total hours squared/100	0.01661**	0.0027
	(2.370)	(0.770)
Total hours cubic/100	-0.0003**	-0.0001
	(-2.300)	(-0.670)
Joint significance test (Chi-squared)	7.47**	0.98
Proportion in make/model	-0.0023	0.0005
	(-0.790)	(0.490)
Proportion as pilot in command	0.0033	0.0008
	(1.370)	(0.650)
Other training		
Pilot's age	0.0003**	0.0000
	(2.640)	(0.160)
Male	0.0013	0.0009
	(0.410)	(0.370)
Landings, last 90 days (thousands)	-0.0006	(0.0029)
	(-0.680)	(2.920)
Scheduled flight	0.0011	-0.0014
	(0.730)	(-1.380)
Visual meteorological conditions	-0.0056**	-0.0011
	(-2.690)	(-1.610)
Night flying	-0.0005	0.0001
	(-0.510)	(0.190)
N. Observations	254	323
Log of Likelihood	-50.55	-42.03

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported. z-statistics in parenthesis

Table 3.13: Description of aviation terms

Term	Description	Comments
Accident	The FAA classifies an accident as "an occurrence associated with the operation of an aircraft which takes place between the time any person boards the aircraft with the intention of flight and all such persons have disembarked, and in which any person suffers death or serious injury, or in which the aircraft receives substantial damage."	
Certificates	There are a number of certificates available to pilots. The most basic is the "Private pilot" certificate. This allows pilots to fly under VFR, and not commercially. An instrument rating would then allow the pilot to fly under IFR, but still not commercially. A commercial certificate allows operation under chapter 135 of the CFR, "Commercial" flights. An air transport certificate allows the pilot to operate under part 121 of the CFR, "Trunk carrier" flights.	
CFR	Code of Federal Regulation	
Commuter	Refers to flights conducted under chapter 135 of the CFR. Small commuter flights, both scheduled (24 percent) and non-scheduled (76 percent). Also includes on-demand air taxi flights. 87 percent of entries are accidents.	Plane capacity can not exceed 30 seats. Some flights require commercial certificate, others require air transport certificate. With few exceptions, instrument rating is also required.
General aviation	Also known as "Part 91" flights, referring to the chapter of the CFR in which these flights are regulated. Mostly private flights, conducted for personal and recreational purposes. No flights are scheduled, virtual all entries in the database are accidents.	Private pilot license required.

Table 3.13 (cont.)

Term	Description	Comments
IFR	Instrument flight rules. Rules which govern flights which are conducted when "instrument meteorological conditions" (IMC) exist. The existence of IMC depends on visibility, cloud cover, time of day and class of airspace. Loosely speaking, when the weather is bad pilots must fly under IFR rules and must file a flight plan.	In practice, most commercial aviation operate under IFR rules regardless of the weather, in order to take advantage of the radar and ATC services available only to IFR flights.
Incident	The FAA classifies and incident as "an occurrence other than an accident, associated with the operation of an aircraft, which affects or could affect the safety of operations."	
Instrument Rating	Rating that enables a pilot to fly under IFR.	To obtain the rating the plot must undergo specific training requirements which require both time in ground and flight training.
Trunk carriers	Refers to flights conducted under chapter 121 of the CFR. both trunk carriers and large commuter flights. Most flights are scheduled (94 percent). Most entries are incidents, only 39 percent of entries are accidents.	Requires air transport certificate and instrument rating.
VFR	Visual flight rules. Rules that govern operations conducted when "visual meteorological conditions" (VMC) exist. All flights are either conducted under IFR or VFR.	

Table 3.14: Probit for pilot error, with trends and interactions

Pilot flight hours (thousands)	General Aviation (Part 91)	Commuters (Part 135)	Trunk Carriers (Part 121)
Total hours (thousands)	-0.0014	-0.1492**	0.0639*
	(-1.11)	(-1.94)	(1.74)
Spline dummy	-0.0726**	...	...
	(-6.02)	...	...
Total hours squared/100	...	0.0094*	-0.0055
	...	(1.75)	(-1.47)
Total hours cubic/100	...	-0.0002	0.0001
	...	(-1.51)	(1.32)
Proportion in make/model	-0.0277	...	-0.0351
	(-1.39)	...	(-0.29)
Proportion as pilot in command	-0.0022	...	0.0277
	(-0.07)	...	(0.20)
Trend	0.0064	-0.0409	0.0157
	(1.50)	(-1.13)	(0.89)
Trends Interactions			
Trend x total hours	-0.0001	0.0017	-0.0069
	(-0.42)	(0.19)	(-1.69)
Trend x spline dummy	0.0018	...	...
	(1.23)	...	...
Trend x total hours squared/100	...	0.0002	0.0005
	...	(0.28)	(1.24)
Trend x total hours cubic/100	...	-0.0001	-0.0001
	...	(-0.63)	(-0.99)
Trend x proportion make/model	-0.0008	...	0.0067
	(-0.31)	...	(0.45)
Trend x proportion pilot in command	-0.0101**	...	0.0010
	(-2.37)	...	(0.06)
Marginal trend effect (dP/dTrend)	-0.0037	-0.0154	-0.0023
Joint significance for trend interactions (Chi-squared)	(6.12)	(4.95)	(5.41)

* indicates significant at the 10% level; ** at the 5% level. dF/dx reported. z-statistics in parethesis.

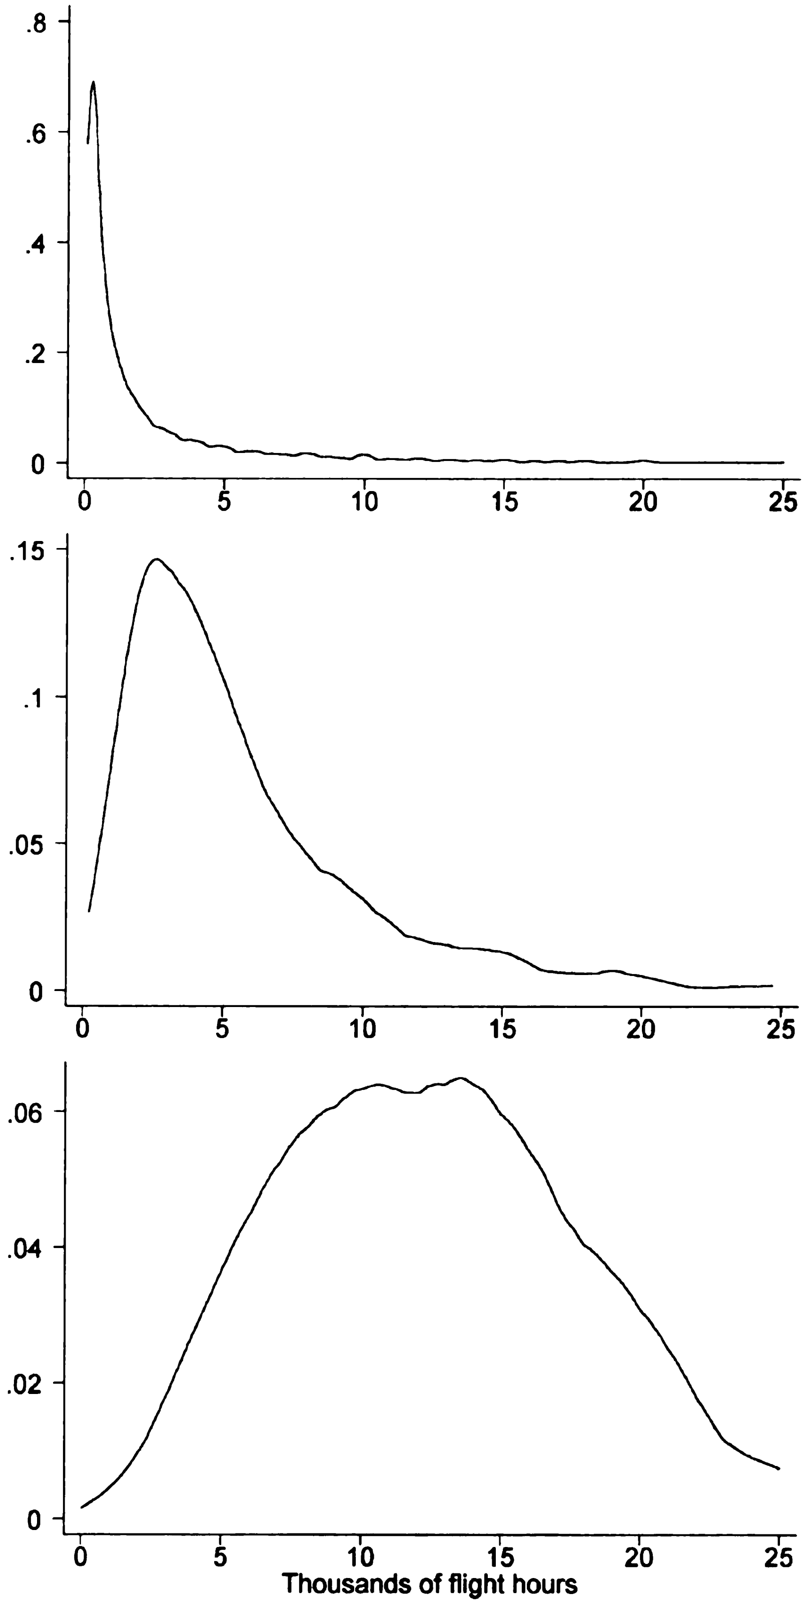

Figure 3.1: Density function for flight hours, general aviation (top), commuters (middle) and trunk carriers (bottom)

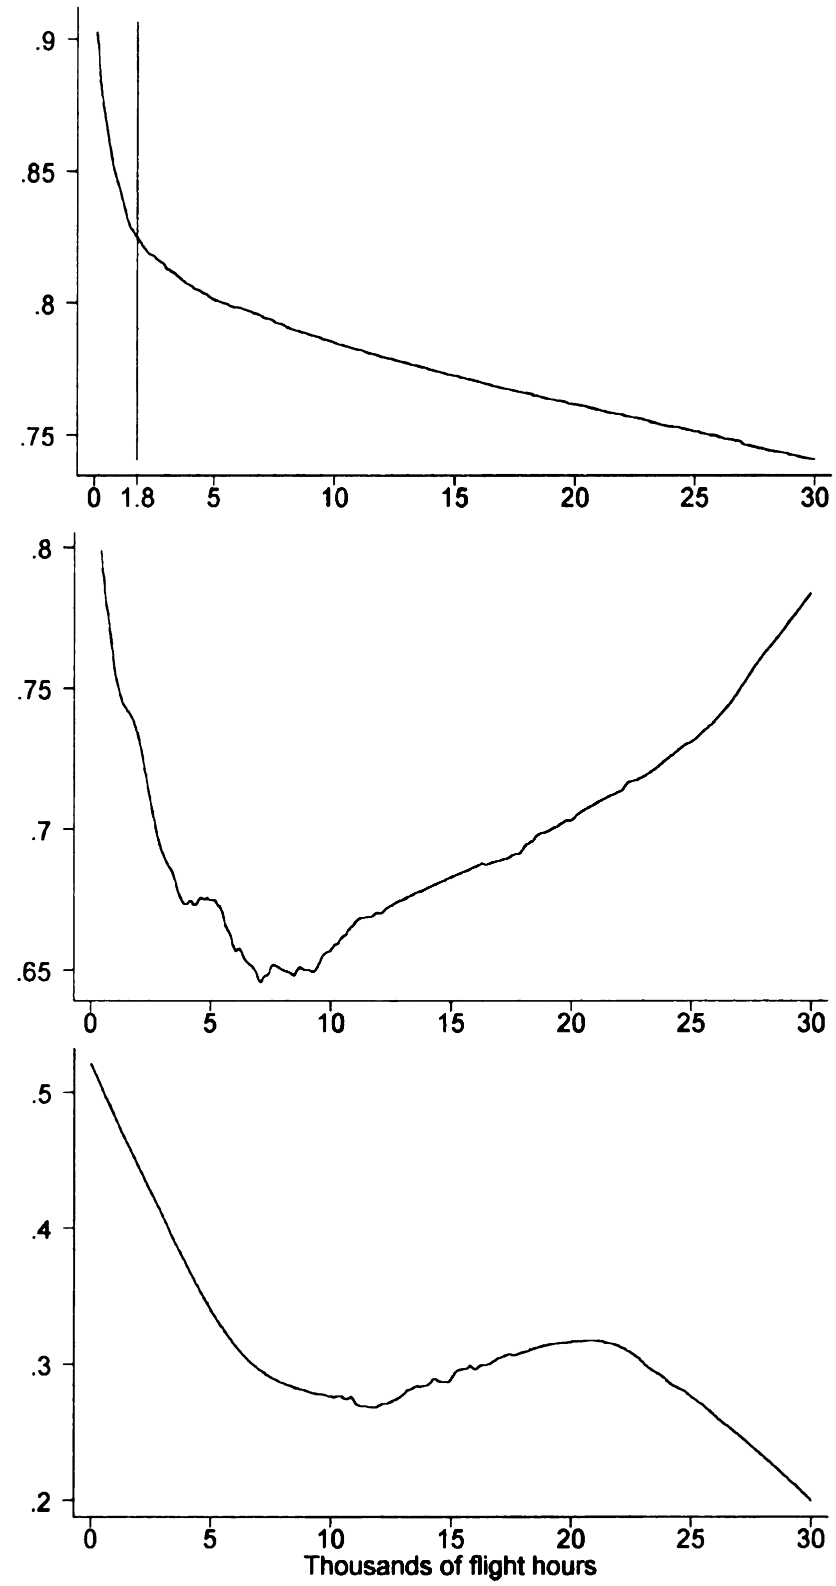

Figure 3.2: Non-parametric estimates of pilot error probability, general aviation (top), commuters (middle) and trunk carriers (bottom). LOWESS, bw=0.8

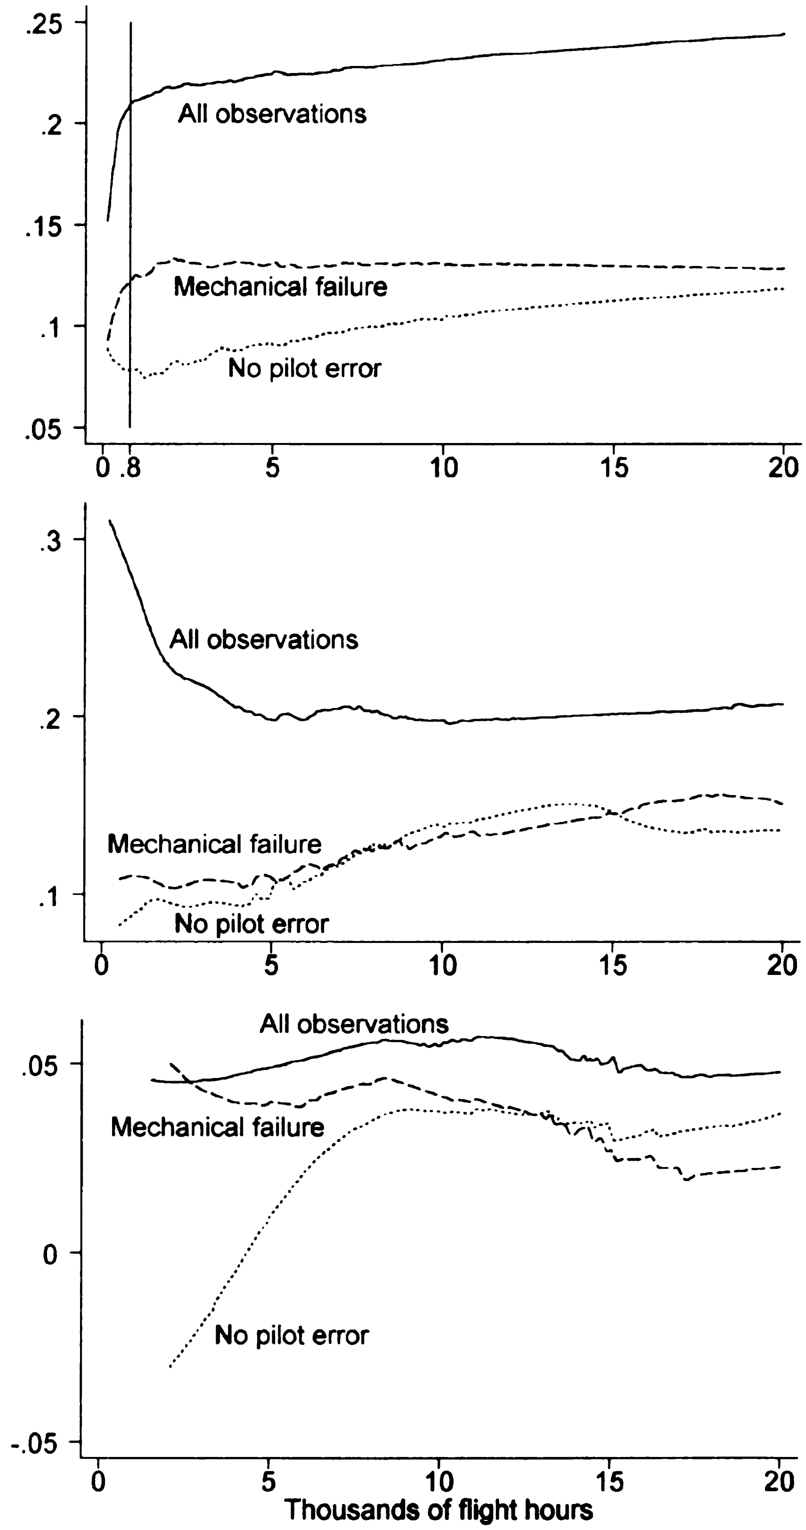

Figure 3.3: Non-parametric estimates of fatality probability, general aviation (top), commuters (middle) and trunk carriers (bottom). LOWESS, bw=0.8

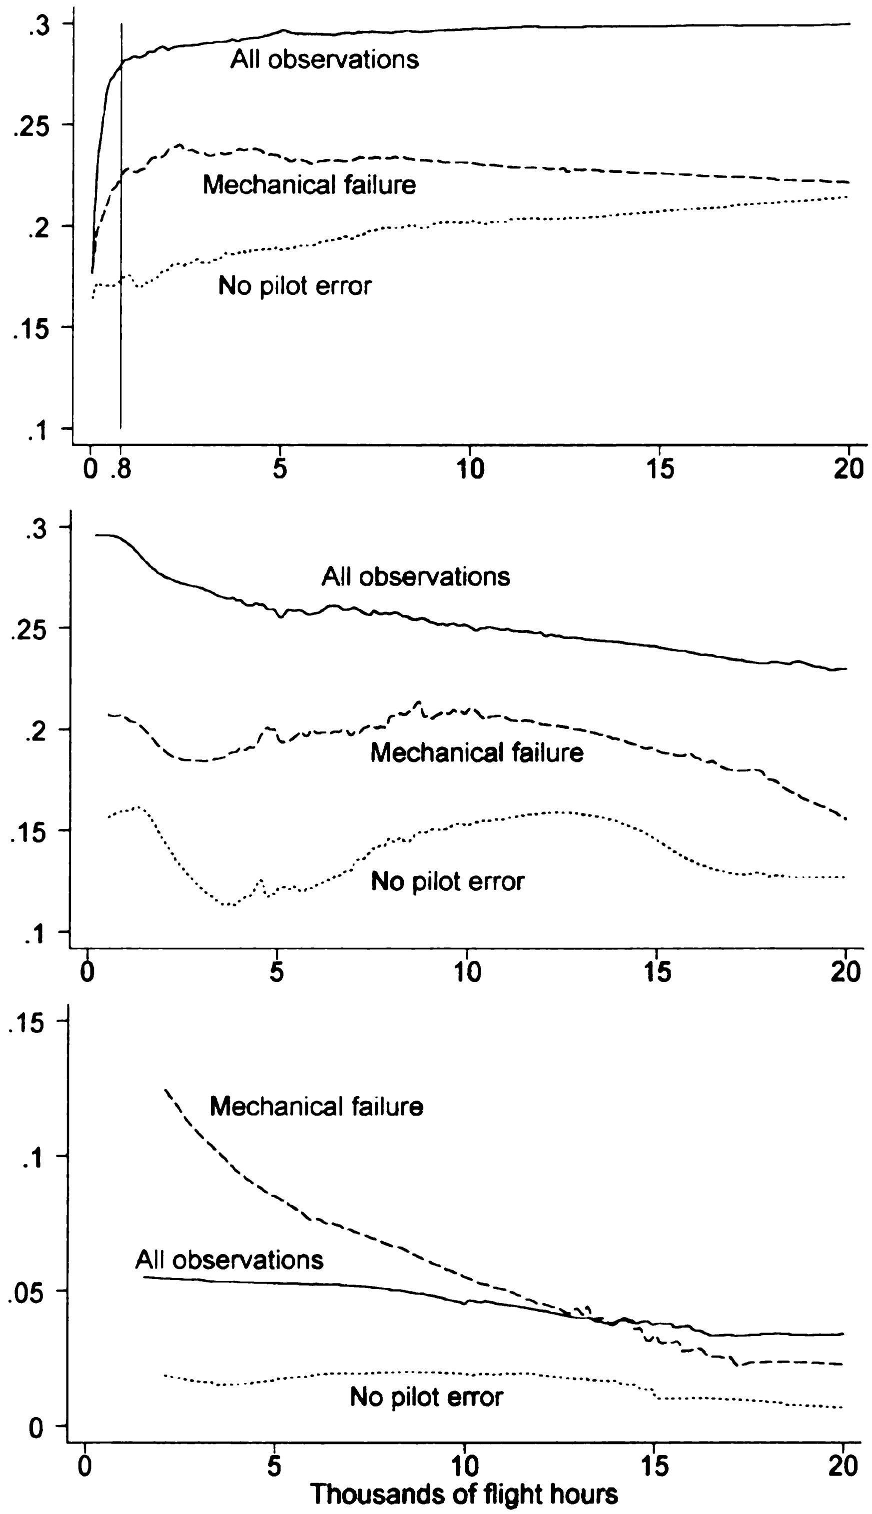

Figure 3.4: Non-parametric estimates of probability destroyed, general aviation (top), commuters (middle) and trunk carriers (bottom). LOWESS, bw=0.8

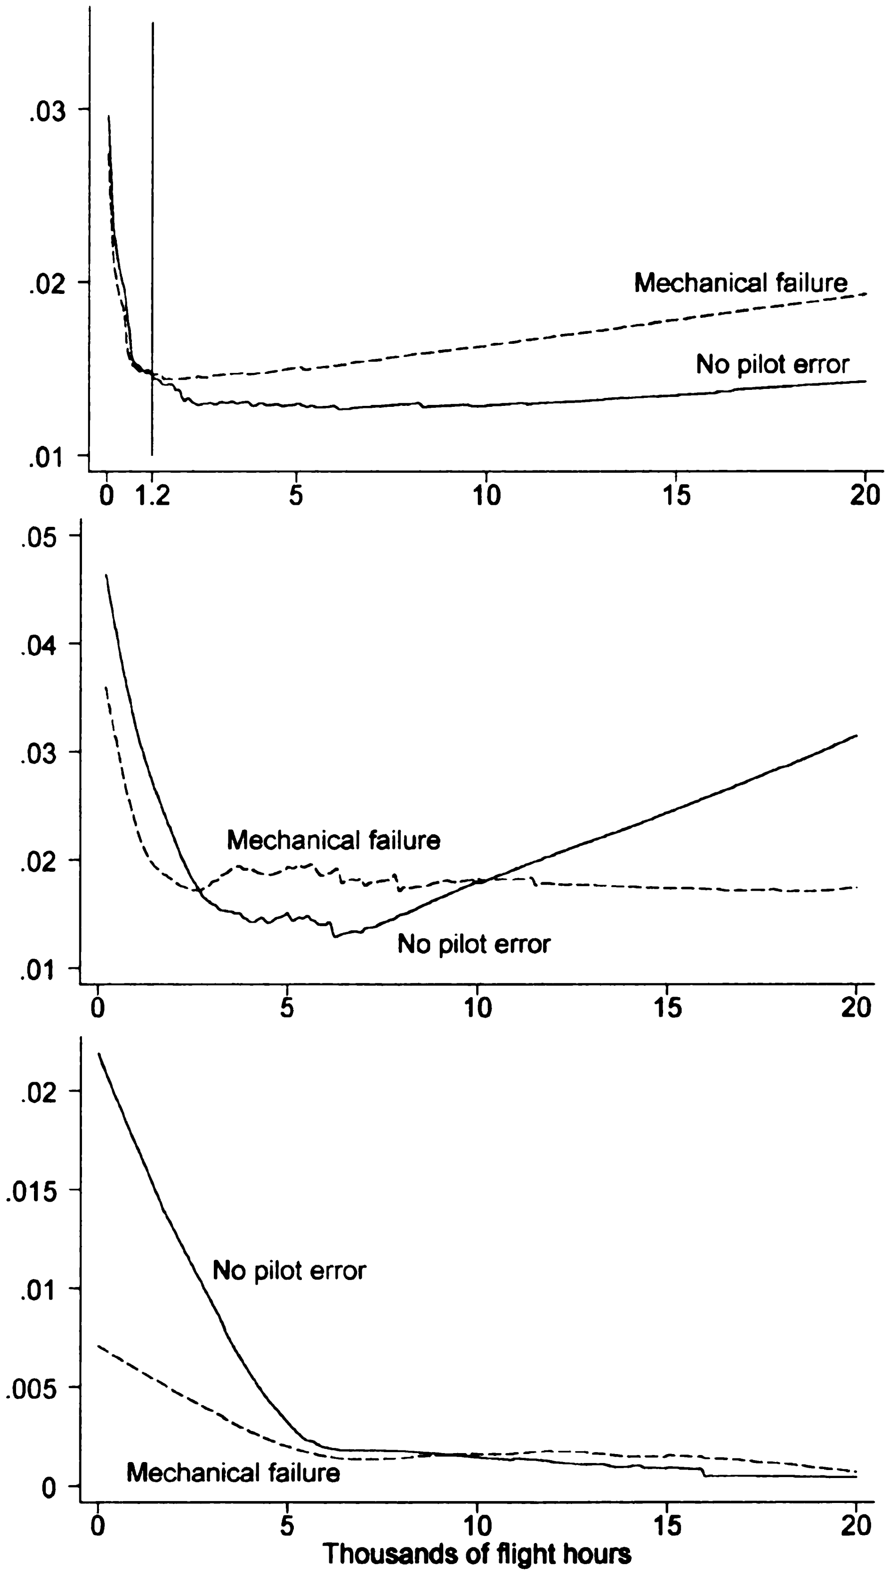

Figure 3.5: Non-parametric estimates of accident probability, general aviation (top), commuters (middle) and trunk carriers (bottom). LOWESS, bw=0.8

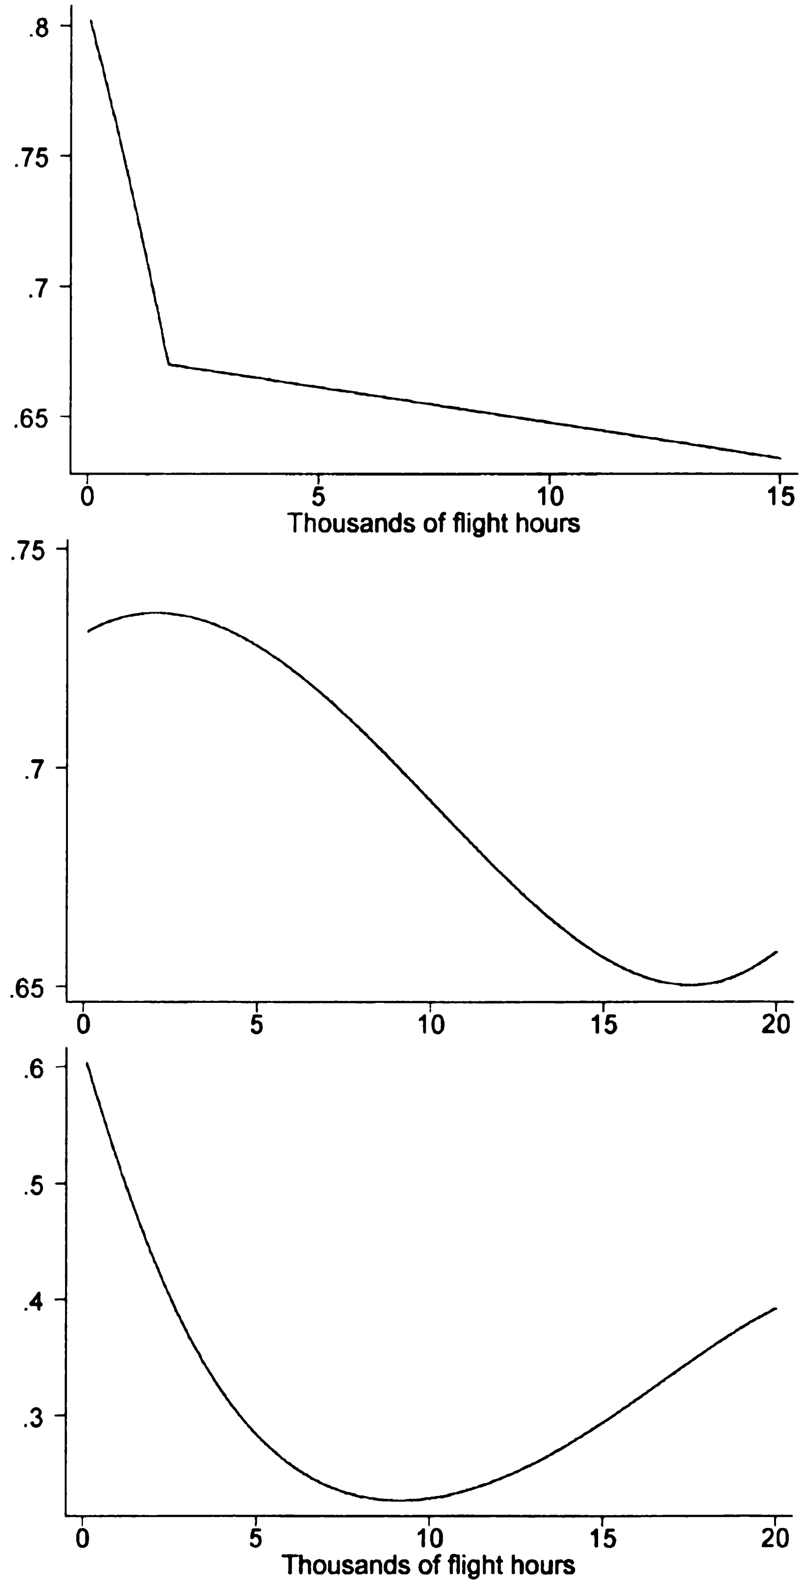

Figure 3.6: Predicted probability of pilot error from probit model, general aviation (top), commuters (middle) and trunk carriers (bottom)