Essays on European Labor Markets

by

Filipa Sá

Licenciatura, Economics, Universidade do Porto (2001)
MSc, Economics, The London School of Economics and Political Science (2002)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2008

Signature of Author .....
Department of Economics
January 12, 2008

Certified by .....
Joshua Angrist
Professor of Economics
Thesis Supervisor

Certified by .....
Olivier Blanchard
Class of 1941 Professor of Economics
Thesis Supervisor

Accepted by .....
Peter Temin
Elisha Gray II Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Essays on European Labor Markets

by

Filipa Sá

Submitted to the Department of Economics
on January 12, 2008, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

Chapter 1 examines whether immigrants gain a comparative advantage relative to natives in highly protected labor markets. This may be the case if immigrants, being new to the country, are less aware of employment protection regulations and less likely to claim their rights. I test this hypothesis drawing on evidence for the EU and on two natural experiments for Spain and Italy. The results suggest that stricter Employment Protection Legislation (EPL) does indeed benefit immigrants relative to natives. Stricter EPL is found to reduce employment and reduce hiring and firing rates for natives. By contrast, it has no effect on most immigrants and may even increase employment rates for those who have been in the country for a longer period.

Chapter 2 is the product of joint work with Marcello Estevão (IMF) and looks at the effect of the 35-hour workweek in France on wages, employment, dual job holdings and happiness. It explores the different timing of implementation of the shorter workweek in large and small firms to measure its causal effect. The results suggest that the reduction in hours did not succeed in increasing employment and generated a series of behavioural responses that are likely to have reduced welfare, as workers and firms tried to avoid the rigidities created by the reform. This suggests that the French government should increase the flexibility of workers and firms in setting hours of work.

Chapter 3 is the product of joint work with Olivier Blanchard (MIT) and Francesco Giavazzi (Universitá Commerciale Luigi Bocconi). Two main forces lie behind the large U.S. current account deficits: an increase in U.S. demand for foreign goods and an increase in foreign demand for U.S. assets. Both have contributed to steadily increasing current account deficits since the mid-1990s, accompanied by a real dollar appreciation until late 2001 and a real depreciation since. We develop a simple model of exchange rate and current account determination based on imperfect substitutability in goods and asset markets and use it to interpret the past and explore alternative future scenarios. We conclude that substantially more depreciation is to come against the yen, the renminbi, and the euro.

Thesis Supervisor: Joshua Angrist
Title: Professor of Economics

Thesis Supervisor: Olivier Blanchard
Title: Class of 1941 Professor of Economics

To my parents and to Paul

Acknowledgments

I am very grateful to my advisors, Josh Angrist and Olivier Blanchard, for their continuous and invaluable support. They have been true mentors and I cannot overstate my debt to them for guiding me at every step of the way. Josh, this dissertation owes very much to you. You helped me from the initial idea to the final draft and taught me principles about empirical research that shaped the way I approach economics. Olivier, you opened me many doors, giving me the opportunity to work with you as a research assistant and a co-author. Thank you for always pushing me forward and making me think harder.

Also at MIT, I would like to thank Francesco Giavazzi for his friendship and advice, and Gary King for helping me submit this thesis from London. I also thank my undergraduate students who brightened my time at MIT.

Three people played a crucial role in my graduate studies. Álvaro Aguiar, my professor at Universidade do Porto, encouraged me to go abroad and do the MSc at the LSE. John Sutton and Peter Diamond played a critical role in bringing me to MIT. I thank them for their generosity.

I gratefully acknowledge financial support from the MIT Shultz Fund and the Portuguese Science and Technology Foundation (through a PhD fellowship).

I owe much to my friends, in Portugal, Boston, Washington, and London. Sara has been my friend since childhood and I thank her for staying in touch during all these years. In Boston, I would like to thank Miriam, Olga, and Rita for making my time at MIT much more enjoyable. In Washington, I met many friends who made my summer there very happy and productive. In London, I thank my colleagues at the Bank of England for their support at the final stage of this dissertation.

I would also like to thank my extended family, my sister Anabela and her two children, Lena and Leo, for never forgetting me, even if far apart.

Finally, I have no words to thank my parents and Paul. My father has made my education his greatest priority and encouraged me to always do my best. My mother has been a constant source of support and was ready to cross the Atlantic every time I needed her in the last five years. Paul, thank you for your love and patience and for helping me to stay calm and focused while finishing this dissertation.

Contents

Chapter 1

Does Employment Protection Help Immigrants? Evidence from European Labor Markets

1.1 Introduction

This paper is concerned with two heavily discussed topics in the literature on European labor markets: immigrant assimilation (how well immigrants do in the labor market), and the effect of Employment Protection Legislation (EPL) on job flows and employment.

EPL includes several elements, such as notification procedures for dismissal, rules regarding the classification of dismissals as 'unfair', the amount of compensation following 'unfair' dismissals, etc. Depending on the reasons behind the firing, regulations in most European countries distinguish between redundancies (due to economic reasons, for example, poor sales, loss of competitiveness, economic recession, etc.) and disciplinary dismissals (due to misconduct, absenteeism, negligence, etc.). Redundancies are generally considered 'fair' dismissals and the

[1]Special thanks goes to my advisors, Josh Angrist and Olivier Blanchard, for their advice and encouragement. For helpful discussions and insightful comments, I thank Paul Bedford, Adriana Kugler, and Giovanni Pica. The MIT Shultz Fund and the Portuguese Science and Technology Foundation (through a PhD fellowship) provided financial support.

worker is entitled to a severance payment. Disciplinary dismissals do not imply any severance payment, but the worker may take the case to the labor authority if he believes the dismissal was 'unfair'. In this situation, one outcome could be compensation through conciliation, where the firm and the worker meet with the labor authority in order to reach a settlement. When conciliation fails, the case goes to the labor courts. If the court decision favors the worker and the judge declares the dismissal 'unfair', then severance payments are increased.

Because EPL increases firms' firing costs, it may impact on job flows and the level of employment. There has been an extensive amount of research exploring these relations. Theoretical models, such as the flow model with endogenous job destruction of Pissarides (2000), show that an increase in employment protection reduces both job creation and job destruction and has an ambiguous effect on employment.

The purpose of this paper is to examine the differential impact of EPL on the labor market outcomes of natives and immigrants, in particular to ask whether immigrants gain a comparative advantage relative to natives in highly regulated labor markets. This could be the case if immigrants are less covered by EPL than natives. Because the analysis here focuses solely on legal immigrants, there is no difference between the legal status of natives and immigrants, so the same employment protection regulations apply to both. However, being new to the country, immigrants may be less informed about their rights and may, for example, not seek to claim compensation for 'unfair' dismissal. Therefore, while the *legal* coverage of EPL is the same for natives and immigrants, the *effective* coverage may be lower for immigrants.

One reason why immigrants may be less informed about their rights is if they are relatively less likely to join unions than natives. Figure 1 reports union membership rates by citizenship for a number of European countries, the US and Canada, taken from the 2003 International Social Survey Programme (ISSP). Union membership varies widely across Europe, and is typically highest in northern Europe. In all countries, natives have higher union membership rates than immigrants. For the US, these numbers can be compared with data from the Current Population Survey (CPS), reported in MPI (2004). In the CPS, 12.9% natives and 10.2% immigrants were union members in 2003. The evidence on union membership supports the hypothesis that there may be an information gap between immigrants and natives regarding labor market regulations.

To test whether immigrants gain a comparative advantage in heavily regulated labor mar-

kets, this paper draws on three pieces of evidence. The first is the EU Labour Force Survey (LFS), which covers 15 European countries from the mid-1990s to 2005.The other two sources of evidence are natural experiments, based on labor market reforms in Spain and Italy. The Spanish reforms were enacted in 1997 and 2001 and introduced a new type of permanent contract with lower dismissal costs for some categories of workers, namely unemployed people between the ages of 16 and 30 and over age 45. The Italian reform was enacted in 1990 and increased dismissal costs for permanent contracts in firms with under 15 employees. The heterogeneity in the application of both these reforms allows the design of natural experiments that can be used to examine the effect of changes in dismissal costs on employment of natives and immigrants.

All these pieces of evidence tell a consistent story and confirm the hypothesis that immigrants obtain a comparative advantage in heavily regulated labor markets. The evidence from the EU LFS shows that EPL (measured by a time-varying index constructed from the OECD synthetic indices) reduces employment for natives but not for immigrants. The results suggest that EPL may even increase employment for immigrants who have been in the country for a longer time. The natural experiment for Spain shows that a reduction in dismissal costs increased permanent employment for natives but not for immigrants. Finally, the evidence for Italy shows that an increase in dismissal costs decreased both hiring and firing, with a larger effect on natives than on immigrants.

The empirical framework used in these three applications follows closely the methodology in Angrist and Kugler (2003), Kugler et al (2002), and Kugler and Pica (2005). Angrist and Kugler (2003) examine how labor and product market rigidities affect the impact of immigration on the employment of natives. Using data from the EU LFS, they find that immigration generates larger job losses for natives in countries with more restrictive regulations. Kugler et al (2002) and Kugler and Pica (2005) analyze the 1997 Spanish reform and the 1990 Italian reform using the same data sources and research design as this paper. However, they do not distinguish between natives and immigrants. This paper extends their analysis by looking at the differential impact of EPL on these two groups of workers.

The paper is structured as follows. Section 1.2 presents a simple model of the effect of EPL on the labor market. Section 1.3 briefly reviews some key empirical studies. Section 1.4 presents evidence based on the EU LFS. Section 1.5 presents the results of a natural experiment

for the 1997 and 2001 Spanish reforms. Section 1.6 looks at the 1990 Italian reform. Section 1.7 concludes.

1.2 Theoretical Framework

To analyze the effect of EPL on the labor market outcomes of natives and immigrants, a simple model is constructed, drawing on the flow model with endogenous job destruction of Pissarides (2000). The presentation here combines elements of the models in Blanchard and Portugal (2001) and Cahuc and Zylberberg (2004), but extends them to consider two types of workers: natives and immigrants. In the simplest version of the model, a particular job is either always filled by natives or always filled by immigrants, i.e., there is no substitution between the two types of workers on the labor demand side. The model can then be extended such that a particular job can be filled by either natives or immigrants over time. In this version, firms observe the worker's type and can choose freely between hiring a native or an immigrant.

1.2.1 Assumptions

There are jobs and workers in the economy. Workers are of two types: natives and immigrants. A match between a job and a worker may be destroyed due to a productivity shock. When a productivity shock occurs, firms have the option to either fire or retain the worker. Employment protection takes the form of a firing cost that the firm must pay if it decides to fire the worker. For simplicity, it assumed that natives and immigrants are in all ways identical, except in the level of the firing cost, which is lower for immigrants. This difference in firing costs has implications for wages and the hiring and firing rates of the two types of workers.

More specifically, the model has the following ingredients:

1. A match between a job and a worker yields productivity ε, which is a random variable with support $]-\infty, \bar{\varepsilon}]$ and cumulative distribution function $G(.)$.

2. Productivity shocks occur with probability λ. When a shock occurs, a new level of productivity is drawn from the distribution $G(.)$ and firms have the option to either fire or retain the worker. Productivity shocks are idiosyncratic, i.e., they affect jobs independently.

3. If the firm fires the worker following a productivity shock, it must pay a firing cost. This cost should not be interpreted as a transfer from the firm to the worker, but rather as pure waste (for example, administrative costs such as notice periods, possible legal procedures, necessary authorizations for dismissal, etc.).[1]

4. The firing cost determines the critical level of productivity, ε_R, below which the firm fires the worker.

5. There are two types of workers in the economy: natives and immigrants. They are in all ways identical, except in the level of the firing cost, which equals f for natives and ϕf, where $0 \leq \phi < 1$, for immigrants. The difference in firing costs between the two groups is not due to differences in their legal coverage (since the same employment protection rules apply to both), but rather to differences in the level of awareness regarding regulations. Immigrants, being new to the labor market and less likely to join unions, are less aware of the employment protection regulations and less likely to claim compensation for dismissal. Therefore, even though *legal* coverage is the same for the two types of workers, *effective* coverage is lower for immigrants.

6. The model is characterized by 'workers waiting at the gate', so that, when a match is destroyed following a productivity shock, firms can hire another worker immediately. All new hires start with productivity ε_s.

7. Employed workers receive a wage w. Wages for each match are determined by symmetric Nash bargaining between the firm and the worker. For simplicity, it is assumed that the wage is set at the beginning of the match.

8. Unemployed workers are entitled to unemployment benefits z and have a probability h of finding a job. Because job creation and job destruction follow Poisson processes, the inverse of h equals the duration of unemployment.

[1] The distinction between interpreting EPL as a transfer from firms to workers or pure waste is important. Lazear (1990) shows that, if employment protection was just a transfer from firms to workers and capital markets were perfect, then the effects of EPL could be undone by a contract including a compensation from the worker to the firm at the time of hiring.

1.2.2 No Substitution on the Labor Demand Side

In this version of the model, a particular job is either always filled by natives (N) or always filled by immigrants (I), i.e., there is no substitution between the two types of workers on the demand side. The equilibrium for each type is characterized by the critical level of productivity ε_R^i, the wage w^i, and the job finding rate h^i, where $i = N, I$. The equilibrium will be derived for natives. Because immigrants are in all ways identical to natives, except for the firing cost, the equilibrium equations for them will be identical to those for natives, except that f should be replaced with ϕf.

It is useful to start by writing the value of a filled vacancy for the firm and a native worker. A job filled with a native worker with current productivity ε gives an expected profit $\Pi_e^N(\varepsilon)$, defined as:

$$r\Pi_e^N(\varepsilon) = (\varepsilon - w^N) + \lambda G(\varepsilon_R^N)[\Pi_e^N(\varepsilon_s) - \Pi_e^N(\varepsilon) - f] + \lambda \int_{\varepsilon_R}^{\bar{\varepsilon}} (\Pi_e^N(\varepsilon') - \Pi_e^N(\varepsilon)) dG(\varepsilon') \quad (1.1)$$

In this equation $(\varepsilon - w^N)$ is the current profit from the match. The second term captures the expected profit if there is a productivity shock and the firm fires the worker, an outcome that occurs with probability $\lambda G(\varepsilon_R^N)$. In that case, the firm loses the current profit and must pay the employment protection cost f. The job is filled with another native worker, with productivity ε_s. The third term captures the expected profit if there is a productivity shock and the firm keeps the worker. In this case, a new level of productivity is drawn from the distribution $G(.)$.

The expected utility of the native worker is defined as:

$$rV_e^N = w^N + \lambda G(\varepsilon_R^N)(V_u^N - V_e^N) \quad (1.2)$$

The second term captures the utility if the worker is fired following a productivity shock. In that case, he becomes unemployed and receives the expected utility of unemployment, V_u^N, which is given by:

$$rV_u^N = z + h^N[V_e^N - V_u^N] \quad (1.3)$$

These three equations can be used to solve for the equilibrium, which is characterized by the critical level of productivity ε_R^N, the wage w^N, and the hiring rate h^N.

(i) **Critical level of productivity ε_R^i**

The firm fires the worker if the value of the match falls below the value of a new match net of employment protection costs. Therefore, the critical level of productivity below which a native worker is fired satisfies:

$$\Pi_e^N(\varepsilon_R^N) = \Pi_e^N(\varepsilon_s) - f$$

Using equation (1.1) and solving:

$$\varepsilon_R^N = \varepsilon_s - f(\lambda + r)$$

This equation describes the effect of an increase in employment protection on the firing rate and gives the first important prediction of the model: if the firing cost increases, the critical level of productivity below which the firm fires the worker decreases. Therefore, *an increase in employment protection reduces the firing rate.*

The equivalent equilibrium equation for immigrants is:

$$\varepsilon_R^I = \varepsilon_s - \phi f(\lambda + r)$$

Because $0 \leq \phi < 1$, it follows that $\varepsilon_R^I > \varepsilon_R^N$. This implies that the firing rate (given by $\lambda G(\varepsilon_R^i)$) is higher for immigrants than for natives. It also implies that *the impact of an increase in employment protection on the firing rate is lower for immigrants than for natives*, since they are effectively less covered by EPL.

(ii) **Wage**

Symmetric Nash bargaining implies that the surplus of the job for the worker equals the surplus for the firm. For jobs filled with native workers, this implies:

$$\begin{aligned} V_e^N - V_u^N &= \Pi_e^N(\varepsilon_s) - (\Pi_e^N(\varepsilon_s) - f) \Rightarrow \\ V_e^N - V_u^N &= f \end{aligned}$$

The left hand side is the value of the job for a native worker. The alternative to being employed is to be unemployed and have expected utility V_u^N. The right hand side is the value of the job for the firm. If the firm fires the worker, it finds a replacement worker, but must pay the cost f.

This expression can be combined with equations (1.2) and (1.3) to solve for the wage paid to natives:

$$w^N = z + (r + \lambda G(\varepsilon_R^N) + h^N)f$$

The equivalent equilibrium equation for the wage paid to immigrants is:

$$w^I = z + (r + \lambda G(\varepsilon_R^I) + h^I)\phi f$$

For a given critical level of productivity ε_R^i and a given job finding rate h^i, *an increase in employment protection increases the bargained wage.* This result is intuitive: employment protection increases the workers' bargaining power as it is more costly for the firm to replace the worker. This effect is weaker for immigrants than for natives as they are effectively less covered by EPL.

(iii) **Hiring rate** h^i

Employment protection increases the cost of labor in two ways: by increasing the firing cost, and by increasing the bargained wage. Firms respond to the increase in costs by hiring fewer workers. So, *an increase in employment protection reduces the hiring rate and increases the duration of unemployment.* Because the effect on firing costs and on wages is larger for natives, *an increase in EPL reduces the hiring rate by more for natives than for immigrants.*

An implication of these predictions is that *the effect of employment protection on the unemployment rate of natives and immigrants is ambiguous.* The unemployment rate equals the

product of the firing rate and unemployment duration. Because employment protection reduces the firing rate and increases unemployment duration, its effect on the unemployment rate is theoretically undetermined and is ultimately an empirical question.

1.2.3 Substitution on the Labor Demand Side

In the model with no substitution on the labor demand side it is assumed that a particular job is either always filled by natives or always filled by immigrants. In practice, however, a particular job can be filled by either natives or immigrants over time. Therefore, it is important to introduce substitution between the two types of workers on the demand side.

The nature of the equilibrium is quite intuitive. Because immigrants have a lower firing cost than natives and are identical to them in all other respects, they are *a priori* more attractive to firms. Firms will only hire natives if their expected cost is the same as that of immigrants, i.e., if their relative wage decreases in order to compensate for their higher expected firing cost. If natives are to accept a lower relative wage in the bargaining process, unemployment must be less attractive to them. Therefore, equilibrium unemployment duration will be higher for natives than for immigrants, i.e., they will have a relatively lower hiring rate.

More formally, the equilibrium can be characterized as follows:

(i) Relative firing rates

The firing rate of each type is determined by the condition that the value of the match equals the value of a new match net of employment protection costs. For natives, this can be written as:

$$\Pi_e(\varepsilon_R^N) = \Pi_e(\varepsilon_s) - f$$

The equivalent condition for immigrants is:

$$\Pi_e(\varepsilon_R^I) = \Pi_e(\varepsilon_s) - \phi f$$

Note that in this version of the model there is only one profit function $\Pi_e(\varepsilon)$ rather than two separate functions for each type, since now the same job may be filled by either natives or

immigrants over time.

Because $0 \leq \phi < 1$ and both types have the same initial productivity ε_s, it follows that $\Pi_e(\varepsilon_R^N) < \Pi_e(\varepsilon_R^I)$. Given that the profit function is increasing in the level of productivity, this implies that $\varepsilon_R^N < \varepsilon_R^I$, i.e., immigrants have a relatively higher firing rate than for natives. Moreover, *an increase in employment protection reduces the firing rate of natives by a relatively larger amount*, since they are effectively more covered by EPL than immigrants.

(ii) Relative wages

Using equation (1.1) and the fact that the two types of workers have the same initial productivity ε_s and the same productivity distribution $G(.)$, it follows that the condition for firms to be willing to hire natives is:

$$w^N - w^I = \lambda G(\varepsilon_R^I)\phi f - \lambda G(\varepsilon_R^N) f < 0$$

This condition says that firms are indifferent between hiring natives and immigrants if the relative wage of natives decreases as to equate the difference in expected firing costs. This implies a lower relative wage for natives than for immigrants, which is probably at odds with reality in most countries.

It is important to emphasize that this is a *ceteris paribus* result, i.e., for simplicity, the model assumes that immigrants and natives are in all ways identical, except for the level of the firing cost. If natives were assumed to have a higher starting productivity, ε_s, and a productivity distribution, $G(.)$, with a higher support than immigrants, then the difference in wages would be equal to the difference in current and expected future productivities minus the difference in expected firing costs. With those assumptions, it could easily be the case that natives would have a higher relative wage than immigrants.

Regardless of what is assumed in terms of relative productivities, an increase in f increases the relative firing cost of natives and makes them relatively less attractive to firms. Therefore, *an increase in employment protection reduces the equilibrium wage of natives relative to immigrants.*

(iii) Relative hiring rates

For natives to be willing to accept a lower relative wage in the bargaining process following an increase in f, their labor market conditions must deteriorate. Therefore, *an increase in*

employment protection increases relative unemployment duration for natives, i.e., it decreases their relative hiring rate.

To close this section it is important to note that, just as in the model with no substitution between the two types of workers, *the effect of an increase in employment protection on the relative unemployment rates of natives and immigrants is ambiguous.* On the one hand, EPL reduces the relative firing rate of natives. On the other hand, it increases relative unemployment duration. Because the unemployment rate equals the product of the firing rate and unemployment duration, its response to an increase in employment protection is ambiguous.

1.3 Results from Empirical Studies

The predictions of the flow model of unemployment have been tested in a number empirical studies, using both macro and micro data. Overall, the macro evidence is inconclusive and the results are very sensitive to the specification used. Evidence based on micro data is more conclusive.

A key reference in the macro literature is the study by Lazear (1990), which looks at the effect of employment protection on the level of employment in 22 OECD countries. Using data from 1956 to 1984, he finds that an increase in EPL reduces employment. Similar results were obtained by Scarpetta (1996) using data for 15 to 17 OECD countries over the period 1983 to 1993. But studies based on other specifications find different results. Bertola (1990) finds that employment protection does not reduce employment for a given wage and does not increase bargained wages. Nickell (1997) investigates the relation between unemployment and the OECD index of employment protection, using data for 20 OECD countries over the periods 1983 to 1988 and 1989 to 1994. He also does not find any significant effect of employment protection on unemployment.

Studies based on macro data suffer from two main problems. First, it is difficult to establish the direction of causality, as it is not clear whether employment protection reduces employment or whether countries with low employment adopt stricter labor market regulations. Second, there is little time-series variation in measures of employment protection, which makes it hard to identify any effect. Reflecting these concerns, a more promising avenue of research are studies

with micro data, in particular those that analyze the effects of particular reforms that change the regulation on employment protection. When there is heterogeneity in the applicability of the reforms, it is possible to construct natural experiments and identify the causal effect of employment protection on employment and job flows.

Two key references in the micro literature are Kugler et al (2002) and Kugler and Pica (2005). Kugler et al (2002) construct a natural experiment to analyze the impact of the 1997 reform in Spain, which introduced a new type of permanent contract, with lower dismissal costs, applicable to certain categories of workers (namely unemployed people between the ages of 16 and 30 and over age 45). They find that, for young workers, the reform increased transitions both from unemployment to permanent employment and from temporary to permanent employment and had little effect on dismissals. As a result, permanent employment increased for young workers following the reform. The results for old workers are less significant. Kugler and Pica (2005) also use a natural experiment to analyze the effect of the 1990 reform in Italy, which increased dismissal costs for firms with under 15 employees. They find that the increase in dismissal costs decreased hiring and firing in small firms relative to large firms.

This paper extends the analysis of Kugler et al (2002) and Kugler and Pica (2005) to examine the differential impact of EPL on immigrants and natives. The study which is closest in spirit to this paper, exploring the interaction between immigration and institutional factors, is Angrist and Kugler (2003). The authors look at how labor and product market rigidities affect the impact of immigration on natives and find that immigration generates larger native job losses in countries with more restrictive institutions. This is consistent with the view that immigrants gain a comparative advantage relative to natives in highly regulated labor markets.

1.4 Evidence from the EU LFS

1.4.1 Data and Descriptive Statistics

The EU data were extracted from the Eurostat Labor Force Survey (LFS) by the German Federal Statistics Office. The dataset comprises time series of immigration and labor market variables aggregated by age, gender, education, country of birth, and, for immigrants, years since arrival. The Data Appendix contains more information about the data.

The countries and years included in the sample are listed in Table 1. Data for most countries cover the period 1995 to 2005. To ensure that the coverage of immigrants in the data is appropriate, the percentage of foreign born in the population in the LFS is compared with similar numbers from the last Census year, reported in OECD(2006). A perfect coincidence between the two series should not be expected because the LFS covers only the population above age 15, whereas the Census covers the entire population. Even so, the two series are remarkably similar, suggesting a good coverage of immigrants in the LFS data.

The definition of immigrant adopted here for the EU LFS analysis is based on nativity, i.e., country of birth. This definition is chosen because the variable years of residence, used as a control in the empirical analysis, is based on nativity. An alternative would be to focus on nationality. Angrist and Kugler (2001) briefly explore the implications of defining immigrants based on nativity or nationality and find that the groups of immigrants roughly coincide.

Some descriptive statistics are reported in Table 2. Participation rates show no clear pattern across countries, but immigrants tend to have lower participation rates than natives, especially the ones who have been in the country for a shorter period of time. Employment-to-population ratios exhibit a similar pattern. Turning to the share of total employment in permanent and temporary contracts, Spain is the country where temporary contracts are most common, accounting for over 20% of employment for natives and over 50% for recent male immigrants. France, Germany, Sweden, Finland, Switzerland, and Portugal also have large shares of workers in temporary contracts. The incidence of temporary employment is larger for women than for men and for immigrants than for natives, declining with years since immigration.

1.4.2 Measures of Employment Protection

Time series data on EPL for permanent and temporary contracts for the period 1985 to 2005 were created using OECD summary indicators of the strictness of EPL in different countries and information on the breaking points of the indicators. The Data Appendix explains how the data were constructed.

Other papers have constructed similar measures. Blanchard and Wolfers (2000) draw on data from Lazear (1990) and from the OECD indicators from the late 1980s and late 1990s to construct a 5-year period index of employment protection from 1960 to 1995. This paper

follows a similar approach, but with three main differences. First, the availability of OECD indicators for 2003 allows the construction of a longer time series. Second, it looks separately at EPL on permanent and temporary contracts, rather than the aggregation of the two. Third, rather than constructing 5-year period indices, this paper explores information on the timing of reforms to adjust the OECD indicators at the exact year when reforms took place.

The evolution of the EPL indices for permanent and temporary contracts in Germany, Spain, France, the UK, and Italy is described in Figures 2 and 3. There is considerable variation in EPL across countries. The UK is the least regulated, whereas Spain and France are the most restrictive. However, there is not much time variation in the indices. Some countries, such as Italy and Germany, have introduced reforms at the margin, increasing the flexibility of temporary contracts, with the introduction and generalization of temporary work agencies, while leaving regulations on permanent contracts virtually unchanged. There are very few examples of EPL reforms for permanent contracts. Spain is an exception, having introduced reforms in 1994 and 1997. The Spanish case will be discussed in the next section. The 1990 Italian reform discussed in section 1.6, which increased 'unfair' dismissal costs for firms with under 15 employees, was not taken into account in the construction of the OECD series.

1.4.3 Estimates of the Effects of Employment Protection

The flow model of unemployment presented in section 1.2 predicts that EPL increases hiring and firing rates and has an ambiguous effect on the unemployment rate. Before turning to a more rigorous analysis of the effects of EPL, it is useful to look at some simple scatter diagrams. Figures 4, 5, and 6 plot the employment protection index for permanent jobs against unemployment duration, employment duration, and the employment rate, for natives and immigrants. The observations are averages across all countries and years in the sample. If we view job creation and job destruction as Poisson processes, unemployment duration is the inverse of the probability of being hired and employment duration is the inverse of the probability of being fired. Therefore, unemployment and employment durations provide information on the hiring and firing rates, respectively.

Figures 4 and 5 suggest that there is a positive relation between EPL and unemployment and employment durations for natives, but not for immigrants. This is consistent with the

predictions of the model: EPL reduces hiring and firing, but the effect is smaller for immigrants, since they are effectively less covered by EPL. Figure 6 suggests that there is no correlation between EPL and the employment rate, which is also consistent with the results of the model.

Although suggestive, these scatter diagrams do not explore the time series, cross-country, and cross-demographic group variation in the data. To do a more formal analysis of the effects of EPL on labor market outcomes, the following general specification is used:

$$y_{ijt} = \alpha EPL_{it} + \delta_j(Years\ residence_j * EPL_{it}) + \beta_i + \gamma_t + \mu X_{ijt} + \varepsilon_{ijt} \tag{1.4}$$

where y_{ijt} is the outcome of interest for country i, year t, and demographic group j (where j denotes gender, age, education and nativity groups as defined in Appendix A.1). EPL_{it} is the EPL index on permanent or temporary contracts and is measured in deviations from cross country means to make the effect interpretable as the effect 'at the mean'. $Years\ residence_j$ is a set of dummy variables, where the omitted variable equals 1 for natives and 0 for immigrants, and the remaining three dummies classify immigrants into different groups depending on years since arrival (5 years or less, 6 to 10 years, and more than 10 years). β_i are country fixed effects and γ_t are year fixed effects. X_{ijt} is a vector of other controls, which includes dummies for gender, age, education, and nativity groups. It also includes interactions of all regressors (including country and year dummies) with the dummies $Years\ residence_j$, so that α can be interpreted as the effect of EPL on native workers and δ_j as the differential effect of EPL on immigrants relative to natives, by years of residence.

To explore the channel through which EPL may impact differently on natives and immigrants, X_{ijt} includes interactions of the education dummies with EPL_{it}. The reason for including these interactions is that immigrants may be less educated than natives and may, therefore, work in less protected jobs. As a result, they would be less affected by EPL due to their lower education and not because of other differences. The inclusion of these controls allows the two effects to be separated.

In this specification, identification comes not only from the time series variation in EPL, which is very limited, but also from the cross-country and cross-demographic group variation. Standard errors are clustered on $country * year$ to control for common errors within these groups. Failure to take this into account would lead to underestimation of the standard errors

and overestimation of the effect of interest, as described in Moulton (1990).

Participation Rates

Before turning to the effect of EPL on employment, equation (1.4) is estimated with the labor force participation rate as the dependent variable. In principle, the effect of EPL on participation is undetermined. On the one hand, EPL reduces the hiring rate, making it harder for an unemployed worker to find a job. The worker may get discouraged and leave the labor force. On the other hand, EPL also reduces the firing rate, increasing the value of having a job and encouraging participation.

The results are reported in the first two columns of Table 3. There is no evidence that EPL on permanent contracts has any effect on participation. The controls have the expected signs: immigrants have higher participation rates than natives, men have higher participation rates than women and participation increases with age and education. EPL on temporary contracts has a very small effect on participation rates.

Employment Rates

Theory does not make clear predictions regarding the effect of EPL on employment and, as discussed above, the empirical evidence on this effect based on macro studies is mixed. The last two columns of Table 3 report the results of estimating equation (1.4) to study the impact of EPL on the employment rate. Interestingly, higher EPL for permanent contracts unambiguously decreases the employment rate of natives. An increase in the EPL index for permanent contracts by one unit relative to the cross country mean reduces the employment rate of natives by approximately four percentage points. However, the effect is positive for immigrants who have been in the country for at least six years and increases on years since arrival. This suggests that higher EPL on permanent contracts may actually improve the job prospects of immigrants. By contrast, EPL on temporary contracts has a small positive effect for natives and no significant effect for immigrants.

To investigate whether the differential effect of EPL on immigrants and natives is due to differences in workers' characteristics, the estimation includes interactions between the education dummies and EPL. The coefficients on these interactions show that EPL has a larger negative

effect on employment for more educated workers. This is not surprising since it is plausible that more educated workers are more aware of their rights under EPL. Even including these controls, the results suggest that EPL on permanent contracts reduces employment for natives and increases it for immigrants with more than six years of residence.

There may be other differences between immigrants and natives, in addition to education, which explain the differential impact of EPL. For example, it is possible that immigrants, being new to the country, are less informed about the legislation (perhaps because they are less unionized than natives) and, therefore, less likely to seek compensation in case of dismissal. The data on unionization rates in Figure 1 support this explanation. Another piece of evidence that would be relevant to further support this hypothesis would be data on the number of complaints or the amount of severance payments received by natives and immigrants. Unfortunately, this information is not available disaggregated by nativity.

The notion that the differential effect of EPL on natives and immigrants is due to an information problem seems to be at odds with the evidence in Table 3 that the effect of EPL on employment becomes positive for immigrants who have been in the country for more than six years and increases on years since arrival. It seems plausible that immigrants who have been in the country for a longer time are more integrated, have a better knowledge of the language, and are also more aware of the labor market regulations than those who have just arrived. Therefore, one would expect that, if the comparative advantage of immigrants was based on an information problem, recent immigrants would be the ones to benefit the most.

However, an increase in years since arrival has two counteracting effects. On the one hand, recent immigrants may be less informed of the legislation and be effectively less covered by EPL. On the other hand, they may also be less attractive to employers, either because they are less productive (for example, because they have a poor knowledge of the language) or are perceived as being less productive or less reliable (since they do not have a history of employment in the country). If the second effect dominates, then it is plausible that an increase in EPL induces firms to hire more immigrants of older arrival cohorts.

Employment and Unemployment Duration

The effect of EPL on the employment rate is the result of two types of flows: flows from employment into unemployment, and flows from unemployment into employment. Both theory and the micro evidence suggest that EPL decreases these flows. On the one hand, by increasing the cost of firing, EPL reduces flows into unemployment. On the other hand, by increasing the costs for firms and the bargaining power of workers, EPL reduces flows into employment. The EU LFS collects no information on job flows, but it contains information on unemployment and employment durations. Since unemployment duration is the inverse of the hiring rate and employment duration is the inverse of the firing rate, this information is sufficient to analyze the effect of EPL on job flows.

The results of estimating equation (1.4) with unemployment and employment duration as dependent variables are shown in Table 4. An increase in EPL for permanent contracts increases unemployment duration, i.e., decreases the hiring rate, for natives. The effect on immigrants is much weaker, and even reverses sign for immigrants with six to ten years of residence. These results are consistent with the model in section 1.2, which predicts that an increase in EPL reduces the relative hiring rate of natives. By contrast, EPL for temporary contracts has no significant effect on the probability of being unemployed for more than one year.

Turning to the effects on employment duration, the results suggest that EPL on permanent contracts increases the probability of being employed for more than five years by approximately the same amount for natives and for immigrants. This result does not match the model in section 1.2, which predicts that an increase in EPL should reduce the relative firing rate of natives. When looking at the effect of EPL on temporary contracts, the results in column (2) show no significant effect on the probability of being employed for more than five years. This is not surprising since this type of EPL works on contracts of short duration. Indeed, as the last column suggests, EPL on temporary contracts does have a significant positive effect on the probability of being employed for more than one year.

The results on employment and unemployment duration are consistent with the findings for the employment rate. EPL on permanent contracts decreases hiring and firing. For natives, the reduction in hiring has a larger magnitude than the reduction in firing, resulting in a negative effect on the employment rate. For immigrants, EPL on permanent contracts has a

much smaller effect on hiring and about the same effect of firing. The combination of these two effects results in a positive effect of permanent EPL on the employment rate of immigrants with more than six years of residence.

As for the employment rate, the differential effect of EPL on employment and unemployment duration across nativity groups cannot be fully explained by differences in education. Even controlling for interactions between the EPL index and the education dummies, the differential effect of EPL on natives and immigrants persists. Again, this suggests that there may be other important differences between the two groups, such as the fact that immigrants are new to the country, less aware of the legislation and less unionized than natives and, therefore, less likely to seek compensation in case of dismissal.

1.5 Evidence from the 1997 and 2001 Spanish Reforms

1.5.1 Institutional Background

In 1984 Spain liberalized the use of temporary contracts which previously could only be used for seasonal jobs. This reform generalized their use to all workers and sectors, resulting in a rapid increase in temporary employment and creating a dual labor market. Spain became the country with the highest proportion of workers with temporary contracts in the EU, as reported in Table 2.

Subsequent governments implemented reforms to mitigate the duality in the Spanish labor market. In 1994, a reform restored the rule that temporary contracts could only be used for seasonal jobs and eliminated the most flexible type of temporary contract: the fixed term contract to promote employment. This reform also acted upon permanent contracts, relaxing their dismissal conditions. However, the 1994 reform is regarded as having been largely ineffective. In practice, firms continued to hire workers under temporary contracts for all types of jobs.

In 1997, again with the intention of reducing duality in the labor market, the newly elected conservative government of José María Aznar implemented a further reform. Instead of reducing dismissal costs for permanent contracts across the board, this reform introduced a new type of permanent contract - 'contract for promoting permanent employment' - with lower levels of compensation for 'unfair' dismissals. While the level of compensation of standard permanent

contracts equals 45 days for each year of service, up to a maximum of 42 monthly payments, the new contract has a level of compensation equal to 33 days for each year of service, up to a maximum of 24 monthly payments.

The new type of permanent contract applied only to certain categories of workers, namely the young unemployed (age 16 to 30), the unemployed for more than one year, the unemployed over age 45, disabled workers, and workers previously hired under temporary contracts. In order to encourage firms to hire under the new permanent contract, the law reduced payroll taxes for newly signed contracts and for conversions from temporary to permanent employment using the new contract. The payroll tax reductions ranged from 40% to 90% and lasted for two years, except for contracts for the unemployed over age 45, where the reductions lasted for the duration of the contract.

The 1997 reform was initially valid for four years. In 2001, the Spanish government decided to adopt the new type of contract on a permanent basis and extended its scope, including new target groups, such as women hired in sectors were they are under-represented, unemployed women hired in the period of 24 months after childbirth, and people earning integration incomes. There was also another reform in 2001 which introduced dismissal costs for temporary contracts, equal to 8 days per year worked. This amount is not significant, but the purpose was again to approximate the costs of temporary and permanent contracts in order to reduce duality in the labor market.

Table 5, reproduced from Kugler et al (2002), lists the reductions in dismissal costs and in payroll taxes introduced by the 1997 and 2001 reforms for different categories of workers. The 1997 reform applied to contracts signed in 1997 and 1998. The 2001 reform extended the reduction in dismissal costs and in payroll taxes to contracts signed in or after 1999. The table shows that the extent of the incentives varies by demographic group, with unemployed workers age 30 to 44 unaffected.

1.5.2 Identification Strategy

The variability in the application of the 1997 and 2001 reforms to different demographic groups allows the design of a natural experiment to study the effect of the reduction in dismissal costs on the Spanish labor market. There are different margins of variability that could be explored:

comparing workers in different age groups, comparing disabled to non-disabled, women under-represented in their occupation to those not under-represented, the long-term unemployed to other unemployed, etc.

Following Kugler et al (2002), this paper explores the variability in age and compares labor market outcomes of workers age 30 to 44 (control group) to workers age 16 to 29 (treatment group 1) and workers age 45 to 59 (treatment group 2) before and after the reform. Workers above age 59 are excluded in order to isolate the effect of the reform of dismissal costs and avoid capturing the effect of simultaneous reforms to the pension system[2]. The advantage of defining groups based on age is that, unlike some other classifications (such as women under-represented in their occupation and the long-term unemployed) self-selection is not a problem.

To estimate the effect of the reform on the employment probability, the following logit specification is used:

$$\begin{aligned} e_{it} &= \Lambda[\alpha X_{it} + \gamma_0 d_t + \gamma_1 q_t + \beta_1 treatment_{1i} + \beta_2 treatment_{2i} + \\ &\quad \delta_1(treatment_{1i} \times post_t) + \delta_2(treatment_{2i} \times post_t)] + \varepsilon_{it} \end{aligned} \tag{1.5}$$

The dependent variable e_{it} is an indicator equal to 1 if the individual is employed under a permanent or temporary contract and 0 if he is unemployed or inactive. Self-employed workers are excluded from the analysis. X_{it} is a vector of controls, including an indicator for head of household, an indicator for being married, and the number of years of schooling. d_t is a set of year dummies and q_t is a set of quarter dummies. $treatment_{1i}$ is an indicator equal to 1 for young individuals (age 16 to 29) and $treatment_{2i}$ is an indicator equal to 1 for old individuals (age 45 to 59).

The coefficients δ_1 and δ_2 are the *difference-in-differences estimates*, which capture the effect of the reduction in dismissal costs on the employment probability of the young and the old. They are the coefficients on the interactions between the treatment indicators and $post_t$, which equals 1 for the period after the 1997 reform and 0 for the period before the reform. Because

[2]See, for example, European Commission Report (2006) for an overview of reforms to the Spanish pension system in the early 2000s.

the reform took place during the second quarter of 1997, this quarter is omitted. Therefore, the period before the reform extends from the first quarter of 1992 to the first quarter of 1997 and the period after the reform extends from the third quarter of 1997 to the fourth quarter of 2006. By going beyond 2001, the results capture the effects of both the 1997 reform and its extension in 2001.

Equation (1.5) is estimated separately for immigrants and natives and for men and women. The distinction between the effect on immigrants and natives is central to the analysis in this paper. The distinction between the effect on men and women is important given that the extension of the reform in 2001 gave special treatment to some categories of women. When the equation is estimated for the subsample of immigrants, the vector of controls X_{it} also includes the number of years of residence and the region of nationality (EU 15, other Europe, Central and South America, North America, Africa, and Asia and Oceania). In all estimations, standard errors are clustered by *year* $*$ *age group*, where *age group* is 16 to 29 (young), 30 to 44 (middle-aged), and 45 to 59 (old). This controls for common errors within these age groups.

The key element of this strategy is that the outcome of the control group (individuals age 30 to 44) is taken as a counterfactual for the experiences of the treatment groups in the absence of the reforms. This strategy is valid if there are no contemporaneous shocks, other than the reforms on dismissal costs, affecting the outcomes of the control and treatment groups during the period of analysis. If the business cycle affects the control and treatment groups differently, for example, then the estimates may capture the effect of the business cycle instead of the effect of the reduction in dismissal costs, producing biased results. To account for this possibility, the specification is changed to control for age group-specific business cycle effects:

$$\begin{aligned} e_{it} = \Lambda[&\alpha X_{it} + \gamma_0 d_t + \gamma_1 q_t + \beta_1 treatment_{1i} + \beta_2 treatment_{2i} + \delta_1(treatment_{1i} \times post_t) + \\ &\delta_2(treatment_{2i} \times post_t) + \lambda_1(treatment_{1i} \times expansion_t) + \lambda_2(treatment_{2i} \times expansion_t)] + \varepsilon_{it} \end{aligned} \tag{1.6}$$

In this specification, the variable $expansion_t$ is an indicator equal to 1 for the expansionary years 1995, and 1998 to 2000, and 0 otherwise. The interactions between the indicators for treatment and the expansion variable control for age group-specific business cycle effects.

Equations (1.5) and (1.6) capture the effect of the reform on the probability of being employed. Because a clear motivation for the reform was to reduce labor market segmentation, encouraging firms to replace temporary contracts with permanent ones, it is important to look not only at the effect on the probability of being employed, but also at the composition of employment. For this purpose, the model can be extended to a multinomial framework in which e_{it} takes three values: 1 if the individual is unemployed or inactive, 2 if he works under a temporary contract, and 3 if he works under a permanent contract. The estimation of the extended model is carried out by multinomial logit.

1.5.3 Data and Descriptive Statistics

The data are drawn from the Spanish economically active population survey (EPA - Encuesta de Población Activa), which collects detailed information on individuals' demographic characteristics and labor market outcomes. The data are at quarterly frequency and cover the period from the first quarter of 1992 to the fourth quarter of 2006. Data starting in the first quarter of 1999 are available on the website of the Spanish Statistics Institute (INE - Instituto Nacional de Estadística). Data for the earlier period were provided on request by the INE.

Because the purpose of this study is to compare the labor market performance of immigrants and natives, it is important to have information on nativity or nationality. It is also important to control for years of residence in Spain. The EPA reports this information from the first quarter of 1992 onwards. The identification of immigrants and natives is based on nationality rather than nativity, since the variable years of residence in Spain is only collected for foreign nationals.

To test whether the coverage of immigrants in the EPA is appropriate, Table 6 compares the percentage of immigrants in the EPA with equivalent numbers from OECD (2006), which are taken from the registry on the number of work permits until 1999 and, from 2000 onwards, from social security data. The table shows that immigrants are over-represented in the EPA relative to the data on work permits, and slightly under-represented relative to the social security data. Overall, though, the coverage of immigrants in the EPA seems reasonably accurate relative to these other sources.

The numbers in Table 6 show a remarkable increase in immigration into Spain, especially

since 2000. By the end of 2006, immigrants represented more than 8% of the Spanish labor force. Looking at the origin of immigrants, Figure 7 shows that, not surprisingly, most Spanish immigrants come from Central and South America. The second largest area of origin is other Europe, with a remarkable increase since 2000, most likely reflecting an increase in immigration from Eastern Europe following the enlargement of the European Union.

Descriptive statistics for the Spanish EPA are reported in Table 7. The probability of being employed on a permanent job increases with age and is larger for natives than for immigrants. The employment probability increased after the reforms for all groups. There also appears to be a change in the composition of employment for native workers affected by the reforms. Natives age 30 to 44 show no increase in their probability of permanent employment, but younger and older natives do appear to have found it easier to find permanent jobs. For these age groups, there was also a simultaneous reduction in the probability of temporary employment, suggesting that firms may have used the new type of contract to convert temporary jobs into permanent positions. This suggests that the 1997 and 2001 reforms had a positive effect on the permanent employment probability of natives and reduced duality in the Spanish labor market. For immigrants, the results are not as clear. The next section explores these effects in more detail, controlling for demographic characteristics.

1.5.4 Employment Effects

The results of estimating equations (1.5) and (1.6) for the employment probability and its decomposition into temporary and permanent employment are reported in Table 8. The evidence suggests that the reforms had no significant effect on the employment probability of young native men relative to middle-aged, but reduced the employment probability of old native men by 4.8 percentage points. Relative to middle-aged native women, the employment probability of young and old native women increased by 1.4 and 1.7 percentage points, respectively. There was no significant effect on the employment probability of immigrants, except for old immigrant women, whose employment probability increased by 5.2 percentage points relative to middle-aged immigrant women. Overall, the results suggest that the reduction in dismissal costs did not increase employment across the board. While it seems to have improved employment prospects for women, it had the opposite effect for old native men.

Turning to the effect of the reforms on permanent and temporary employment, the results show a large and statistically significant reduction in the temporary employment probability and an increase in the permanent employment probability for young natives and for old native women. For old native men, however, there is a reduction in both the temporary and the permanent employment probabilities. The effect on immigrants is less clear. There is evidence that the reform reduced the temporary employment probability for young immigrants and increased it for old immigrants, but the effects on the permanent employment probability are insignificant.

These results are consistent with the findings in Kugler et al (2002), which suggest that the 1997 reform increased permanent employment for young natives relative to the middle-aged, but had an insignificant effect on old natives. As discussed in section 1.2, theory makes no clear predictions regarding the effect of EPL on the level of employment, since it affects hiring and firing rates in the same direction. Using a panel version of the EPA which allows following the same individuals over time and measuring transitions, Kugler et al find that, for young natives, the 1997 reform increased transitions both from unemployment and from temporary employment to permanent employment and had little effect on dismissals. The increase in the hiring rate without a simultaneous increase in the firing rate explains the increase in permanent employment for young natives following the reform. The results in Kugler et al cannot be used to benchmark the effects on immigrants, since their study does not distinguish workers by nationality.

1.6 Evidence from the 1990 Italian Reform

1.6.1 Institutional Background

A significant reform in EPL in Italy took effect in May 1990 and updated legislation from 1970. The 1970 law introduced much stricter dismissal regulation for large firms than for small firms. While large firms (those above 15 employees) were required to pay foregone wages and rehire the worker in case of 'unfair' dismissal[3], small firms were exempt. The 1990 reform reduced this

[3] Aternatively, the worker may opt for a financial compensation (equal to 15 monthly payments) instead of being rehired.

discrepancy between large and small firms, introducing dismissal costs for firms with less than 15 employees. If the dismissal is considered to be 'unfair', small firms must either rehire the worker or pay compensation. The latter varies between 2.5 and 6 monthly payments, depending on the tenure of the dismissed worker. The rules for large firms were not changed. The different treatment given to large and small firms in the 1990 reform introduces variation which can be explored to study the effect of dismissal costs on employment.

The 1970 law was at the root of much tension in the early 2000s. In October 2001, the recently elected center-right government of Silvio Berlusconi proposed a series of labor market reforms, including the reform of the 1970 regulation on 'unfair' dismissals. The proposal was to abolish compulsory reinstatement of unfairly dismissed workers in favor of financial compensation. This generated controversy among trade unions and left-wing movements, culminating with the murder of the mentor of the proposed reforms, Marco Biagi, in March 2002, attributed to the terrorist group Brigate Rosse. Given the turmoil, the reform of the dismissal regulations was not implemented.

While there have not been major reforms on employment protection for permanent contracts in Italy (with the exception of the 1990 reform), there has been an increase in flexibility in the use of temporary contracts. In 1987 the government liberalized the use of temporary contracts, subject to collective agreements specifying certain target groups. Flexibility at the margin was further increased in 1997 with the liberalization of temporary help agencies.

1.6.2 Identification Strategy

The differential treatment given to small and large firms in the 1990 reform is explored to estimate the impact of dismissal costs on hiring and firing rates. Because the reform increased dismissal costs for firms with less than 15 employees and did not change costs for firms with more than 15 employees, a natural experiment can be constructed. Large firms are taken as the control group and their outcomes interpreted as the counterfactual, i.e., what would have happened in small firms if dismissal costs had not increased. The effect of the increase in dismissal costs is measured by comparing the change in outcomes before and after the reform in small relative to large firms.

Since firm size may itself be affected by the reform, the classification of firms as small and

large follows the strategy in Kugler and Pica (2005): small firms are those with less than 15 employees and large firms are those with more than 15 employees *in all years before the reform.* Therefore, the analysis excludes firms which crossed the 15 employees threshold before the reform.[4]

The following logit specification is used to estimate the effect of the 1990 reform on hiring and firing rates:

$$\begin{aligned} y_{ijt} &= \Lambda[\alpha_0 post_t + \alpha_1 small_j + \alpha_2 immigrant_i + \alpha_3(immigrant_i \times post_t) \\ &\quad +\alpha_4(immigrant_i \times small_j) + \alpha_5(small_j \times post_t) \\ &\quad +\alpha_6(immigrant_i \times smallj \times post_t) + \beta X_{ijt}] + \varepsilon_{it} \end{aligned} \tag{1.7}$$

The dependent variable y_{ijt} is an indicator equal to 1 if firm j hired (fired) worker i at time t and 0 otherwise. $post_t$ is an indicator equal to 1 for the period post 1990 and 0 for the period pre 1990. $small_j$ is an indicator equal to 1 if firm j had less than 15 employees in all years before 1990 and 0 otherwise. $immigrant_i$ is an indicator equal to 1 if worker i is an immigrant. X_{ijt} is a vector of controls, including worker characteristics (age, skill level, and log of wage) and firm characteristics (sector of activity, and region). All the controls are also interacted with the $immigrant_i$ indicator.

The effect of the reform on natives is given by the coefficient α_5 and the effect on immigrants is given by the sum of α_5 and α_6. Therefore, α_6 captures the differential effect of the reform on immigrants relative to natives. Because the 1990 law refers to permanent rather than temporary employment, the sample is limited to job spells under permanent contracts. In all specifications, standard errors are clustered by $post * firm\ size$.

As for the analysis of the Spanish reform, it is important to control for group-specific business cycle effects. If the business cycle affects small and large firms differently, the results will be biased and will be capturing the effect of the business cycle as well as the effect of the increase in dismissal costs. To account for this possibility, the following specification is used, where the

[4] This correction is important as there is evidence that firms change behaviour in response to the threshold. For example, Borgarello et al (2004) find that firms just below15 employees are more reluctant to hire than those more distant from the threshold.

variable $recession_t$ is an indicator equal to 1 during the recession years of 1992, 1993, and 1996:

$$\begin{aligned} y_{ijt} = \Lambda[&\alpha_0 post_t + \alpha_1 small_j + \alpha_2 immigrant_i + \alpha_3(immigrant_i \times post_t) \\ &+\alpha_4(immigrant_i \times small_j) + \alpha_5(small_j \times post_t) + \alpha_6(immigrant_i \times smallj \times post_t) \\ &+\lambda_1 recession_t + \lambda_2(small_j \times recession_t) + \lambda_3(recession_t \times immigrant_i) \\ &+\lambda_4(small_j \times recession_t \times immigrant_i) + \beta X_{ijt}] + \varepsilon_{it} \end{aligned} \tag{1.8}$$

1.6.3 Data and Descriptive Statistics

The data are drawn from the Work Histories Italian Panel (WHIP), which contains information on individual job spells and is based on administrative records from the Italian Social Security Institute (INPS)[5]. The reference population includes Italians and foreigners who have worked in Italy at some point during their careers. A sample of 1:90 is extracted from this population, including about 740,000 observations. Workers in the public sector are not sampled.

The dataset is a matched employer-employee panel and contains detailed information on the characteristics of the worker (including year of birth, gender, country of birth, and skill level), the job (including the dates when the job started and ended, the annual wage, and the number of days worked in the year) and the firm (including the exact number of employees, the year when the firm was established, its region, and sector of activity). It covers the years 1986 to 1999 and the definition of immigrants is based on country of birth, i.e., nativity.

Table 9 reports descriptive statistics. Small firms have higher hiring and firing rates than large firms. Moreover, there has been a larger decline in the hiring rate in small firms than in large firms after the reform for all demographic groups. Also, while the firing rate increased in large firms after the law, it did not change much in small firms. This is a first indication that the increase in dismissal costs for small firms reduced hiring and firing rates relative to large firms. However, to confirm this result, it is important to control for worker and firm characteristics.

[5]The WHIP data were provided on request by Laboratorio Revelli - Centre for Employment Studies. http://www.laboratoriorevelli.it/whip.

1.6.4 Effect on hiring and firing

The results of estimating equations (1.7) and (1.8) for the hiring and firing rates are reported in Table 10. The coefficients on $Small\ firm * Post$ capture the effect of the increase in dismissal costs on natives and the sum of the coefficients on $Small\ firm * Post$ and $Immigrant * Small\ firm * Post$ captures the effect on immigrants. The results suggest that the reform reduced hiring and firing rates for native men and women. The effects are quite large and very significant. The probability of being hired decreased by 3.4 percentage points for native men and 6.1 percentage points for native women following the reform. The probability of being fired decreased by 6.5 percentage points for native men and 5.9 percentage points for native women.

Turning to the effect on immigrants, the evidence suggest that the increase in dismissal costs also reduced their hiring and firing rates, but by a much smaller amount than for natives. The probability of being hired decreased by only 0.7 percentage points for immigrant men and 3.1 percentage points for immigrant women. The decrease in the probability of being fired was of 2.8 and 3.6 percentage points for immigrant men and women, respectively. This smaller impact on immigrants is consistent with the predictions of the model, since they are effectively less covered by EPL than natives.

1.7 Summary and Conclusions

This paper revisits the literature on the labor market effects of EPL, exploring its differential effect on natives and immigrants. Across three empirical exercises, it tells a consistent story: because immigrants are less protected by EPL than natives, they gain a comparative advantage in the labor market.

The results are summarized in Table 11. The evidence from the EU LFS suggests that stricter EPL on permanent contracts reduces employment of natives, but has a positive effect on employment of immigrants with more than six years of residence. This comes from a combination of the effects on hiring and firing rates: for natives, EPL reduces both hiring and firing, but the impact on hiring is larger, resulting in a reduction on the employment rate; for immigrants, EPL has no effect on hiring and reduces firing, resulting in an increase on the employment rate.

The analysis of the 1997 and 2001 Spanish reforms suggests that a reduction on EPL on permanent contracts increases permanent employment for natives but not for immigrants. The evidence from the 1990 Italian reform suggests that an increase in EPL reduces hiring and firing for natives, but has a much smaller effect on immigrants.

The consistency of the results from these three experiments is reassuring, since some estimations overcome the shortcomings of others. In particular, country-specific evidence has advantages over cross-country evidence as it allows disentangling the effects of EPL *per se* from the effects of the interaction of EPL with other labor market features, such as unemployment benefits, minimum wages, and early retirement schemes. Cross-country evidence does not separate these effects because of the lack of time series variation for many of these features. Country-specific evidence, however, separates these effects because different institutions interacting with EPL do not have the same cross-sectional variation as EPL.

Within the country-specific evidence, the experiment for Spain is probably superior to the one for Italy. This is because self-selection between the treatment and control groups is absent in Spain, where the treatment was applied as a function of age, but is likely to be present in Italy, where workers may move between large and small firms and employers may change their hiring and firing decisions in order to keep the size of the firm below a certain threshold.

The results point to an interesting interaction between labor market institutions and immigrant assimilation. Angrist and Kugler (2003) show that the displacement effect of immigrants on natives is smaller in countries with more flexible institutions. Our findings are supportive of that conclusion and suggest that, in countries with restrictive labor markets, immigrants gain a comparative advantage relative to the natives because they are less protected. There appears to be a silver lining to labor market rigidities after all.

Figure 1. Union Membership Rates by Citizenship

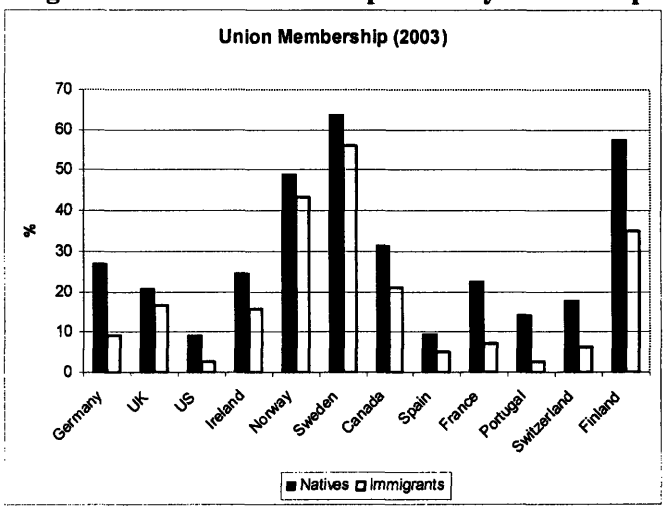

SOURCE: International Social Survey Programme (2003)

Table 1. Coverage of the EU LFS

	LFS coverage	% of immigrants in the population	
		LFS 2005	OECD 2004
Austria	1995-2005	14.9	13
Belgium	1995-2005	13	11.4
Denmark	1995-2005	6.2	6.3
Finland	1995-2005	2.6	3.2
France	1995-2005	11.7	10 (a)
Germany	1999-2005	n.a.	12.9
Greece	1995-2005	6.7	10.3 (b)
Ireland	1999-2005	10.3	11
Netherlands	1996-2005	12.2	10.6
Norway	1996-2005	7.5	7.8
Portugal	1995-2005	6.1	6.7
Spain	1995-2005	10.1	5.3 (b)
Sweden	1995-2005	13	12.2
Switzerland	2005	24.7	23.5
UK	1995-2005	10.4	9.3

SOURCE: EU Labor Force Survey and OECD (2006)
NOTE: Definition of immigrants is based on nativity. OECD (2006) numbers come from the last Census in or before 2004. (a) Data from 1999 Census. (b) Data from 2001 Census.

Table 2. Characteristics of Immigrants and Natives

	Men				Women			
	Natives	Immigrants			Natives	Immigrants		
		0-5 years residence	6-10 years residence	>10 years residence		0-5 years residence	6-10 years residence	>10 years residence
A. Labor Force Participation rate								
Austria	88.33	83.86	92.74	87.91	71.76	53.99	70.90	69.52
Belgium	85.06	78.10	82.67	77.94	66.14	45.08	48.43	49.52
Denmark	90.87	74.97	72.10	82.02	82.33	47.66	56.40	72.22
Finland	86.13	80.72	85.48	87.40	81.13	52.82	63.05	78.09
France	87.88	73.77	85.83	87.30	74.18	40.36	53.31	63.76
Germany	90.44	79.80	90.69	89.18	76.50	49.65	62.67	66.59
Greece	88.52	91.48	95.40	92.24	55.42	61.54	64.29	60.67
Ireland	87.99	78.19	86.34	89.55	60.16	57.06	58.31	62.55
Netherlands	91.68	66.26	80.24	82.73	72.85	44.64	51.24	62.03
Norway	89.26	83.67	86.31	84.66	81.41	59.40	68.86	75.94
Portugal	88.24	92.49	84.23	86.11	71.29	75.34	67.65	75.57
Spain	87.21	93.21	92.64	89.07	56.25	69.98	67.03	58.27
Sweden	89.05	69.83	80.96	82.92	84.92	54.37	67.25	76.63
Switzerland	94.34	94.04	93.77	89.71	86.44	69.18	70.61	76.01
UK	90.01	74.80	84.11	87.14	74.32	53.59	57.25	66.30
B. Employment/Population								
Austria	84.72	74.76	84.32	79.63	68.79	47.86	66.04	64.35
Belgium	80.38	64.38	67.08	66.39	60.32	34.67	37.89	40.68
Denmark	86.30	61.98	62.61	73.45	76.92	38.50	47.27	65.61
Finland	77.07	59.71	63.22	72.84	72.68	36.52	47.02	66.26
France	80.65	54.66	66.54	73.70	65.68	27.27	37.62	52.72
Germany	83.11	63.66	76.47	78.21	69.84	38.44	54.43	58.83
Greece	83.05	82.17	89.28	83.52	47.21	48.26	52.40	49.83
Ireland	80.94	70.25	78.59	81.61	56.00	50.82	53.66	57.79
Netherlands	89.52	56.33	70.90	76.57	70.11	37.55	45.73	57.76
Norway	86.54	72.96	74.14	79.91	78.92	53.71	61.56	72.53
Portugal	84.57	85.37	75.52	80.04	67.23	64.91	59.64	69.19
Spain	76.92	81.86	82.46	80.28	44.05	59.20	55.93	49.09
Sweden	83.36	53.86	61.70	73.15	80.44	41.26	54.42	69.52
Switzerland	89.45	86.84	83.33	83.99	82.22	58.68	61.21	69.91
UK	83.65	65.08	73.36	78.54	70.73	47.19	51.83	61.67
C. Permanent employment/Employment								
Austria	77.3	76.8	84.5	84.7	79.2	79.7	86.5	87.1
Belgium	77.6	73.4	73.3	75.2	77.1	69.1	72.2	74.2
Denmark	79.9	73.9	72.1	75.7	84	65.9	64	79.7
Finland	70.2	59.7	65.3	69.8	71.9	50.9	51.6	63.2
France	74.8	60.6	69	75.5	77.6	58.6	73.3	80.2
Germany	74.8	65.7	77.1	79.1	79.4	66.7	76.7	81.8
Greece	49.5	67.1	70.9	64.9	52.3	63.1	63.8	65.8
Ireland	70.9	80.7	70.4	70.2	86.7	81.7	77.7	80.6
Netherlands	77.1	61.2	68	76	75.6	55.7	65.8	75.7
Norway	80.4	75.4	73.4	79	81.8	68.6	73.1	78.9
Portugal	59.7	39.1	43.3	62.4	60.2	47.9	54.9	66.9
Spain	53.8	28.8	36.4	38.4	53.8	43.7	43	43.5
Sweden	74.3	57.2	62.7	69.7	78.9	51.2	58.7	76.7
Switzerland	73.3	73.5	78.5	79.9	72	73.9	76.2	82.6
UK	77.7	74.2	73.3	70.5	85.4	69.3	78.9	79.9
D. Temporary employment/Employment								
Austria	6.5	15	9.2	4.3	6.9	13	8.5	4.5
Belgium	4.7	12	9	4.3	9.1	16.8	10.8	8.7
Denmark	7.8	21.3	17	8.1	10.4	20.3	27.7	12
Finland	11.4	31.8	20.1	13.8	18.6	55.2	40.4	32
France	10.4	29.2	18.3	10.6	13.2	31.1	20.4	13.5
Germany	10.7	27.4	16	7.4	11.5	27.4	17.6	8.7
Greece	5.5	25.8	17.3	9.6	7.9	28.9	22.1	12.1
Ireland	2.9	8.1	3.4	2.2	5.3	11.8	5.5	5
Netherlands	8.8	31.9	22.7	11.8	14	37.7	25.5	15
Norway	7.3	16.1	15.1	8.3	11.9	24.3	18	13.4
Portugal	10.7	45.5	32.7	16.7	13.1	42.5	26	17.9
Spain	22.9	57	40.5	24.8	28.1	48.1	37.6	23.8
Sweden	9.8	29	19.4	10.8	14.7	39.5	32.6	15.6
Switzerland	18	17.7	13.1	6.8	20.6	18.5	14.9	8.3
UK	4.5	17	9.5	4.5	6.4	22.7	10.7	7

SOURCE: EU Labor Force Survey

NOTE: Averages over the period of data coverage. Sample restricted to population aged 20-59. Definition of immigrants is based on nativity.

Figure 2. Evolution of EPL Index for Permanent Contracts 1985-2005

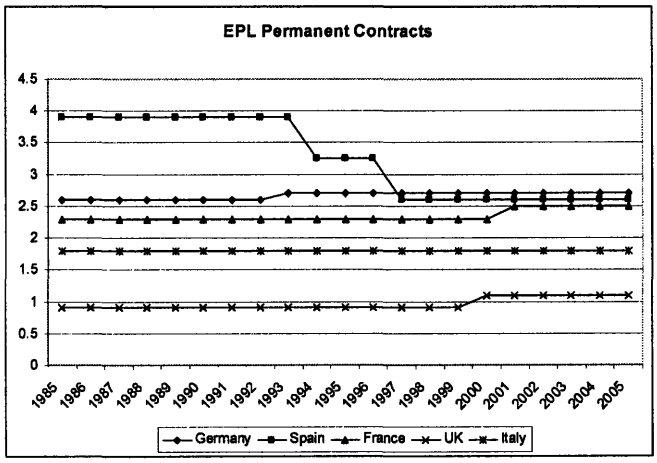

SOURCE: See Data Appendix

Figure 3. Evolution of EPL Index for Temporary Contracts 1985-2005

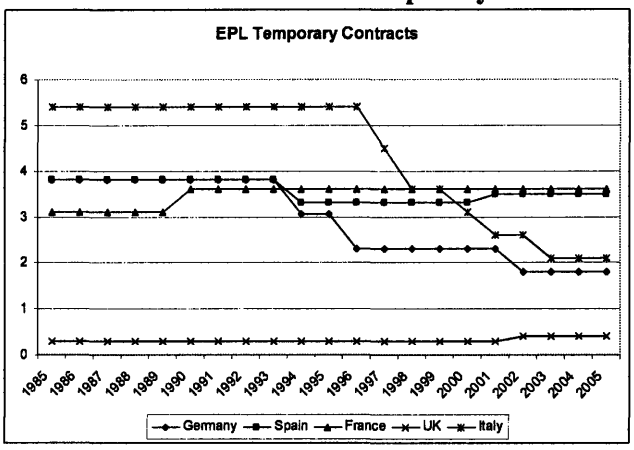

SOURCE: See Data Appendix

Figure 4. Correlation between EPL and Unemployment Duration

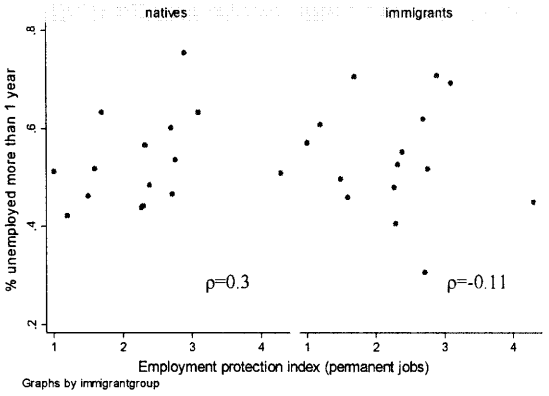

Figure 5. Correlation between EPL and Employment Duration

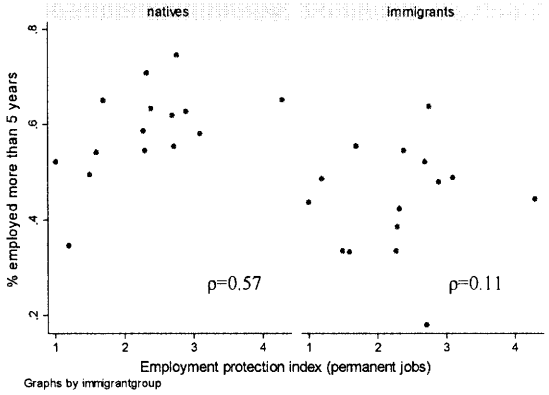

Figure 6. Correlation between EPL and the Employment Rate

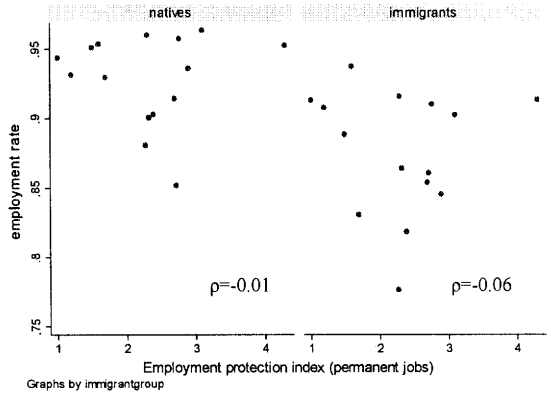

SOURCE: for labor market variables, EU Labor Force Survey; for employment protection index, see Data Appendix.
NOTE: ρ denotes the correlation coefficient between the series in the two axes.

Table 3. Effect of EPL on Participation and Employment

	Labor Force Participation Rate		Employment Rate	
	(1)	(2)	(1)	(2)
EPL permanent	-0.006	-	-0.038*	-
	(0.012)		(0.023)	
EPL temporary	-	0.009**	-	0.013***
		(0.004)		(0.004)
Years residence 0-5	0.111***	0.109***	-0.019	-0.029
	(0.030)	(0.026)	(0.027)	(0.027)
Years residence 6-10	0.175***	0.224***	0.015	0.037*
	(0.033)	(0.027)	(0.023)	(0.020)
Years residence >10	0.074***	0.080***	-0.093***	-0.060***
	(0.018)	(0.018)	(0.021)	(0.014)
Years residence 0-5*EPL permanent	0.015	-	0.003	-
	(0.030)		(0.047)	
Years residence 6-10*EPL permanent	0.041	-	0.064**	-
	(0.036)		(0.028)	
Years residence >10*EPL permanent	0.003	-	0.108***	-
	(0.023)		(0.032)	
Years residence 0-5*EPL temporary	-	-0.002	-	-0.013
		(0.013)		(0.015)
Years residence 6-10*EPL temporary	-	0.023**	-	-0.013
		(0.010)		(0.011)
Years residence >10*EPL temporary	-	-0.011*	-	-0.015**
		(0.006)		(0.006)
Medium education	0.143***	0.141***	0.048***	0.047***
	(0.005)	(0.005)	(0.003)	(0.003)
High education	0.232***	0.231***	0.070***	0.069***
	(0.006)	(0.006)	(0.004)	(0.004)
Medium education*EPL permanent	-0.040***	-	-0.016***	-
	(0.006)		(0.003)	
High education*EPL permanent	-0.036***	-	-0.025***	-
	(0.007)		(0.004)	
Medium education*EPL temporary	-	-0.024***	-	-0.008***
		(0.003)		(0.003)
High education*EPL temporary	-	-0.014***	-	-0.005
		(0.004)		(0.003)
Female	-0.121***	-0.122***	-0.024***	-0.024***
	(0.005)	(0.005)	(0.003)	(0.003)
Age 25-39	0.274***	0.273***	0.089***	0.089***
	(0.011)	(0.011)	(0.005)	(0.005)
Age 40-59	0.213***	0.213***	0.118***	0.117***
	(0.010)	(0.010)	(0.007)	(0.007)

SOURCE: EU Labor Force Survey

NOTE: EPL is measured as deviation from cross country mean; N=4428; standard errors adjusted for clustering by country*year; regressions also control for country and year dummies, country*years residence, year*years residence, female*years residence, age*years residence and education*years residence. OLS estimates. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.

Table 4. Effect of EPL on Employment and Unemployment Duration

	Unemployed >1 year		Employed > 5 years		Employed > 1 year	
	(1)	(2)	(1)	(2)	(1)	(2)
EPL permanent	0.132***	-	0.097**	-	0.002	-
	(0.043)		(0.045)		(0.042)	
EPL temporary	-	-0.012	-	0.008	-	0.023***
		(0.018)		(0.005)		(0.006)
Years residence 0-5	0.038	0.083	-0.078**	-0.150***	-0.135***	-0.092*
	(0.057)	(0.062)	(0.032)	(0.045)	(0.048)	(0.049)
Years residence 6-10	0.171**	0.100	0.011	0.032	0.218***	0.190***
	(0.075)	(0.074)	(0.087)	(0.082)	(0.043)	(0.042)
Years residence >10	-0.032	-0.100**	-0.050**	-0.047**	-0.060***	-0.059***
	(0.035)	(0.039)	(0.020)	(0.018)	(0.022)	(0.017)
Years residence 0-5*EPL permanent	0.043	-	-0.053	-	0.175***	-
	(0.085)		(0.071)		(0.060)	
Years residence 6-10*EPL permanent	-0.195**	-	0.043	-	-0.079	-
	(0.097)		(0.086)		(0.057)	
Years residence >10*EPL permanent	-0.081*	-	0.048	-	0.034	-
	(0.048)		(0.036)		(0.026)	
Years residence 0-5*EPL temporary	-	0.036	-	-0.040***	-	-0.027*
		(0.029)		(0.015)		(0.015)
Years residence 6-10*EPL temporary	-	0.029	-	-0.010	-	0.007
		(0.030)		(0.020)		(0.013)
Years residence >10*EPL temporary	-	-0.014	-	0.012	-	-0.010*
		(0.015)		(0.008)		(0.006)
Medium education	-0.043***	-0.041***	0.018***	0.019***	0.027***	0.028***
	(0.006)	(0.006)	(0.004)	(0.003)	(0.002)	(0.002)
High education	-0.084***	-0.083***	-0.021***	-0.021***	-0.001	-0.001
	(0.010)	(0.010)	(0.004)	(0.004)	(0.004)	(0.003)
Medium education*EPL permanent	0.029***	-	0.002	-	0.001	-
	(0.008)		(0.004)		(0.003)	
High education*EPL permanent	0.041***	-	0.007*	-	-0.007**	-
	(0.011)		(0.004)		(0.003)	
Medium education*EPL temporary	-	0.014***	-	-0.010***	-	-0.008***
		(0.005)		(0.003)		(0.002)
High education*EPL temporary	-	0.019***	-	-0.001	-	-0.007***
		(0.006)		(0.003)		(0.003)
Female	0.062***	0.062***	-0.030***	-0.030***	-0.022***	-0.022***
	(0.005)	(0.005)	(0.002)	(0.002)	(0.002)	(0.002)
Age 25-39	0.182***	0.183***	0.429***	0.430***	0.308***	0.308***
	(0.008)	(0.008)	(0.005)	(0.005)	(0.006)	(0.006)
Age 40-59	0.292***	0.292***	0.707***	0.707***	0.415***	0.414***
	(0.011)	(0.011)	(0.006)	(0.006)	(0.009)	(0.009)

SOURCE: EU Labor Force Survey

NOTE: Dependent variable: for unemployment duration, fraction of unemployed who have been in unemployment for more than one year; for employment duration, fraction of employed who have been on the current job for more than one year/more than five years. EPL is measured as deviation from cross country mean. N=4317. Standard errors adjusted for clustering by country*year. Regressions also control for country and year dummies, country*years residence, year*years residence, female*years residence, age*years residence and education*years residence. OLS estimates.

Table 5. Labor Market Reforms in Spain after 1997

	Dismissal costs under existing permanent contracts	Dismissal costs under new permanent contracts	Payroll tax reductions for newly hired workers under permanent contracts in 1997-1998	Payroll tax reductions for newly hired workers under permanent contracts in 1999
Unemployed aged 30-44 years	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	None	None
Young unemployed workers (under 30 years of age)	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 33 days' wages per year of seniority with a maximum of 24 months' wages	40% of employer contributions for 24 months	35% of employer contributions for 12 months, 25% for another 12 months
Unemployed workers above 45 years of age	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 33 days' wages per year of seniority with a maximum of 24 months' wages	60% of employer contributions for 24 months, 50% thereafter	45% of employer contributions for 12 months, 40% for another 12 months
Long-term unemployed (over 1 year of registered unemployment)	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 33 days' wages per year of seniority with a maximum of 24 months' wages	40% of employer contributions for 24 months	40% of employer contributions for 12 months, 30% for another 12 months
Workers employed under temporary contracts	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 33 days' wages per year of seniority with a maximum of 24 months' wages	50% employer contributions for 24 months, 20% for another 12 months	None
Women hired under temporary contracts or long-term unemployed hired in occupations with low weight of female employment	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 33 days' wages per year of seniority with a maximum of 24 months' wages	60% employer contributions for 24 months, 20% for another 12 months	45% employer contributions for 24 months, 40% for another 12 months
Workers hires under training contracts	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 33 days' wages per year of seniority with a maximum of 24 months' wages	50% employer contributions for 24 months, 20% for another 12 months	25% employer contributions for 24 months
Workers above 45 years of age hired under temporary contracts	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 33 days' wages per year of seniority with a maximum of 24 months' wages	60% employer contributions for 24 months, 20% for another 12 months	60% employer contributions for 24 months, 20% for another 12 months
Disabled workers	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 45 days' wages per year of seniority with a maximum of 42 months' wages	Fair dismissals: 20 days' wages per year of seniority with a maximum of 12 months' wages Unfair dismissals: 33 days' wages per year of seniority with a maximum of 24 months' wages	70%-90% for the whole employment spell	70%-90% for the whole employment spell

SOURCE: reproduced from Kugler et al (2002)

Table 6. Coverage of Immigrants in the Spanish EPA

Year	% Immigrants in Labor Force	
	Spain LFS	OECD (2006)
1995	1.08	0.8
1996	1.11	1.0
1007	1.18	1.1
1008	1.38	1.2
1999	1.58	1.1
2000	1.79	2.5
2001	2.56	3.4
2002	3.52	4.5
2003	4.67	5.2
2004	5.54	6.3
2005	6.62	n.a.
2006	8.34	n.a.

SOURCE: Spanish economically active population survey (EPA) and OECD (2006)
NOTE: OECD (2006) data comes from the number of valid work permits. EU workers are not included. From 2000 on, data refers to the number of workers who are registered in the Social Security system. The source of the data is the Ministry of Labor and Social Security.

Figure 7. Origin of Immigrants - Spain

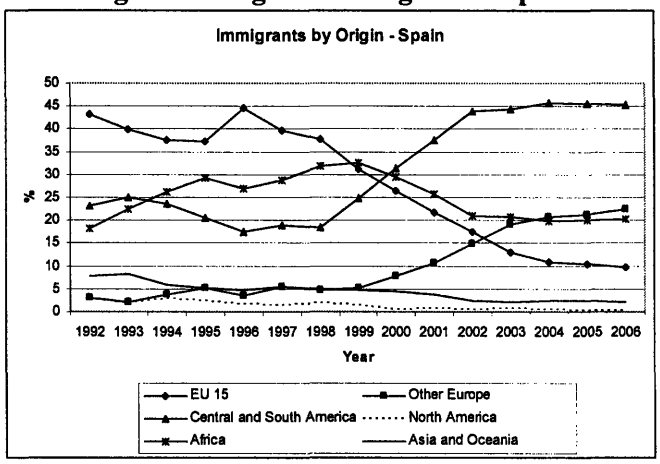

SOURCE: Spanish economically active population survey (EPA)

Table 7. Descriptive Statistics for the Spanish EPA

	Age 16-29		Age 30-44		Age 45-59	
	Pre-reform	Post-reform	Pre-reform	Post-reform	Pre-reform	Post-reform
			A. Native Men			
Employment probability	39.25	49.58	78.09	84.54	67.1	74.25
Permanent employment probability	36.23	43.06	75.06	74.92	84.62	85.08
Temporary employment probability	63.77	56.94	24.94	25.08	15.38	14.92
			B. Native Women			
Employment probability	25.97	36.51	36.68	50.83	19.16	32.46
Permanent employment probability	37.33	42.4	73.11	70.9	80.02	82.2
Temporary employment probability	62.67	57.6	26.89	29.1	19.98	17.8
			C. Immigrant Men			
Employment probability	48.29	68.72	69.08	83.46	52.66	71.46
Permanent employment probability	25.4	26.19	48.63	40.62	71.33	50.8
Temporary employment probability	74.6	73.81	51.37	59.38	28.67	49.2
			D. Immigrant Women			
Employment probability	24.78	47.1	31.95	56.11	22.27	47.09
Permanent employment probability	35.4	35.7	51.8	47.06	73.25	56.77
Temporary employment probability	64.6	64.3	48.2	52.94	26.75	43.23

SOURCE: Spanish economically active population survey (EPA)

Table 8. Employment Effects of the 1997 and 2001 Spanish Reforms

	(1)			(2)		
	P(employment)	Log[P(temporary)/ P(non-employment)]	Log[P(permanent)/ P(non-employment)]	P(employment)	Log[P(temporary)/ P(non-employment)]	Log[P(permanent)/ P(non-employment)]
			A. Native Men			
Age 16-29	-0.197***	-0.202***	-1.527***	-0.197***	-0.209***	-1.512***
	(0.010)	(0.041)	(0.055)	(0.010)	(0.042)	(0.056)
Age 45-59	-0.156***	-1.289***	-0.488***	-0.160***	-1.296***	-0.509***
	(0.010)	(0.045)	(0.046)	(0.011)	(0.046)	(0.047)
Age 16-29*Post	-0.007	-0.177***	0.165***	-0.007	-0.180***	0.180***
	(0.005)	(0.034)	(0.034)	(0.005)	(0.034)	(0.029)
Age 45-59*Post	-0.048***	-0.158***	-0.252***	-0.051***	-0.164***	-0.268***
	(0.008)	(0.036)	(0.046)	(0.008)	(0.036)	(0.043)
Age 16-29*Expansion	-	-	-	0.001	0.029	-0.094**
				(0.005)	(0.023)	(0.038)
Age 45-59*Expansion	-	-	-	0.017**	0.035	0.113**
				(0.009)	(0.029)	(0.050)
			B. Native Women			
Age 16-29	-0.186***	-0.091	-1.542***	-0.183***	-0.095	-1.508***
	(0.016)	(0.075)	(0.094)	(0.016)	(0.076)	(0.095)
Age 45-59	-0.116***	-0.985***	-0.375***	-0.115***	-00980***	-0.374***
	(0.008)	(0.048)	(0.041)	(0.009)	(0.051)	(0.043)
Age 16-29*Post	0.014*	-0.128***	0.239***	0.017**	-0.128***	0.264***
	(0.008)	(0.038)	(0.052)	(0.007)	(0.037)	(0.046)
Age 45-59*Post	0.017**	-0.084*	0.129***	0.017**	-0.081	0.130***
	(0.008)	(0.048)	(0.039)	(0.008)	(0.052)	(0.036)
Age 16-29*Expansion	-	-	-	-0.018***	0.021	-0.180***
				(0.006)	(0.022)	(0.050)
Age 45-59*Expansion	-	-	-	-0.0004	-0.029	-0.002
				(0.005)	(0.044)	(0.026)
			C. Immigrant Men			
Age 16-29	-0.082***	-0.171**	-0.915***	-0.078***	-0.176**	-0.867***
	(0.018)	(0.081)	(0.118)	(0.018)	(0.079)	(0.125)
Age 45-59	-0.178***	-1.227***	-0.833***	-0.168***	-1.209***	-0.765***
	(0.023)	(0.114)	(0.118)	(0.024)	(0.121)	(0.132)
Age 16-29*Post	-0.008	-0.239***	0.216*	-0.01	-0.231***	0.199*
	(0.016)	(0.079)	(0.111)	(0.016)	(0.075)	(0.117)
Age 45-59*Post	0.006	0.332***	0.023	0.004	0.333***	0.008
	(0.019)	(0.117)	(0.124)	(0.019)	(0.119)	(0.131)
Age 16-29*Expansion	-	-	-	-0.013	0.016	-0.239**
				(0.018)	(0.085)	(0.100)
Age 45-59*Expansion	-	-	-	-0.039**	-0.086	-0.334***
				(0.019)	(0.119)	(0.092)
			D. Immigrant Women			
Age 16-29	-0.127***	-0.239***	-0.907***	-0.125***	-0.244***	-0.908***
	(0.022)	(0.062)	(0.200)	(0.022)	(0.062)	(0.203)
Age 45-59	-0.138***	-1.093***	-0.291***	-0.139***	-1.107***	-0.278***
	(0.017)	(0.093)	(0.100)	(0.018)	(0.094)	(0.103)
Age 16-29*Post	-0.002	-0.117*	0.134	-0.002	-0.115*	0.134
	(0.022)	(0.062)	(0.197)	(0.022)	(0.061)	(0.198)
Age 45-59*Post	0.052**	0.593***	0.046	0.052**	0.061***	0.044
	(0.022)	(0.108)	(0.111)	(0.022)	(0.106)	(0.110)
Age 16-29*Expansion	-	-	-	-0.011	0.023	0.005
				(0.014)	(0.052)	(0.092)
Age 45-59*Expansion	-	-	-	0.005	0.082	-0.081
				(0.016)	(0.073)	(0.073)

SOURCE: Spanish economically active population survey (EPA)

NOTE: Logit marginal effects. Regressions include year and quarter dummies, an indicator for head of household, an indicator for being married, number of years of schooling and, for immigrants, years of residence and region of nationality (EU 15, other Europe, Central and South America, North America, Africa, and Asia and Oceania).

Table 9. Descriptive Statistics for the Italian WHIP

	Small firms		Large firms	
	Pre-reform	Post-reform	Pre-reform	Post-reform
		A. Native Men		
Hiring rate	37.86	11.06	27.17	9.50
Firing rate	28.98	28.13	14.59	21.71
		B. Native Women		
Hiring rate	35.13	9.68	28.44	12.07
Firing rate	26.54	27.00	16.45	23.82
		C. Immigrant Men		
Hiring rate	49.32	27.11	38.27	23.94
Firing rate	40.18	40.65	24.84	32.62
		D. Immigrant Women		
Hiring rate	51.59	26.44	41.10	23.83
Firing rate	43.65	40.69	30.47	33.49

SOURCE: Work Histories Italian Panel (WHIP)
NOTE: Sample restricted to workers with permanent contracts.

Table 10. Effect of the 1990 Italian Reform on Hiring and Firing Rates

	Hiring Rate		Firing Rate	
	(1)	(2)	(1)	(2)
			A. Men	
Post	-0.141***	-0.135***	0.057***	0.090***
	(0.013)	(0.014)	(0.011)	(0.011)
Small firm	0.0002	0.0001	0.104***	0.103***
	(0.010)	(0.010)	(0.012)	(0.011)
Immigrant	0.087	0.091	0.528***	0.537***
	(0.242)	(0.241)	(0.194)	(0.185)
Immigrant*Post	0.060***	0.064***	-0.011**	-0.016***
	(0.010)	(0.011)	(0.005)	(0.005)
Immigrant*Small firm	-0.005	-0.005	-0.032***	-0.032***
	(0.005)	(0.005)	(0.008)	(0.008)
Small firm*Post	-0.034***	-0.034***	-0.065***	-0.073***
	(0.003)	(0.003)	(0.002)	(0.002)
Immigrant*Small firm*Post	0.027***	0.017***	0.037***	0.031***
	(0.004)	(0.003)	(0.003)	(0.003)
Recession	-	-0.015***	-	-0.086***
		(0.001)		(0.001)
Small firm*Recession	-	0.0003	-	0.043***
		(0.0005)		(0.0004)
			B. Women	
Post	-0.132***	-0.129***	0.043***	0.075***
	(0.009)	(0.009)	(0.006)	(0.006)
Small firm	-0.013*	-0.013*	0.080***	0.079***
	(0.007)	(0.007)	(0.008)	(0.008)
Immigrant	0.156	0.157	0.616***	0.611***
	(0.209)	(0.212)	(0.082)	(0.084)
Immigrant*Post	0.040***	0.044***	-0.003	-0.004
	(0.011)	(0.011)	(0.005)	(0.006)
Immigrant*Small firm	-0.008***	-0.008***	-0.044***	-0.044***
	(0.002)	(0.002)	(0.003)	(0.003)
Small firm*Post	-0.061***	-0.062***	-0.059***	-0.066***
	(0.003)	(0.004)	(0.003)	(0.003)
Immigrant*Small firm*Post	0.030***	0.023***	0.023***	0.028***
	(0.006)	(0.006)	(0.007)	(0.006)
Recession	-	-0.008***	-	-0.085***
		(0.0007)		(0.0008)
Small firm*Recession	-	0.004***	-	0.036***
		(0.001)		(0.0009)

SOURCE: Work Histories Italian Panel (WHIP)
NOTE: logit marginal effects. Regressions also control for log(wage), age, sector of activity, region where firm is located, skill level and interactions of all the covariates with the immigrant dummy.

Table 11. Summary of results

Experiment:	EU		Spain		Italy	
	Increase in EPL index		Reduction in EPL for young and old workers		Increase in EPL for small firms	
Effect on:	Natives	Immigrants	Natives	Immigrants	Natives	Immigrants
Employment rate	-	+ (a)				
Permanent employment			+	insignificant		
Hiring rate	--	insignificant			--	-
Firing rate	-	-			--	-

NOTE: (a) for immigrants with more than 6 years of residence.

Data Appendix

The Eurostat EU Labour Force Survey

The Eurostat EU Labour Force Survey (LFS) covers the 25 EU member states, as well as Iceland and Norway. The data are collected by national statistical agencies, following a common procedure developed by Eurostat, ensuring comparability across countries. The data used in this paper were constructed on request by the German Federal Statistics Office and contain information on immigration and labor market variables aggregated by the following groups:

- gender: male, female.
- age: 15 to 19, 20 to 24, 25 to 29, 30 to 34, 35 to 39, 40 to 44, 45 to 49, 50 to 54, and 55 to 59.
- education: low (lower secondary), medium (upper secondary), and high (third level).
- nativity: native born, foreign born (aggregated by years of residence: 0 to 5, 6 to 10, more than 10).

Eurostat implements guidelines to ensure that the data are statistically reliable. In particular, it computes thresholds of population counts below which the data are unreliable. All unreliable observations were dropped from the dataset used in this paper.

Construction of EPL Time Series

The time series for EPL on permanent and temporary contracts used in this paper cover the period 1985 to 2005 and were constructed from the OECD indicators on the strictness of EPL for the late 1980s, the late 1990s, and 2003, using information on the timing of reforms which introduced breaking points in the series. The indicators and details on how they are constructed are available in OECD (1999) and OECD (2004). The version of the indicators used in this paper is version 1, which does not include information on collective dismissals, but covers a longer period of time.

The OECD indicators aggregate information in a number of dimensions. For permanent contracts, they contain information on severance payments, notification procedures, the definition of 'fair' and 'unfair' dismissals, the amount of compensation following 'unfair' dismissals,

among other measures. For temporary contracts, they contain information on the valid cases for use of this type of contracts, the maximum number of fixed term contracts permitted, their maximum cumulative duration, as well as information on temporary work agencies.

To transform the EPL indices for the late 1980s, the late 1990s, and 2003 into a time series from 1985 to 2005, the information on the breaking points of the EPL time series given in table 2.A2.6 in OECD (2004) is used. This table lists a number of reforms for each country which had an impact on the EPL indices for permanent and temporary contracts. The starting point of the time series, 1985, takes the values of the OECD indices for the late 1980s. The other trivial point is 2003, where the series take the same values as the OECD indices for 2003. Between those dates, the series are constructed using information on the breaking points.

It is useful to give some illustrative examples. For instance, Belgium adopted a reform in 1997 which reduced restrictions on temporary contracts. This was the only reform on regulation on temporary contracts in Belgium during the period considered. Therefore, the series on EPL for temporary contracts in Belgium takes the value of the late 1980s from 1985 to 1996. In 1997 it falls to the value in the late 1990s (which equals the value in 2003 as there were no reforms after 1997).

Sometimes it is necessary to interpolate. For example, Spain had two reforms on regulation on permanent contracts: one in 1994, which relaxed the procedural requirements for dismissals for economic reasons, and another in 1997, which introduced a new type of contract, applicable to certain categories of workers, with reduced compensation for 'unfair' dismissals. Given these reforms, the series on EPL for permanent contracts in Spain takes the value of the late 1980s from 1985 to 1993. In 1997 it falls to the value in the late 1990s. Between 1994 and 1996, it is a linear interpolation of the values in the late 1980s and the late 1990s.

Bibliography

[1] Angrist, J. and A. Kugler (2003), 'Protective or Counter-Productive? Labour Market Institutions and the Effect of Immigration on EU Natives.', *The Economic Journal*, Vol. 113, No. 488 (June).

[2] Bertola, G. (1990), 'Job Security, Employment, and Wages', *European Economic Review*, Vol. 34, pp. 851-886.

[3] Blanchard, O. and P. Portugal (2001), 'What Hides Behind an Unemployment Rate: Comparing Portuguese and US Labor Markets', *The American Economic Review*, Vol. 91, No. 1 (March 2001), pp. 187-207.

[4] Blanchard, O. and J. Wolfers (2000), 'The Role of Shocks and Institutions in the Rise of European Unemployment: the Aggregate Evidence', *The Economic Journal*, Vol. 110, No. 462 (March), Conference Papers.

[5] Borgarello, A., P. Garibaldi, and L. Pacelli (2004), 'Employment Protection Legislation and the Size of Firms', *Il Giornale degli Economisti*, No. 1.

[6] Cahuc, P. and A. Zylberberg (2004), 'Labour Economics', The MIT Press, Cambridge, Massachusetts, and London, England.

[7] European Commission Report (2006), 'Synthesis Report on Adequate and Sustainable Pensions. Country Summaries: Spain'.

[8] Kugler, A., J. Jimeno, and V. Hernanz (2002), 'Employment Consequences of Restrictive Permanent Contracts: Evidence from Spanish Labor Market Reforms', IZA Discussion Paper, No. 657, November 2002.

[9] Kugler, A. and G. Pica (2005), 'Effects of Employment Protection on Worker and Job Flows: Evidence from the 1990 Italian Reform', IZA Discussion Paper, No. 1743, September 2005, forthcoming in *Labour Economics.*

[10] Lazear, E. (1990), 'Job Security Provisions and Employment', *The Quarterly Journal of Economics*, Vol. 105, No. 3, pp. 699-726.

[11] Migration Policy Institute - MPI (2004), 'Fact Sheet #7: Immigrant Union Members - Numbers and Trends'.

[12] Moulton, B. R. (1990), 'An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units', *Review of Economics and Statistics*, Vol. 72, No. 2, pp. 334-338.

[13] Nickell, S. (1997), 'Unemployment and Labor Market Rigidities: Europe versus North America', *Journal of Economic Perspectives*, Vol. 11, No. 3, pp. 55-74.

[14] OECD (1999), Employment Outlook.

[15] OECD (2004), Employment Outlook.

[16] OECD (2006), International Migration Outlook: SOPEMI.

[17] Pissarides, C. (2000), 'Equilibrium Unemployment Theory', second edition, The MIT Press, Cambridge, Massachusetts, and London, England.

[18] Scarpetta, S. (1996), 'Assessing the Role of Labour Market Policies and Institutional Settings on Unemployment: a Cross-Country Study', *OECD Economic Studies*, No. 26.

Chapter 2

The 35-Hour Workweek in France: Straightjacket or Welfare Improvement?

2.1 Introduction

France's workweek is one of the shortest in the world. The latest reduction in the 'standard' workweek (i.e., the number of weekly hours paid at the regular wage rate) was promoted by Lionel Jospin's socialist government, who reduced it from 39 to 35 hours in the early 2000s. While the purpose of this policy was to decrease unemployment, seven years later the unemployment rate remains high (above 8%), making the impact of the 35-hour workweek on the French economy a topic of heated discussion during the 2007 presidential campaign. The socialist candidate, Ségolène Royal, criticised the law not because she believes that the French should work longer hours, but because she thought that employers had too much power and

[0]This chapter is joint work with Marcello Estevão (IMF). It has been presented at the 46th Panel Meeting of **Economic Policy** in Lisbon. We are grateful to Joshua Angrist, David Autor, Olivier Blanchard, and the referees for insightful comments. We thank Georgios Panos for help with data. Paul Bedford, Giuseppe Bertola, Dora Costa, Luc Everaert, Giuseppe Moscarini, Andrei Shleifer and seminar participants at MIT, the IMF, the Yale School of Management Behavioral Science Conference, the 9th IZA European Summer School in Labor Economics, the Bank of England, the EALE Conference 2006, and the Paris School of Economics provided valuable suggestions. Filipa Sá acknowledges financial support from the Portuguese Science and Technology Foundation through a PhD fellowship.

flexibility in implementing the shorter workweek. On the other hand, the victorious centre-right candidate, Nicolas Sarkozy, sees the 35-hour workweek as 'the worst mistake France has ever made'. During the campaign, he promised to relax the policy by exempting all overtime from payroll charges and income tax in order to encourage workers to work longer hours and increase their income. He believes that the short workweek has hurt France's competitiveness, making it less attractive to foreign investors.

The purpose of this paper is to analyse the effect of the 35-hour workweek on employment and workers' welfare. Our findings shed light on some of the issues raised during the presidential campaign and also relate to a broader debate on the reasons why Europeans work less than Americans. While in the 1960s annual hours per person employed were about the same in Europe and in the United States, today the average American works approximately 400 hours more per year than the average European. The reasons for this difference in working hours have been a topic of recent discussion, with some studies attributing it to differences in tax rates, while others highlight differences in preferences for work and leisure or in labour market regulations and unionization between the two continents.[1]

After discussing the main theories about how reductions in working hours may increase employment and welfare in section 2.2, we look specifically at the effects of the 35-hour workweek in France in section 2.3. We start by describing the institutional features of the laws which implemented the shorter workweek. We then explain the research strategy used and discuss it in the context of the existing literature on the impact of the shorter workweek in France. We describe the dataset and present the empirical results. Section 2.4 contains a brief description of other studies on work-sharing experiments in Germany, Sweden and Canada. Finally, section 2.5 offers brief policy conclusions.

2.2 The Rationale Behind Hours Reductions

There are two arguments commonly used to justify reductions in working hours. First, reducing working hours may create more jobs through work sharing. Second, it may improve workers' welfare by increasing their leisure time. This section discusses these arguments. It also provides

[1] Key references in this debate are Alesina et al. (2005), Blanchard (2004), and Prescott (2004).

predictions for how key economic decisions may be affected by institutionally-driven restrictions in working hours. These predictions are tested in section 2.3 using the experiment of the 35-hour workweek in France.

2.2.1 Job Creation Through Work Sharing

The idea of work sharing as an employment creation policy is simple: if the production of goods and services in an economy is fixed, then a reduction in hours can re-distribute the fixed amount of work across more people, increasing employment. In spite of its intuitive appeal, economists and policy makers are sceptical about the success of work sharing as it is rooted in the so-called "lump-of-labour fallacy": the false premise that the amount of output in the economy is fixed.

The theoretical literature on work sharing suggests that the employment effect of a reduction in hours is ambiguous and may actually be negative. These basic effects can be shown with the help of a simple model. Let us assume perfect substitutability between hours (H) and workers (N) in the production function, given by: $Y = F(HN, K)$ and that, first, individuals do not work overtime. Employment has a fixed cost, a, which does not vary with hours of work (e.g. training, day-care provision, and other overhead costs). The hourly wage is w and the cost of labour is given by: $aN + wHN$. Let us also assume that output is fixed. Firms are initially employing N_0 workers, who work the standard number of hours $\bar{H}_0$. There is a mandated reduction in standard hours to $\overline{H}_1 < \overline{H}_0$. Because output is fixed, firms reduce hours to $\overline{H}_1$ and increase employment to keep output constant. The reduction in hours increases employment to $N_1 = (\bar{H}_0 N_0)/\bar{H}_1 > N_0$. Thus, with fixed output *and* no overtime, this model predicts that a reduction in hours of work per person unambiguously increases employment. This is the basic intuition behind the enthusiasm of policymakers with legislated workweek reductions.

However, as pointed out by Calmfors and Hoel (1988) and the subsequent literature following them, small modifications in this simple setup will make the final employment effect ambiguous. First, let us expand the cost of labour to include the possibility of overtime: $aN + wH_0N + w(1+o)(H - H_0)N$, where o>0 stands for the overtime premium. A reduction in the standard workweek leaves the marginal cost of an extra hour of overtime, $w(1+o)N$, unchanged but raises the marginal cost of hiring an additional worker, $a + wH_0 + w(1+o)(H - H_0)$, causing firms to hire more hours and fewer people. In addition, if firms are allowed to vary output,

there is a negative scale effect: facing a higher cost of production (because the fixed cost of labour is diluted over a smaller number of hours) firms decide to produce less, decreasing both employment and hours. The net effect on employment becomes even more ambiguous if we introduce another production input in the analysis, say "capital", which would raise the possibility of increased labour costs causing a substitution away from labour toward capital.

Houpis (1993), Kapteyn et al (2000 and 2004), and the references cited by them, show that the effectiveness of legislated working-time reduction to increase employment also hinges crucially on the response of wages. At the same hourly wage, a reduction in hours would decrease monthly income. If workers demand higher hourly wages as compensation for the reduction in hours, possibly keeping monthly income unchanged, the increase in costs may offset any potential positive effect of work sharing on employment. Kapteyn et al (2000 and 2004) test this hypothesis using aggregate data for the OECD. They find that a reduction in working hours has a positive direct effect on employment. However, taking into account the increase in wages, the long-run effect becomes small and insignificant.

Wage reactions to a legislated reduction in the workweek is indeed a crucial issue in the French case. The 35-hour workweek law guaranteed the monthly earnings of workers receiving the minimum wage ("salaire minimum interprofessionnel de croissance" or SMIC). For this group of workers, hourly wages increased to compensate them for the reduction in hours and keep their monthly income unchanged. This may reduce the employment effect of the reduction in hours.

2.2.2 Cooperation in the Presence of Positive Spillovers in Leisure

The theoretical results discussed so far limit the effects of the legislated reductions in the workweek to direct employment decisions without taking into account general equilibrium effects or the existence of market inefficiencies that could be corrected by government policies. Very few theoretical studies focus on these issues. Marimon and Zilibotti (2000) is one exception. Using a general equilibrium framework with search-matching frictions, they show that the conditions for obtaining even small employment effects from legislated workweek reduction are rather restrictive. In particular, to raise the effectiveness of the policy, productive factors which complement labor, such as capital, should not be able to adjust to the policy intervention. That

might explain why some proponents of legislated workweek reductions would like these policies to be implemented at the widest possible scale (e.g., at the European Union level). Moreover, the authors show that the output losses associated with these policies could be quite large. However, they do not take into account social coordination problems or the possibility that workers like restrictions on working time as a means to prevent employers from exploiting some type of yardstick competition mechanism to induce them to overwork.

We believe that these externalities could, in principle, provide a beneficial role to policy action. For instance, mandated reductions in the workweek may improve workers' welfare if there are positive spillovers in leisure, i.e. interdependencies between the choices of hours of different workers, with each worker being better off when others are working fewer hours. This may be the case if there is a 'rat race' in the workplace and individuals want to be viewed as hard working. In this case, workers are better off when their colleagues work fewer hours because they gain a comparative advantage in the 'rat race' and can improve their career prospects. The decentralized equilibrium, in which each worker chooses hours taking the other workers' choices as given, is characterized by inefficiently high hours. If workers could coordinate their actions and collectively choose lower working hours, they would achieve a better, cooperative, equilibrium. A law imposing an upper bound on hours may provide this type of coordination. This idea is discussed in Landers et al (1996), who describe an organizational setting with positive spillovers in leisure. They consider a situation in which there are two types of workers: those who prefer to work short hours and those who prefer to work long hours. Workers participate in the labour market for two periods. They may be promoted or not in the second period depending on their performance in the first period. Firms cannot observe a worker's type, but take hours as an indicator and promote in the second period those workers who have worked long hours in the first period. This generates a 'rat race' in which workers with a preference for short hours may have an incentive to work long hours in the first period in order to be perceived as hard-working types. To reduce this adverse selection and ensure that they promote workers who really prefer long hours, firms may find it optimal to increase hours of work, raising the cost that workers with a preference for short hours have to bear when disguising themselves as hard-working types. The authors show that this type of equilibrium is characterized by inefficiently long hours and find evidence in support of this conclusion using

data from two large law firms. Indeed, law firms use long hours as an indicator when deciding whether to promote associates to partners. As a result, associates have a tendency to work long hours to stay ahead of their colleagues. This situation is inefficient and workers could be better off with a coordinated reduction in hours.

More generally, the argument can be made using a simplified representation of the model proposed in Cooper and John (1988). First, assume that a representative worker i has an utility function given by $U(wH_i)+V(1-H_i,1-\hat{H})$, where H_i is the number of hours worked by worker i, $\hat{H}$ is the number of hours worked by all other workers, and w is the hourly wage. Utility is concave in consumption and leisure, $U_{11} < 0, V_{11} < 0$. Using this very simple setup, we can characterize two types of equilibrium:

- In a decentralized equilibrium, each worker maximizes his utility taking other workers' hours as given. We focus on symmetric equilibria, where everyone is working the same number of hours ($H_i = \widehat{H}$). The first order condition to the utility maximization problem in a symmetric decentralized equilibrium is $wU'(wH_i) - V_1(1 - H_i, 1 - H_i) = 0$.

- In a cooperative equilibrium, each worker internalizes the effect of his hours on other workers' utility. The first order condition in the symmetric cooperative equilibrium is $wU'(wH_i) - V_1(1 - H_i, 1 - H_i) - V_2(1 - H_i, 1 - H_i) = 0$. With positive spillovers in leisure $V_2(1 - H_i, 1 - H_i) > 0$ and the cooperative equilibrium is characterized by a lower level of hours than the decentralized equilibrium.

By not taking into account the effect of their working hours on other workers' utility, workers choose inefficiently high hours in the decentralized equilibrium. A coordinated reduction in hours (for example, imposed by law) would be welfare improving and this simple model illustrates the possible role of public policies in helping society to achieve a welfare-enhancing equilibrium.. However, in the absence of positive spillovers in leisure a mandated reduction in hours would be introducing a constraint and could make some workers worse off. For those workers who would, in the decentralized equilibrium, choose hours below the mandated upper bound, there is no welfare effect. But for those who would choose to work more, the law introduces a distortion and makes them worse off.

The simple notion of an externality produced by the free-market equilibrium workweek provides an interesting twist on the labor supply side of the economy, adding a welcome complexity to the labor demand-oriented discussion of the previous subsection. For instance, it is possible that, in the absence of strong externalities, workers decide to work overtime or look for a second job to avoid a reduction in income that could be associated with a reduced workweek, an effect not studied in previous empirical papers. If present in the data, such an effort to circumvent the intent of the law will certainly limit its effect on employment. In Section 2.3 we control for individual characteristics that can be associated with a preference for working longer hours. One of these characteristics is "gender", since it is plausible that the 'rat race' equilibrium described above may be more relevant for men than for women.

The model with leisure complementarities can be modified to illustrate another labour supply effect that would undermine the effectiveness of work-sharing policies. There might be direct complementarities between leisure and income such that the workers might not be able to enjoy longer leisure hours if they are cash constrained. In that case, leisure and income would not be separable in the utility function as sketched above and the cross derivative of the utility function would be positive, at least in the relevant range of incomes. In this modified setup, even in the absence of leisure externalities, workers that face a reduction in their workweek may get an additional job to supplement their income or bargain for increased hourly wages, instead of enjoying the extra leisure.

Finally, and very important for our analysis, the reduction in the workweek may lower the benefits of being a part-time worker vis-à-vis being a full-time worker. In this case, a legislated workweek reduction could induce some part time workers to become full time, which would limit the positive employment effects of the law. This effect could have been particularly important in France, as authorities ended a social security rebate associated to the hiring of part-time when the 35-hour workweek law was enacted. The possible changed incentives to work part time versus full time in France after the reduced standard workweek is discussed in Oliveira and Ulrich (2002). The authors find that part-time employees working between 20 and 30 hours in firms that reduced working time, saw their chances of getting a full-time position increase. To abstract from this type of effect, we limit the sample to full time employees in the empirical analysis in section 2.3.

2.3 The 35-Hour Workweek in France

2.3.1 Institutional Background

The workweek in France has been shortened by a sequence of laws since the early 1980s. In 1982, François Mitterrand's socialist government reduced the length of the workweek from 40 to 39 hours. In 1998, a new socialist government, led by Lionel Jospin, further reduced the workweek to 35 hours through two laws: Aubry I (June 1998), which set the length of the workweek at 35 hours beginning in February 2000 for firms employing more than 20 people and in January 2002 for smaller firms; and Aubry II (January 2000), which introduced more detailed legal provisions regarding overtime. Aubry I gave small firms more time to reduce the workweek as it was understood that it would be more difficult for them to put into practice the necessary changes to implement a shorter workweek. To ease that transition, the law reduced the overtime premium for small firms and increased their annual limit on overtime work compared with large firms. This way, small firms could continue operating on a 39-hour basis paying the difference with a reduced overtime premium.

The purpose of the workweek reduction was to create more jobs during a period of high unemployment (11.5 percent in 1997). Employees were expected to bear a small part of the cost of the working time reduction, continuing to earn roughly the same monthly income—in line with the unions' slogan "35 hours pays 39." To attenuate the negative effects on profitability, the government offered rebates on firms' social security contributions. The rebates declined with the employee's monthly income and were largest for workers receiving the minimum wage. In addition, unions accepted a more flexible accounting of overtime work from a weekly to an annual basis and the working time reduction was expected to increase productivity. The official argument was that productivity increases together with cuts in social security contributions might even lead to a reduction in labour costs, so that firms would not need to cut monthly wages for the policy to be sustainable. To protect low-wage individuals, the law guaranteed the monthly earnings of workers receiving the minimum wage.

Another noteworthy feature of these laws is the treatment given to managers. Recognizing the autonomy of their work, the Aubry laws gave them more flexibility in the negotiation of hours. Aubry II classified managers in three categories: managers integrated in a team,

autonomous managers (including researchers, engineers, investment bankers, etc), and directors. The first two categories of managers could sign an agreement with their employers to establish a regular work length in hours or days on a weekly, monthly or annual basis, with additional hours or days being paid at an overtime premium. Directors were fully exempt from the 35-hour workweek.

2.3.2 Research Strategy

We are interested in studying the effect of the 35-hour workweek along several dimensions, motivated by the theoretical discussion in section 2.2. As we have seen, theory does not make clear-cut predictions about the effects of reductions in hours on employment and welfare. Therefore, whether the 35-hour workweek succeeded in creating more jobs and in making French workers happier is ultimately an empirical question.

To measure the effect of the law, we use the fact that it was applied earlier in large firms to construct a natural experiment. Firms with more than 20 employees had to implement the 35-hours workweek by February 2000, while firms with less than 20 employees had until January 2002 to do so. Therefore, we use workers in large firms (20 to 49 employees) as the treatment group and workers in small firms (less than 20 employees) as the control group.

The effect of interest is captured by the difference between the outcome of the treatment group after the law and before the law, and the corresponding difference for the control group – the difference-in-differences (DD) estimator. Formally, we use the following specification to estimate the effect of the treatment on outcome variable y_i for individual i:[2]

$$y_{it} = \alpha + \gamma d_t + \beta X_{it} + \delta_0 treatment_i + \delta_1 (treatment \times post)_{it} + \epsilon_{it}$$

where y_{it} is the outcome of interest for individual i at year t. X_{it} is a set of controls capturing observable differences in the characteristics of the control and treatment groups that affect the outcome of interest. $d_t$is a vector of year effects controlling for time-specific changes in y_{it}. $treatment_i$ is an indicator equal to 1 if the individual belongs to the treatment group and 0 if she belongs to the control group. This variable captures time-invariant differences between

[2]For a more technical and comprehensive discussion of this and other strategies for identifying causal effects, see Angrist and Krueger (1999).

the two groups.

$(treatment \times post)_{it}$ is the interaction of the indicator for treatment with an indicator equal to 1 after the treatment was implemented. The coefficient δ_1 is the DD estimator measuring the causal effect of the law on the outcome variable.

This strategy requires fairly weak identifying assumptions to be valid. The key identifying assumption is that there are no contemporaneous shocks, other than the treatment, affecting the outcomes of the control and treatment groups during the period we analyze. If this assumption holds, the behaviour of individuals in the control group gives information about how individuals in the treatment group would have behaved if they had not been treated.

There is an element which may bias our results. If the business cycle affects the control and treatment groups differently, our estimates may be capturing the effect of the business cycle instead of the effect of the treatment. This is important in the case of the 35-hour workweek because it was implemented during a period when the French economy was booming. If the boom affects the control and treatment groups differently, the DD estimator will be a biased measure of the effect of the law. To control for this possible bias, we follow the strategy in Kugler and Pica (2005) and estimate the following alternative specification, in which $GDP_t$is the growth rate of real GDP and the extra interaction term captures group-specific business-cycle effects:

$$y_{it} = \alpha + \gamma d_t + \mu(GDP_t \times treatment_i) + \beta X_{it} + \delta_0 treatment_i + \delta_1(treatment \times post)_{it} + \epsilon_{it}$$

Throughout, we report least-squares estimates, but probit or logit estimates give similar results. In all estimations, standard errors are clustered by year $\times$ treatment group cells. By clustering the standard errors we correct for the fact that there may be common errors within these groups. Failure to take this into account would lead to underestimation of the standard errors and overestimation of the effect of interest, as shown in Moulton (1990).

2.3.3 Our Approach in the Context of the Literature on Hours Reductions in France

The approach used in this paper differs from the existing literature on the effects of reductions in hours in France both in terms of methodology and scope of the analysis.

Several studies in France estimate the effect of the 35-hour workweek on employment. They can be categorized in two groups: ex-ante simulations, which predict the effect of the 35-hour workweek before it was implemented; and ex-post evaluations. A survey of some of these studies can be found in Gubian et al (2005).

Ex-ante simulations are based on either macro or micro models. Macro models make assumptions about the evolution of productivity, hourly wages, and capital usage following the reduction in hours and predict the evolution of employment under these assumptions. An example of this type of studies in Dares-BDF-OFCE (1998), which found that a reduction in the workweek from 39 to 35 hours could potentially generate 700,000 additional jobs. Micro models specify a functional form for the production function and endogeneize the evolution of wages using hedonic models. Hedonic models explain wages by looking at the different characteristics of jobs, including hours of work. Gaps between the workers' desired and actual hours lead to an increase in wages to compensate them for not working their desired hours. Micro models tend to predict more modest effects on employment than macro econometric simulations, because wages would need to increase to compensate workers for working shorter hours. This increase in labour costs reduces the positive impact of the reduction in hours on employment.

Ex-post evaluations follow an approach more similar to ours and compare firms that reduced hours (treatment group) with those that did not (control group). However, some of these evaluations have methodological flaws and do not control for differences in characteristics between the two groups. Others apply econometric techniques to account for those differences, which is important to ensure that the control group is a good counterfactual for the behaviour of the treatment group in the absence of the reduction in hours.[3] Most of these studies find a strong positive effect on employment. A potential problem is that they ignore small firms. They restrict the analysis to large firms and compare those that reduced hours earlier with those that

[3] An example of this type of ex post evaluations is the study by Crépon et al (2005).

did it later. By leaving out small firms, these studies ignore potentially useful information. Moreover, by dividing large firms into two groups depending on the timing of implementation of the shorter workweek, these studies generate a bias: both groups of large firms were affected by the law and their decision to implement it sooner or later could be related to unobservable variables (e.g. productivity). The results would be capturing the effect of those unobservable differences rather than the effect of the shorter workweek.

The study that is closest to ours in terms of methodology is Crépon and Kramarz (2002). It analyses the effect on transitions from employment to unemployment of the earlier law that reduced the workweek from 40 to 39 hours in 1982. The authors explore the variation in hours worked to design a 'natural experiment', comparing workers who were already working less than 40 hours before the law (control group) with those who were working 40 hours or more (treatment group). They find that the reduction in hours increased the probability of making a transition from employment to unemployment between 2.3 and 3.9 percentage points. This is an indication that the reduction in hours may have reduced employment, even though the authors do not attempt to estimate the effect on the *level* of employment, i.e., taking into account also possible transitions from unemployment to employment.[4]

The effect of the 35-hour workweek on workers' welfare has received much less attention than the effect on employment. Most studies are based on surveys asking employees affected by the law whether they consider that their situation has improved as a result of the reduction in hours. Typically, the majority of employees say that their situation has improved. However, because these studies do not compare the outcomes of treated individuals to the outcomes of a control group, they fail to capture any causal effect.[5]

In this paper we take a broad approach and look at the effects of the workweek reduction on

[4] The same approach could be applied mechanically in the evaluation of the 35-hour workweek law by defining the treatment group as individuals working more than 35 hours per week before the law was announced and the control group as individuals working less than 35 hours before the law was announced. However, while the difference between having a workweek of 40 hours or more or 39 hours or less could be viewed as marginal in terms of the type of individual who would choose one schedule over the other, the same cannot be argued about the two types of individuals separated by the most recent law. Indeed, individuals working less than 35 hours per week can be significantly different from those working 39 hours or more. In particular, they may have a stronger desire to work part-time or have a more flexible schedule. The dynamics of part-time choice and the decision to hold more than one job when the workweek in the main job becomes significantly smaller may generate behavioral changes that complicate the identification of the effects studied in this paper. We will return to this important issue later.

[5] An example of this type of study is Cette et al (2005).

several margins closely related to workers' welfare. We use information on wages, dual-job holdings, employment, worker transitions from large to small firms, and workers' satisfaction with their working hours to have a fuller assessment of the effects of the law.

2.3.4 Data and Descriptive Statistics

To implement our empirical strategy, we need information on labour market outcomes and on firm size, for a period of time spanning the implementation of the 35-hour workweek in large and small firms. We use data from the French labour force survey (Enquête Emploi) from 1993 to 2000. This survey is conducted in March of each year, with the exception of 1999 when it was done in January. It has information on several demographic characteristics, as well as on labour market status, wages, hours of work, tenure, etc. The sample is renewed by a third every year, so the same individual can be followed for three consecutive years. This dataset is matched with firm-level data from the French Registry of Firms (SIRENE), containing information on firm size.

The classification of workers into the control and treatment groups needs to be done *before* the law was enacted. This is because workers may move from large to small firms (or vice-versa) as a result of the law and firms just above the 20-employees cut-off may reduce size to avoid having to implement the shorter workweek earlier. These behavioural responses would bias our results. For this reason, we exploit the fact that the same individual can be followed for three consecutive years and construct several three year panels: 1993-1995, 1994-1996, 1995-1997, 1996-1998, 1997-1999, and 1998-2000. The period before the law covers the years 1993 to 1998 (since the law was announced in June 1998 and the Enquête Emploi was conducted in March 1998) and the period after covers 1999 and 2000. Individuals are classified in the base year, which always falls within the period before the law was enacted. It is possible that employees do not change behaviour in 1999 as they may not have had enough time to adjust. This would bias our results towards underestimating the effect of the law. We take a conservative approach and see if we find a significant effect already in 1999.[6]

[6] Notice that data characteristics limit how far in time our analysis can go. All information after June 1998 is tainted by the possible endogenous response of economic agents to the law and, thus, cannot be used to "control" the start of the experiment. Because we can only follow people for three years, the experiment proposed here can only produce a snapshot of the effect of the law in March 1999 and March 2000. Further research needs to

Table 1 reports descriptive statistics by hours worked, firm size and year for key individual characteristics. The sample is limited to employees age 15 to 64 with positive net monthly income, and excludes self-employed individuals. Small firms have less than 20 employees. Large firms have between 20 and 49 employees. We impose a cut-off at 49 to ensure more homogeneity between the two groups, so we can be more confident that the behaviour of small firms, after controlling for individuals' observable characteristics, can be taken as the counterfactual for the behaviour of large firms in the absence of the law. We experimented with different cut-off levels and the results are not sensitive to the cut-off choice. The table shows some differences between workers in large and small firms: workers in large firms earn higher wages, have longer tenure, and work fewer hours in the second job (when they have one) than workers in small firms. There are also differences between individuals working longer and shorter hours: individuals working longer hours tend to be male, earn higher wages, have longer tenure, and work fewer hours in the second job than those working shorter hours. Our estimations control for these differences in observable characteristics.

2.3.5 Results

Hours Distribution and Wages

The introduction of the 35-hour workweek had a clear impact on weekly hours of work, as shown in Table 2. In 1999 most employees were still working 39 hours. After that, there was a big increase in the proportion of employees working 35 hours, led by large firms. In 2002 more than 45 percent of employees in large firms and almost 35 percent of employees in small firms worked on a 35-hour basis.

To examine more carefully the effect of the law on hours worked, we use the DD estimator with employees in small firms as the control group and employees in large firms as the treatment group. We report the results in table 3 using a sample in which part-time workers are excluded to facilitate the interpretation of future results. Results including part-time workers have the same flavour and, actually produce parameter estimates with smaller standard errors. All estimates control for possible interactive business cycle effects. The DD estimators are negative,

be designed to study the effects of the law beyond that.

suggesting that the law reduced hours, but only significantly for women. However, even for women the impact of the reduction of the workweek on hours seems surprisingly small (around 25 minutes, compared with the legal reduction of 4 hours). One reason for this is that working time includes not only regular time but also overtime hours. It may happen that firms reduce regular hours and increase overtime, reducing the effect on total hours.

Table 4 shows the results of applying the same technique to measure the effect on hourly wages and monthly incomes in the main job, again studying the separate effects on men and women. We also present results separately for workers receiving between 10% below and 10% above the minimum wage and workers receiving between 10% and 80% above the minimum wage. We are interested in this distinction because the law mandated that the monthly earnings of workers receiving the minimum wage should stay constant.[7] The results suggest that the law led to an increase in hourly wages for men working at and above the minimum wage. In contrast, the average hourly wage earned by women remained constant (even declined a bit in the group of workers earning more than the minimum wage). Monthly income is estimated to have remained the same after the introduction of the law for all different groups, with some chance that it actually declined for women earning more than the minimum wage (although, the parameter estimate has a high standard error).[8]

The increase in hourly wages for men has implications for the work-sharing hypothesis. We have seen in the theoretical discussion that an increase in hourly wages makes it less likely that a reduction in hours increases employment. This is because the increase in labour costs induces a scale effect, with a negative impact on both hours and jobs. As we shall see, this is indeed the case: the flow to unemployment of employed men in large firms increased significantly after the 35-hour workweek was enacted. Women employed in large firms did not face an increase in their chances of losing a job.

[7] We exclude workers with earnings more than 80% above the minimum wage as we want to ensure more homogeneity within the group. Moreover, as we have discussed, the law gave more flexibility in the negotiation of hours to workers in managerial positions. These workers are more likely to be at the top of the earnings distribution.

[8] The results are insensitive to estimation with individual fixed effects to account for unobservable differences among individuals.

Dual-Job Holdings

A margin of adjustment which sheds light both on the work-sharing hypothesis and on the different supply-side hypotheses on workweek choice (externality in leisure or complementarity between income and leisure) is the proportion of workers with more than one job. If individuals respond to the reduction in hours by working in a second job, it is less likely that unemployment will decline, because some jobs will be filled by individuals who are already employed elsewhere. At the same time, the fact that workers are looking for a second job suggests that they have a stronger preference for income relative to leisure. In that case the reduction in hours would not be working, for instance, as a coordination mechanism to encourage workers to increase their leisure time.

When excluding part-time workers, Table 5 shows that the law increased dual job holdings by 1.2 percentage points among women but not among men. If part-time workers are included, the probability of holding an extra job actually seems to decline after the law. This is consistent with the effect discussed in the theoretical section: once the workweek is reduced, part-timers find easier to become full-timers and, thus, some abandon their second job.

So far, the results presented tell a consistent story: full-time men in large firms have not reduced much their workweek right after the law was implemented, have got higher hourly wages, and about unchanged monthly income, therefore, having no reason to take up a second job. On the other hand, women have reduced their workweek by more than men and possibly began earning smaller monthly incomes after the law was enacted. Thus, they had a higher incentive (and more free time) to accept a second job. If men's hourly wages and women's dual-job take-up had remained unchanged, arguments for strong preferences for lower hours (that could not be attained by the market because of the externalities discussed here) or for the absence of relevant complementarities between leisure and cash would be more persuasive.

Employment

To test the work-sharing hypothesis directly, we are interested in the effect of the 35-hour workweek on employment. We look both at flows in and out of employment and at the level of employment.

To measure the effect of the law on transitions from employment to unemployment, we

restrict the sample to employees working more than 35 hours before the law as those were the ones for whom the law was binding. We compare the probability of becoming unemployed for workers initially working at large firms (treatment group) relative to workers initially working at small firms (control group). The results are reported in Table 6. The law increased transitions from employment to unemployment for men by 2.8 percentage points, while it did not affect transitions for women. This is consistent with the results found so far: as hourly wages increased for men, they became more expensive, giving firms an incentive to fire them. Women, on the other hand, accepted the same (or even a slightly reduced) hourly wage, which did not put additional cost pressures on firms.

To measure the effect of the law on transitions in the reverse direction, i.e. from unemployment to employment, we test whether unemployed workers are more likely to find a job in large firms or in small firms as a result of the reduction in hours. Table 7 reports the results. The first column reports the log of the probability of working at a large firm at $t+2$ relative to being unemployed. The second column reports the equivalent relative probability for small firms. We are interested in comparing the coefficients on $unemployed \times post1999$ in the two columns. The comparison of these two coefficients tells us whether unemployed individuals are more likely to find a job at a large firm or at a small firm after the 35-hour workweek was implemented.

The results suggest that after 1999 there was an increase in the probability of finding a job, which was larger in large firms than in small firms for women (although, we reject the equality of the coefficients on $unemployed \times post1999$ only at the 13 percent level of significance) but not for men. Again, this is consistent with the previous estimation, which showed no increase in hourly wages for women in firms that reduced the workweek (large firms).

The evidence on transitions in and out of employment suggests that the 35-hours workweek changed the composition of the labour force, with large firms letting some more expensive workers (men) go. The group that has not earned increased hourly wages (women) has not faced higher transitions to unemployment and had a marginal increase in flows into jobs. This is strong evidence of the importance of wage reactions to determine the final effect of workweek reduction laws on employment, a result consistent with previous findings in the literature.

However, to measure the impact of the law on employment it is important to look not only at transitions in and out of employment but also at the *level* of employment. Unfortunately,

the data do not allow measuring this effect in a clear way. If we were to apply the DD strategy to study the effect of the law on the level of employment, we would like to have treatment and control groups which are composed of both employed and unemployed individuals. Then, we would test whether the fraction of employed individuals increased by more for the treatment group than for the control group after the law. The problem is that the source of variation that we can explore to define the treatment and control groups (variation in firm size) implies that, by definition, all individuals in the two groups are employed. Indeed, employees in large and small firms are all, by definition, employed. Thus, we do not have suitable treatment and control groups.

In spite of this technical difficulty, we can still shed light on the effect of the 35-hour workweek on the level of employment by comparing the evolution of employment in small and large firms. To do so, we look at how the probability of working in large firms relative to being unemployed changed after the implementation of the reduction in hours comparing with the equivalent probability for small firms. Table 7 reports the results. The probability of being employed relative to being unemployed increased after 1999 in large and small firms by approximately the same amount, as the two coefficients on *post*1999 are not statistically different for neither men nor women (although the difference is a bit larger in the second case). Thus, even though employment increased after the law, it did not increase more in large firms relative to small firms. This suggests that the law did not increase employment, at least by 2000. An alternative way of making the same point is to plot the log of the probability of working in a large or in a small firm divided by the probability of being unemployed. Figure 1 shows that the log odds of employment by firm size are essentially parallel, suggesting that the 35-hour workweek had no effect on the level of employment.

Transitions from Large to Small Firms

Transitions of workers from firms affected by the law in the first instance (large firms) to firms where the adjustment in hours was delayed (small firms) could provide some evidence on whether there are positive spillovers in leisure or complementarities between leisure and cash: in the absence of these effects we should see an increase in these transitions as individuals try to circumvent the restrictions imposed by the 35-hour workweek. This may not be the case if

workers anticipate that the 35-hour workweek will also be implemented in small firms with a two-year lag and decide that it is not worthwhile to move. Nevertheless, this is a margin of adjustment we should explore. To do that, we look at employees working in large firms and compare the transitions to small firms of those who were working more than 35 hours (treatment group) relative to those who were working 35 hours or less (control group) before the law. The results are reported in Table 9. The law had no effect on transitions from large to small firms for men. For women, there is evidence of a small increase (0.4 percentage points), but the result is not highly significant (p-value of 13%).

Satisfaction with Hours of Work

Just as we measured the direct effect of the 35-hour workweek on employment, we would like to have a direct measure of its effect on workers' satisfaction. One way of doing this is to look at subjective measures of satisfaction with hours of work. Because this is not the type of information that can be found in labour force surveys, we need to use an alternative data source.

We use data from the French section of the ECHP, Waves 1 to 8, covering the years 1994 to 2001. This survey has a panel structure, following the same individuals over time. It has information on satisfaction with hours of work and with the amount of leisure time.[9] It also has information on firm size, which allows us to classify individuals into the treatment group if they were working in a large firm in the last survey before the law was implemented and 0 if they were working in a small firm.

The results, reported in table 10, suggest that lower working hours did not make workers happier with their work and leisure times. The workweek reduction reduced satisfaction with working hours for both men and women. Men also became less satisfied with their amount of leisure time.

[9] Satisfaction with hours of work and with the amount of leisure time is measured on a scale from 1 to 6. There is some skepticism about the quality of subjective measures of happiness and satisfaction. But, in most cases, friends and colleagues of the individual give ratings which are strongly related to the way people rate themselves, as reported in Diener and Suh (1999). This is reassuring and suggests that these measures provide useful information on individuals' welfare.

2.3.6 Summing up

To close this section it is useful to take stock of the empirical evidence. We found that the 35-hours workweek had the following effects:

- hours of work decreased less for men than for women. This may be explained by men choosing to work more overtime.
- hourly wages increased for men (probably due to an overtime premium). There is no evidence of an increase in wages for women.
- transitions out of employment increased for men as they became more expensive and did not change for women.
- the overall level of employment does not seem to have been affected.
- dual-job holdings increased for women, suggesting a desire to work more than the mandated number of hours.
- and satisfaction with working hours and the amount of leisure time did not increase

These findings tell a consistent story and provide strong evidence against the work-sharing hypothesis and some hints that positive spillovers in leisure are not important and that complementarities between income and leisure may exist. The effect on hours of work suggests that the law was indeed binding, even if the actual reduction in hours was substantially smaller than the amount legislated (because some workers, specially men, might have chosen to work overtime). The evidence on hourly wages and on dual job holdings suggests that there are other margins of adjustment, which reduced any potential positive impact of the reduction in hours on employment. In particular, the importance of wage reactions to determine the impact of workweek reduction laws is clear in our results: negative employment effects exist only for the group that has obtained an increase in hourly wages (men). The group that has not seen increased hourly wages (women) looked more actively to obtain a second job, an alternative way to keep total hours worked and monthly income unchanged, instead of enjoying the additional leisure time.

The evidence on employment and on satisfaction with working hours provides more direct tests of the work-sharing hypothesis and of the hypothesis of positive spillovers or complementarities between income and leisure, respectively. The evidence on employment shows that the law changed the composition of the labour force, with no evident impact on the overall level of employment. The evidence on satisfaction with working hours suggests that there are no positive spillovers in leisure and, on average, workers may need more income to enjoy extra leisure. This is, to some extent, not surprising as we have seen that positive spillovers arise from a 'rat race' in the workplace. This type of 'rat race' is likely to be more important for professions with high earnings and high responsibility, such as managerial positions. But, these professions (e.g. lawyers, investment bankers, engineers, researchers, etc) had a large degree of flexibility in adjusting their hours, with directors being completely exempt from the 35-hour workweek. Thus, it is not surprising that the evidence for the existence of positive spillovers is not strong. The professionals who could have benefited more from a coordinated reduction in hours were the ones least affected by it.

2.4 Work-Sharing Experiments in Other Countries

The evidence presented in the last section suggests that the 35-hour workweek failed to increase employment and make workers happier with their hours of work. To what extent can this negative evaluation be extended to work-sharing policies in other countries?

The literature evaluating the success of work-sharing experiments is not very extensive. In spite of that, there are studies applying empirical strategies similar to ours to evaluate work-sharing experiments in Germany, Sweden and Canada.

Germany started reducing standard hours in 1985 on an industry-by-industry basis, with the purpose of raising employment. Hunt (1998 and 1999) exploits the cross-industry variation in standard hours' reductions to study their impact on actual hours worked, wages and employment. She finds that the reductions in hours decreased employment in the period 1984–94. This may be explained by the large increase in hourly wages, which was enough to offset the decline in actual hours worked and keep monthly incomes unchanged.

Sweden reduced working time by 5% for a particular class of shift workers. This reduction

happened gradually in the mid-1980s. Skans (2004) evaluates this policy comparing workers affected by the reduction in hours to workers unaffected by it. He finds that there was little implementation of the working-time reduction, with actual hours falling by only about 35% of the reduction in standard hours. Hourly wages increased sharply, almost enough to offset the loss in earnings due to the decline in actual hours. He does not look at the effects on employment.

In Canada, there was a mandated reduction of the workweek from 44 to 40 hours in Quebec between 1997 and 2000 with the aim of increasing employment. This reduction applied only to non-unionized hourly paid workers. Skuterud (2007) looks at the effect of this policy on employment. He uses both DD and DDD estimation. In the DD estimation he compares non-unionized hourly paid workers in Quebec (treatment group) to non-unionized hourly paid workers in Ontario, where there was no reduction in the workweek (control group). In the DDD estimation he exploits another source of variation to control for province-specific factors unrelated to the work-sharing policy. He compares non-unionized hourly paid workers (treatment group) to non-unionized salaried workers (control group) in Quebec (experimental province) and in Ontario (non-experimental province). In other specifications, he also exploits cross-industry variation, comparing industries where hours of worked were affected relatively more to those affected relatively less. His findings suggest that the policy failed to increase employment either at the provincial level or within industries that were affected relatively more.

The evidence for Canada is particularly striking because its work-sharing experiment has a set of conditions which are particularly suitable for it to succeed in creating more jobs. Skuterud lists five conditions. First, Canada has a less regulated labour market than most European countries. Therefore, it is less likely that unions or the government will impose full wage compensation for the reduction in hours. Second, the reduction in hours was applied to non-unionized hourly paid workers, who are disproportionately unskilled and have high rates of unemployment. They have little bargaining power in wage negotiations and so are unlikely to be able to obtain wage compensations. Third, fixed costs of unemployment for these workers (training costs, benefits, etc) are low, reducing the importance of the scale effect of the reduction in hours. Fourth, the difference in skills between these workers and the unemployed is small, facilitating the substitution between hours and jobs. Finally, survey evidence suggests that

these workers have a preference for shorter hours. So, they are less likely to look for a second job. Even with all these suitable characteristics, work sharing failed to create more jobs for these workers.

The cross-country evidence for Germany, Sweden and Canada is consistent with our findings for France and confirms that work-sharing policies have little, if any, potential to increase employment.

2.5 Policy Implications

Our analysis suggests that the 35-hour workweek in France failed to create more jobs and promoted a series of behavioural changes suggestive that many workers were less happy with their working hours. Available evidence for other countries, such as Germany, Sweden and Canada, also shows that reductions in hours failed to increase employment.

The fallacy underlying the intuitive and simple idea of work sharing is that it fails to consider additional margins of adjustment other than hours and jobs. In particular, it seems too strong to assume that hourly wages will remain unchanged by the law. Even for groups of workers to whom this assumption applies, other margins of adjustment, like holding an extra job, could decrease the potential of legislated workweek reductions to create more jobs. At the same time, the idea of positive spillovers in leisure, motivated by a 'rat race' equilibrium in the workplace, seems to be relevant in some professions, in particular the ones where long hours and high salaries are the rule. But, because the shorter workweek was applied much more flexibly to workers in managerial positions, with directors being fully exempt, its potential to provide a mechanism for offsetting the 'rat race' equilibrium was largely diminished.

Our results suggest that the 35-hour workweek should be discontinued and workers and firms should be free to choose the length of the workweek. This would eliminate the costly side effects of the law (in terms of individuals' welfare) that we examined in this paper, such as the increase in movements in and out of employment and in dual job holdings. Relaxing the law along the lines initially proposed by Nicolas Sarkozy, exempting overtime from payroll charges and income tax, is unlikely to yield the same results as discontinuing the law altogether, at least for two reasons. First, overtime hours are paid with a premium. It is not clear whether

the tax savings by firms will be enough to bring the cost of overtime hours in line with the cost of regular hours. Second, the two policies (discontinuing the law or relaxing it by exempting overtime from taxes) are clearly not equivalent from the point of view of the government budget.

Table 1: Summary Statistics

	35 Hours or Less				More than 35 Hours			
	Large firms		Small firms		Large firms		Small firms	
	1993-1997	1998-2000	1993-1997	1998-2000	1993-1997	1998-2000	1993-1997	1998-2000
Percentage female	85.460	79.279	85.162	79.447	34.072	32.602	35.999	33.800
	(35.257)	(40.567)	(35.550)	(40.420)	(47.397)	(46.886)	(48.000)	(47.308)
Average net monthly income (francs)	5671.961	5036.295	4578.945	4234.521	7922.348	9450.667	8526.168	9094.425
	(21667.760)	(5036.926)	(18968.900)	(2621.242)	(4645.550)	(36615.400)	(31722.900)	(36429.050)
Percentage with tenure less than 1 year	18.600	13.694	20.077	20.207	8.378	10.269	14.201	16.567
	(38.918)	(34.409)	(40.060)	(40.166)	(27.707)	(30.362)	(34.907)	(37.182)
Average hours in primary job	23.920	23.955	22.703	23.069	40.564	40.541	41.190	40.992
	(7.032)	(7.561)	(7.401)	(7.327)	(4.436)	(4.644)	(5.731)	(5.270)
Average hours in second job	0.704	0.845	0.925	1.110	0.125	0.075	0.136	0.108
	(3.588)	(3.888)	(4.066)	(4.535)	(1.680)	(1.277)	(1.714)	(1.185)

Notes: Table reports group means by hours worked, firm size and time period. Standard errors are reported in parentheses. The sample is limited to employees under age 64 with positive net monthly income and excludes self-employed individuals.
Source: French labour force survey (Enquête Emploi), 1993-2000.

Table 2: Distribution of Usual Weekly Hours (%)

	Small firms					Large firms				
	Above 39	39	35-39	35	Below 35	Above 39	39	35-39	35	Below 35
1999	21.49	49.02	2.25	2.46	24.78	16.44	47.43	8.66	5.48	15.88
2000	21.39	44.64	2.2	7.2	24.57	14.46	30.66	8.64	24.43	15.76
2001	21.7	41.33	3.27	11.86	21.82	13.49	18.9	8.47	35.68	13.53
2002	17.35	21.54	4.75	33.73	22.62	13.38	9.99	9.71	45.42	15.42

Source: Enquête Emploi, 1993-2000.

Table 3: Usual Weekly Hours

	Men	Women
Treatment	-0.594***	-0.243*
	(0.100)	(0.137)
Treatment*Post 1999	-0.170	-0.443*
	(0.104)	(0.247)
Married	0.375***	-0.349***
	(0.085)	(0.087)
Children under 6	0.040	-0.559***
	(0.052)	(0.115)
Age	0.146***	-0.047*
	(0.027)	(0.025)
Tenure less than 1 year	-0.144	-0.419***
	(0.155)	(0.139)
N	25642	15292

Notes: Robust standard errors adjusted for clustering by year × treatment. Regressions include year dummies, region effects, education, age-squared, and GDP× treatment. Sample excludes part time workers. *** significant at the 1% level, ** significant at the 5% level, * significant at the 10% level.
Estimated equations:

$$y_{it} = \alpha + \gamma d_t + \beta X_{it} + \delta_0 Treatment_i + \delta_1 (Treatment \times post1999)_{it} + \varepsilon_{it}$$

y_{it} is usual weekly hours. d_t is a set of time dummies. $post1999_t$ is an indicator equal to 1 if year t corresponds to 1999 or after and 0 otherwise. X_{it} is a vector of control variables. $Treatment_i$ is an indicator equal to 1 if the worker works at a large firm at time t and 0 if he works at a small firm.
Least squares estimates.
Sources: Enquête Emploi, 1993-2000; authors' calculations.

Table 4: Wages and Monthly Income

	Log Hourly Wage		Log Monthly Income	
	Men	Women	Men	Women
A. Workers between 10% below and 10% above minimum wage				
Treatment	0.010***	0.014**	-0.002	-0.001
	(0.002)	(0.005)	(0.002)	(0.003)
Treatment*Post 1999	0.010***	-0.006	0.002	0.001
	(0.001)	(0.007)	(0.002)	(0.002)
Married	0.009**	0.004	0.008***	0.005*
	(0.003)	(0.005)	(0.002)	(0.003)
Children under 6	-0.006**	0.003	-0.003**	-0.001
	(0.002)	(0.003)	(0.001)	(0.002)
Age	0.005***	0.004***	0.003***	0.003***
	(0.001)	(0.001)	(0.0004)	(0.0007)
Tenure less than 1 year	-0.005	-0.002	-0.005**	0.001
	(0.004)	(0.007)	(0.002)	(0.004)
N	7640	3794	7641	3794
B. Workers between 10% and 80% above minimum wage				
Treatment	0.021***	-0.016***	0.008**	-0.006***
	(0.002)	(0.005)	(0.003)	(0.002)
Treatment*Post 1999	0.014**	-0.016*	0.0002	-0.005
	(0.006)	(0.009)	(0.002)	(0.003)
Married	0.024***	0.025***	0.026***	0.007
	(0.005)	(0.007)	(0.004)	(0.004)
Children under 6	-0.003	0.010	-0.004**	-0.008**
	(0.002)	(0.007)	(0.001)	(0.003)
Age	0.019***	0.033***	0.013***	0.023***
	(0.001)	(0.002)	(0.001)	(0.002)
Tenure less than 1 year	-0.024***	-0.035**	-0.012*	-0.027***
	(0.007)	(0.015)	(0.006)	(0.008)
N	10602	5319	10603	5320

Notes: Robust standard errors adjusted for clustering by year × treatment. Hourly wage is monthly wage/(4.33×usual weekly hours). Section A. shows the results for a sample of employees under age 64 with net monthly income between 10% below and 10% above the minimum wage. Section B. shows the results for a sample of employees under age 64 with net monthly income between 10% and 80% above the minimum wage. Regressions include year dummies, region effects, education, age-squared, and GDP× treatment. Sample excludes part time workers.

Estimated equations:

$$y_{it} = \alpha + \gamma d_t + \beta X_{it} + \delta_0 Treatment_i + \delta_1 (Treatment \times post1999)_{it} + \varepsilon_{it}$$

y_{it} is log hourly wage or log monthly income. d_t is a set of time dummies. $post1999_t$ is an indicator equal to 1 if year t corresponds to 1999 or after and 0 otherwise. X_{it} is a vector of control variables. $Treatment_i$ is an indicator equal to 1 if the worker works at a large firm at time t and 0 if he works at a small firm.
Least squares estimates.
Sources: Enquête Emploi, 1993-2000; authors' calculations.

Table 5: Dual Job Holdings

	Including Part Time		Excluding Part Time	
	Men	Women	Men	Women
Treatment	-0.010***	0.002	-0.006***	0.004
	(0.001)	(0.002)	(0.001)	(0.002)
Treatment*Post 1999	-0.005*	-0.0003	-0.0008	0.012***
	(0.002)	(0.006)	(0.002)	(0.003)
Married	0.002	-0.024***	0.004***	-0.008***
	(0.002)	(0.003)	(0.001)	(0.002)
Children under 6	-0.001	-0.004**	-0.001	0.004***
	(0.001)	(0.002)	(0.001)	(0.001)
Age	0.002**	0.004***	0.001	0.001
	(0.001)	(0.001)	(0.0007)	(0.0007)
Tenure less than 1 year	-0.007**	-0.009	-0.002	0.003
	(0.002)	(0.007)	(0.002)	(0.003)
N	27263	24112	25935	14716

Notes: Robust standard errors adjusted for clustering by year × treatment. Regression includes year dummies, region effects, education, net monthly income, age-squared and GDP× treatment.
Estimated equation:
$y_{it} = \alpha + \gamma d_t + \beta X_{it} + \delta_1 Treatment_i + \delta_2 (Treatment \times post1999)_{it} + \varepsilon_{it}$
y_{it} is an indicator equal to 1 if the worker has a second job and 0 otherwise. d_t is a set of time dummies. $post1999_t$ is an indicator equal to 1 if year t corresponds to 1999 or after and 0 otherwise. X_{it} is a vector of control variables. $Treatment_i$ is an indicator equal to 1 if the worker works at a large firm at time t and 0 if he works at a small firm.
Least squares estimates.
Sources: Enquête Emploi, 1993-2000; authors' calculations.

Table 6: Transitions from Employment to Unemployment

	Men	Women
Treatment	-0.011**	0.006
	(0.004)	(0.006)
Treatment*Post 1999	0.028***	-0.003
	(0.005)	(0.005)
Married	-0.016**	-0.005
	(0.007)	(0.006)
Children under 6	0.003	-0.007
	(0.003)	(0.006)
Age	0.0003	0.0004
	(0.002)	(0.002)
Tenure less than 1 year	0.082***	0.074***
	(0.008)	(0.014)
Seasonal contract	0.003	0.174**
	(0.047)	(0.062)
N	8542	4779

Notes: Robust standard errors adjusted for clustering by year × treatment. Regression includes year dummies, region effects, education, age-squared, and GDP× treatment. The sample is limited to employees working more than 35 hours at time t (the first year in each panel).
Estimated equation:

$$y_{it+2} = \alpha + \gamma d_{t+2} + \beta X_{it} + \delta_0 Treatment_i + \delta_1 (large\ firm \times post1999)_{it+2} + \varepsilon_{it+2}$$

y_{it+2} is an indicator equal to 1 if the individual is unemployed at time $t+2$ (the last year in each panel). d_{t+2} is a set of time dummies. $post1999_{t+2}$ is an indicator equal to 1 if year $t+2$ corresponds to 1999 or after and 0 otherwise. X_{it} is a vector of control variables. $Treatment_i$ is an indicator equal to 1 if the worker works at a large firm at time t and 0 if he works at a small firm.
Least squares estimates.
Sources: Enquête Emploi, 1993-2000; authors' calculations.

Table 7: Transitions from Unemployment to Employment

	Men		Women	
	log[P(large)/ P(unemployed)]	log[P(small)/ P(unemployed)]	log[P(large)/ P(unemployed)]	log[P(small)/ P(unemployed)]
Unemployed	-2.875***	-2.576***	-2.462***	-2.253***
	(0.176)	(0.218)	(0.207)	(0.208)
Unemployed*post1999	1.878***1.891***		1.602***	1.449***
	(0.186)	(0.187)	(0.289)	(0.228)
Unemployed*unemployed less than a year	0.511***	0.447***	0.669***	0.501***
	(0.097)	(0.165)	(0.177)	(0.124)
Married	0.878***	0.806***	0.132**	0.171***
	(0.089)	(0.060)	(0.067)	(0.061)
Children under 6	-0.094	-0.054	-0.335***	-0.299***
	(0.084)	(0.091)	(0.091)	(0.069)
Age	0.106***	0.051**	0.059*	0.060**
	(0.033)	(0.026)	(0.036)	(0.031)
p-value for equality of coefficients on unemployed*post1999	0.893		0.126	
N	7016		7110	

Notes: Robust standard errors adjusted for clustering by year × unemployed in base year. Comparison group is the unemployed. Estimates give the change in the log odds ratio for a one-unit change in the independent variable. Regression includes year dummies, region effects, education and age-squared.
Estimated equation:

$$P(status_{it+2} = k) = \Phi(\alpha + \gamma d_{t+2} + \beta X_{it} + \delta_0 unemp_{it} + \delta_1 (unemp_{it} \times post1999_{t+2}) + \delta_2 (unemp_{it} \times unemp less 1 year_{it}))$$

$status_{it+2}$ takes three values (k=1, 2, 3): 1 if the individual is unemployed, 2 if he is employed at a large firm, and 3 if he is employed at a small firm at time $t+2$ (the last year in each panel). d_{t+2} is a set of time dummies. $post1999_{t+2}$ is an indicator equal to 1 if year $t+2$ corresponds to 1999 or after and 0 otherwise. X_{it} is a vector of control variables. $unemp_{it}$ is an indicator equal to 1 if the individual is unemployed at time t (the first year in each panel) and 0 otherwise. $unemp less 1 year_{it}$ is an indicator equal to 1 if the individual has been unemployed for less than one year at time t and 0 otherwise.
Multinomial logit estimates.
Sources: Enquête Emploi, 1993-2000; authors' calculations.

Table 8: Employment

	Men		Women	
	log[P(large)/ P(unemployed)]	Log[P(small)/ P(unemployed)]	log[P(large)/ P(unemployed)]	log[P(small)/ P(unemployed)]
Post1999	0.143***	0.153***	0.038	-0.010
	(0.038)	(0.033)	(0.038)	(0.031)
Married	1.162***	1.074***0.360***		0.452***
	(0.036)	(0.031)	(0.033)	(0.027)
Children under 6	-0.123***	-0.070***	-0.280***	-0.288***
	(0.027)	(0.024)	(0.030)	(0.023)
Age	0.165***	0.106***	0.095***	0.076***
	(0.010)	(0.008)	(0.010)	(0.008)
p-value for equality of coefficients on post1999	0.755		0.185	
N	46719		44425	

Notes: The comparison group is the unemployed. Estimates give the change in the log odds ratio for a one-unit change in the independent variable. Regression includes year dummies, region effects, education and age-squared.
Estimated equation:
$P(status_{it} = k) = \Phi(\alpha + \gamma post1999_t + \beta X_{it})$

$status_{it}$ takes three values (k=1, 2, 3): 1 if the individual is unemployed, 2 if he is employed at a large firm, and 3 if he is employed at a small firm at time t. $post1999_t$ is an indicator equal to 1 if year t corresponds to 1999 or after and 0 otherwise. X_{it} is a vector of control variables.
Multinomial logit estimates.
Sources: Enquête Emploi, 1993-2000; authors' calculations.

Table 9: Transitions from Large to Small Firms

	Men	Women
Treatment	-0.022***	-0.031***
	(0.007)	(0.004)
Treatment*Post 1999	0.0006	0.005
	(0.009)	(0.003)
Married	0.008**	0.007**
	(0.003)	(0.003)
Children under 6	0.002	-0.007*
	(0.003)	(0.003)
Age	-0.006**	-0.004
	(0.002)	(0.003)
Tenure less than 1 year	0.062***	0.074***
	(0.006)	(0.012)
N	18671	14563

Notes: Robust standard errors adjusted for clustering by year × treatment. Regression includes year dummies, region effects, education, net monthly income, age-squared, and GDP× treatment. The sample is limited to employees working at large firms at time t (the first year in each panel).
Estimated equation:

$$y_{it+2} = \alpha + \gamma d_{t+2} + \beta X_{it} + \delta_0 treatment_i + \delta_1 (treatment \times post1999)_{it+2} + \varepsilon_{it+2}$$

y_{it+2} is an indicator equal to 1 if the worker works at a small firm at time $t+2$ (the last year in each panel) and 0 if he works at a large firm. d_{t+2} is a set of time dummies. $post1999_{t+2}$ is an indicator equal to 1 if year $t+2$ corresponds to 1999 or after and 0 otherwise. X_{it} is a vector of control variables. $treatment_i$ is an indicator equal to 1 if the worker works more that 35 hours at time t and 0 otherwise.
Least squares estimates.
Sources: Enquête Emploi, 1993-2000; authors' calculations.

Table 10: Satisfaction with Hours of Work and Amount of Leisure Time

	Men		Women	
Satisfaction with:	Working Hours	Amount of Leisure Time	Working Hours	Amount of Leisure Time
Treatment	-2.375**	0.217***	-3.139***	0.023
	(1.186)	(0.058)	(0.993)	(0.182)
Post	0.908***	0.098***	0.648***	0.088**
	(0.055)	(0.014)	(0.058)	(0.042)
Treatment*Post	-1.610**	-0.114***	-1.773***	0.006
	(0.645)	(0.033)	(0.559)	(0.117)
Health problems	0.027	-0.167***	0.061	-0.136
	(0.167)	(0.061)	(0.261)	(0.253)
Married	-0.096	-0.105	-0.161	-0.201***
	(0.112)	(0.090)	(0.185)	(0.038)
Age	-0.002	-0.003	0.009***	-0.010*
	(0.004)	(0.005)	(0.004)	(0.006)
Children under 12	0.017	-0.231**	-0.153**	-0.659***
	(0.045)	(0.111)	(0.073)	(0.074)
N	1016		562	

Notes: Robust standard errors adjusted for clustering by treatment × Post. Regressions include region effects, education, net monthly income, and GDP× treatment. The sample is limited to full time employees.
Estimated equations:

$$P(y_{it} = k) = \Phi[\alpha post_t + \delta_0 treatment_i + \delta_1 (treatment \times post)_{it} + \beta X_{it}]$$

y_{it} represents self-reported levels of satisfaction with hours worked and the amount of leisure time, on a scale from 1 (not satisfied) to 6 (fully satisfied). $post_t$ is an indicator equal to 1 if year t falls within the period after the law and 0 otherwise. X_{it}

is a vector of control variables. $treatment_i$ is an indicator equal to 1 if the worker works in a large firm just before the law was implemented and 0 if he works in a small firm just before the law was implemented.
Ordered Probit estimates.
Sources: ECHP for France, Waves 1 to 8 (1994-2001); authors' calculations.

Figure 1. Log employment probabilities by firm size

A. Men

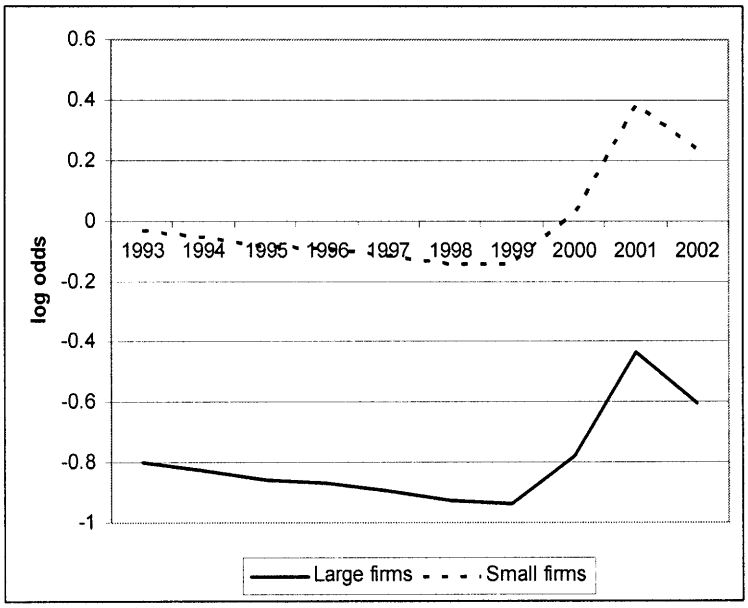

A. Women

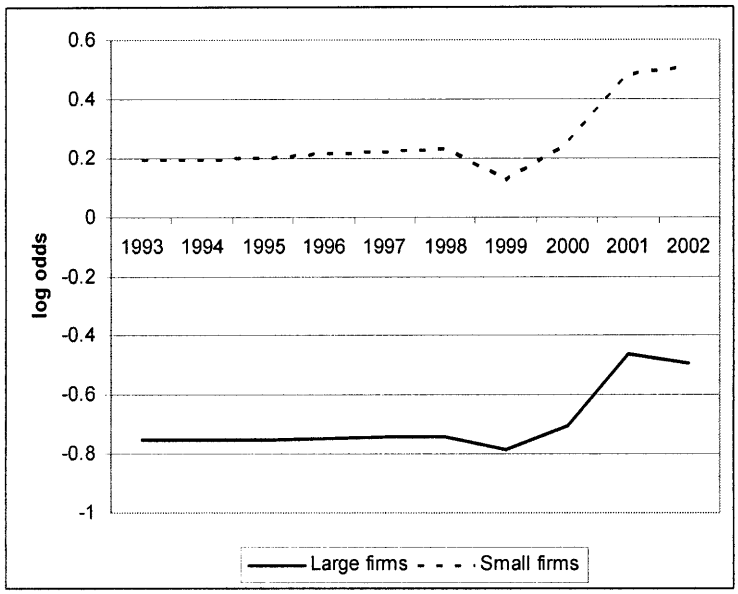

Notes: The two series are $\log[\frac{\Pr ob(employed\ large\ firm)}{\Pr ob(unemployed)}]$ and $\log[\frac{\Pr ob(employed\ small\ firm)}{\Pr ob(unemployed)}]$.

Sources: Enquête Emploi, 1993-2002; authors' calculations.

Bibliography

[1] Alesina, A., E. Glaeser, and B. Sacerdote (2005), 'Work and Leisure in the U.S. and Europe: Why So Different?', *NBER Macroeconomics Annual*, Vol. 20, Issue 1, pp. 1–64.

[2] Altonji, J. and C. Paxson (1988), 'Labour Supply Preferences, Hours Constraints, and Hours-Wage Trade-offs', *Journal of Labour Economics*, Vol. 6, No. 2, pp. 254-276.

[3] Angrist, J. and A. Krueger (1999), 'Empirical Strategies in Labour Economics', *Handbook of Labour Economics*, Vol. 3A, pp.1227-1366, Ashenfelter and Card (editors), Elsevier Science, New York.

[4] Blanchard, O. (2004), 'The Economic Future of Europe', *Journal of Economic Perspectives*, Vol. 18, No. 4, pp. 3–26.

[5] Calmfors, L, and M. Hoel (1988), 'Work Sharing and Overtime', *Scandinavian Journal of Economics,* Vol. 90 (1), pp. 45-62.

[6] Cette, G, N. Dromel, and D. Méda (2005), 'Les determinants du jugement des salaries sur la RTT', *Économie et Statistique*, Paris: INSEE, No. 37677, June, pp. 117-151.

[7] Cooper, R., and A. John (1988), 'Coordinating Coordination Failures in Keynesian Models', *The Quarterly Journal of Economics*, Vol. 103, No. 3, August, pp. 441–63.

[8] Crépon, B., and F. Kramarz (2002), 'Employed 40 Hours or Not Employed 39: Lessons from the 1982 Mandatory Reduction of the Workweek', *Journal of Political Economy*, Vol. 110, No. 6, pp. 1355–89.

[9] Crépon, B., M. Leclair, and S. Roux (2005), 'RTT, productivité et employ: nouvelles estimations sur données d'entreprises', *Économie et Statistique*, Paris: INSEE, No. 37677, June, pp. 55-89.

[10] DARES-BDF-OFCE (1998), 'L'impact macroéconomique d'une politique de RTT: l'approche par les modèles macroéconometriques', DARES document de travail no. 17.

[11] Diener, E., and E. Suh (1999), 'National Differences in Subjective Well-Being', in *Well-Being: the Foundations of Hedonic Psychology,* ed. by D. Kahneman, E. Diener and N. Schwarz, New York: Russell Sage Foundation, pp. 434–50.

[12] Gubian A., S. Jugnot, F. Lerais, and V. Passeron (2005), 'Les effets de la RTT sur l'emploi: des simulations *ex ante* aux évaluations *ex post*', *Économie et Statistique*, Paris: INSEE, No. 37677, June, pp. 25-54.

[13] Houpis, G. (1993), 'The Effects of Lower Hours of Work on Wages and Employment', *Discussion Paper*, No. 131, CEP, The London School of Economics and Political Science, London.

[14] Hunt, J. (1998), 'Hours Reductions as Work-Sharing', *Brookings Papers on Economic Activity*, No. 1, pp. 339–81.

[15] Hunt, J. (1999), 'Has Work-Sharing Worked in Germany?', *The Quarterly Journal of Economics*, Vol. 114, No. 1, pp. 117–48.

[16] Kapteyn, A., A. Kalwij, and A. Zaidi (2000), 'The Myth of Work sharing', *Department of Economics Discussion Paper Series*, University of Oxford, No. 32.

[17] Kapteyn, A., A. Kalwij, and A. Zaidi (2004), 'The Myth of Work sharing', *Labour Economics*, Vol. 11, pp. 293-313.

[18] Kugler, A., and G. Pica (2005), 'Effects of Employment Protection on Worker and Job Flows: Evidence from the 1990 Italian Reform', *Discussion Paper*, Bonn: Institute for the Study of Labour (IZA), No. 1743.

[19] Landers, R., J. Rebitzer, and L. Taylor (1996), 'Rat Race Redux: Adverse Selection in the Determination of Work Hours in Law Firms', *The American Economic Review*, Vol. 86, No. 3, pp. 329-348.

[20] Marimon, R. and F. Zilibotti (2000), 'Employment and Distributional Effects of Restricting Working Time', *European Economic Review*, Vol. 44, No. 7, pp. 1291-1326.

[21] Moulton, B. (1990), 'An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units', *The Review of Economics and Statistics*, Vol. 72 , No. 2, pp. 334-338.

[22] Oliveira, A. and V. Ulrich, (2002), 'L'Effet des 35 Heures sur la Durée du Travail des Salariés à Temps Partiel', *Document d'études*, No. 62, DARES, Ministère des Affaires Sociales, du Travail, et de la Solidarité, September.

[23] Prescott, E.C. (2004), 'Why Do Americans Work So Much More Than Europeans?', *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 28, No. 1, pp. 2–13.

[24] Skans, O. (2004), 'The Impact of Working-Time Reductions on Actual Hours and Wages: Evidence from Swedish Register-Data', *Labour Economics*, Vol. 11, pp. 647-665.

[25] Skuterud, M. (2007), 'Identifying the Potential of Work-Sharing as a Job-Creation Strategy', *Journal of Labour Economics*, Vol. 25, No. 2, pp. 265-287.

Chapter 3

International Investors, the U.S. Current Account, and the Dollar

3.1 Introduction

Two main forces underlie the large U.S. current account deficits of the past decade. The first is an increase in U.S. demand for foreign goods, partly due to relatively faster U.S. growth, and partly due to shifts in demand away from U.S. goods toward foreign goods. The second is an increase in foreign demand for U.S. assets, starting with high foreign private demand for U.S. equities in the second half of the 1990s, and later shifting to foreign private and then central bank demand for U.S. bonds in the 2000s. Both forces have contributed to steadily increasing current account deficits since the mid-1990s, accompanied by a real dollar appreciation until late 2001 and a real depreciation since. The depreciation accelerated in late 2004, raising the issues of whether and how much more is to come and, if so, against which currencies: the euro, the yen, or the Chinese renminbi.

We address these issues by developing a simple model of exchange rate and current account

[0]This chapter is joint work with Olivier Blanchard (MIT) and Francesco Giavazzi (Universitá Commerciale Luigi Bocconi). It has been published in *Brookings Papers on Economic Activity*, 1:2005. An earlier version of this paper with the same title was circulated as MIT working paper WP 05-02, January 2005. We thank Ben Bernanke, Ricardo Caballero, Menzie Chinn, William Cline, Guy Debelle, Kenneth Froot, Pierre-Olivier Gourinchas, Søren Harck, Maurice Obstfeld, Hélène Rey, Roberto Rigobon, Kenneth Rogoff, Nouriel Roubini, and the participants at the Brookings Panel conference for comments. We also thank Suman Basu, Nigel Gault, Brian Sack, Catherine Mann, Kenneth Matheny, Gian Maria Milesi-Ferretti, and Philip Lane for help with data.

determination, which we then use to interpret the recent behavior of the U.S current account and the dollar and explore what might happen in alternative future scenarios. The model's central assumption is that there is imperfect substitutability not only between U.S. and foreign goods, but also between U.S. and foreign assets. This allows us to discuss the effects not only of shifts in the relative demand for goods, but also of shifts in the relative demand for assets. We show that increases in U.S. demand for foreign goods lead to an initial real dollar depreciation, followed by further, more gradual depreciation over time. Increases in foreign demand for U.S. assets lead instead to an initial appreciation, followed by depreciation over time, to a level lower than before the shift.

The model provides a natural interpretation of the recent behavior of the U.S. current account and the dollar exchange rate. The initial net effect of the shifts in U.S. demand for foreign goods and in foreign demand for U.S. assets was a dollar appreciation. Both shifts however imply an eventual depreciation. The United States appears to have entered this depreciation phase.

How much depreciation is to come, and at what rate, depends on how far the process has come and on future shifts in the demand for goods and the demand for assets. This raises two main issues. First, can one expect the deficit to largely reverse itself without changes in the exchange rate? If it does, the needed depreciation will obviously be smaller. Second, can one expect foreign demand for U.S. assets to continue to increase? If it does, the depreciation will be delayed, although it will still have to come eventually. While there is substantial uncertainty about the answers, we conclude that neither scenario is likely. This leads us to anticipate, in the absence of surprises, more dollar depreciation to come at a slow but steady rate.

Surprises will, however, take place; only their sign is unknown. We again use the model as a guide to discuss a number of alternative scenarios, from the abandonment of the renminbi's peg against the dollar, to changes in the composition of reserves held by Asian central banks, to changes in U.S. interest rates.

This leads us to the last part of the paper, where we ask how much of the dollar's future depreciation is likely to take place against the euro, and how much against Asian currencies. We extend our model to allow for four "countries": the United States, the euro area, Japan, and China. We conclude that, again absent surprises, the path of adjustment is likely to

be associated primarily with an appreciation of the Asian currencies, but also with a further appreciation of the euro against the dollar.

3.2 A Model of the Exchange Rate and the Current Account

Much of economists' intuition about joint movements in the exchange rate and the current account is based on the assumption of perfect substitutability between domestic and foreign assets. As we shall show, introducing imperfect substitutability changes the picture substantially. Obviously, it allows one to think about the dynamic effects of shifts in asset preferences. But it also modifies the dynamic effects of shifts in preferences with respect to goods.

We are not the first to insist on the potential importance of imperfect substitutability. Indeed, the model we present builds on an older (largely and unjustly forgotten) set of papers by Paul Masson, Dale Henderson and Kenneth Rogoff, and, especially, Pentti Kouri[1]. These papers relax the interest parity condition and instead assume imperfect substitutability of domestic and foreign assets. Masson and Henderson and Rogoff focus mainly on issues of stability; Kouri focuses on the effects of changes in portfolio preferences and the implications of imperfect substitutability between assets for shocks to the current account.

The value added of this paper is in allowing for a richer description of gross asset positions. By doing this, we are able to incorporate into the analysis the "valuation effects" that have been at the center of recent empirical research on gross financial flows[2], and that play an important role in the context of U.S. current account deficits. Many of the themes we develop, including the roles of imperfect substitutability and valuation effects, have also been recently emphasized by Maurice Obstfeld[3].

[1] Masson (1981); Henderson and Rogoff (1982); Kouri (1983). The working paper version of the paper by Kouri dates from 1976. One could argue that there were two fundamental papers written that year, the first by Dornbusch (1976), who explored the implications of perfect substitutability, and the other by Kouri explored the implications of imperfect substitutability. The Dornbusch approach, with its powerful implications, has dominated research since then. But imperfect substitutability seems central to the issues we face today. Branson (1985) provides a survey of this early literature.

[2] See, in particular, Gourinchas and Rey (2005) and Lane and Milesi-Ferretti (2002, 2004).

[3] Obstfeld (2004). We limit our analysis of valuation effects to those originating from exchange rate movements. Valuation effects can and do also arise from changes in asset prices, particularly stock prices. The empirical analysis of a much richer menu of possible valuation effects has recently become possible, thanks to the data on gross financial flows and gross asset positions assembled by Lane and Milesi-Ferretti.

3.2.1 The Case of Perfect Substitutability

To see how imperfect substitutability of assets matters, it is best to start from the well-understood case of perfect substitutability. Consider a world with two "countries": the United States and a single foreign country comprising the rest of the world. We can think of the U.S. current account and the exchange rate as being determined by two relations. The first is the uncovered interest parity condition:

$$(1+r) = (1+r^*)\frac{E}{E^e_{+1}}$$

where r and r^* are U.S. and foreign real interest rates, respectively (asterisks denote foreign variables), E is the real exchange rate defined as the price of U.S. goods in terms of foreign goods (so that an increase in the exchange rate denotes an appreciation of the dollar), and E^e_{+1} is the expected real exchange rate in the next period. The condition states that expected returns on U.S. and foreign assets must be equal.

The second relation is the equation giving net debt accumulation:

$$F_{+1} = (1+r)F + D(E_{+1}, z_{+1})$$

where $D(E,z)$ is the trade deficit. The trade deficit is an increasing function of the real exchange rate (so that $D_E > 0$). All other factors—changes in total U.S. or foreign spending, as well as changes in the composition of U.S. or foreign spending between foreign and domestic goods at a given exchange rate—are captured by the shift variable z. We define z such that an increase worsens the trade balance ($D_Z > 0$). F is the net debt of the United States, denominated in terms of U.S. goods. The condition states that net debt in the next period is equal to net debt in the current period times 1 plus the interest rate, plus the trade deficit in the next period.

Assume that the trade deficit is linear in E and z, so that $D(E,z) = \theta E + z$. Assume also, for convenience, that U.S. and foreign interest rates are equal ($r^* = r$) and constant. From the interest parity condition, it follows that the expected exchange rate is constant and equal to the current exchange rate. The value of the exchange rate is obtained in turn by solving out the net debt accumulation forward and imposing the condition that net debt does not grow at

a rate above the interest rate. Doing this gives

$$E = -\frac{r}{\theta}\,[F_{-1} + \frac{1}{1+r}\sum_{0}^{\infty}(1+r)^{-i}\,z^{e}_{+i}\,]$$

The exchange rate depends negatively on the initial net debt position and on the sequence of current and expected shifts to the trade balance.

Replacing the exchange rate in the net debt accumulation equation in turn gives

$$F_{+1} - F = [z - \frac{r}{1+r}\sum_{0}^{\infty}(1+r)^{-i}\,z^{e}_{+i}\,]$$

That is, the change in the net debt position depends on the difference between the current shift and the present value of future shifts to the trade balance.

For our purposes these two equations have one main implication. Consider an unexpected, permanent increase in z at time t - say, an increase in the U.S. demand for Chinese goods (at a given exchange rate) - by Δz. Then, from the two equations above,

$$E - E_{-1} = -\frac{\Delta z}{\theta}; \qquad F_{+1} - F = 0$$

In words: permanent shifts lead to a depreciation large enough to maintain current account balance. By a similar argument, shifts that are expected to be long lasting lead to a large depreciation and only a small current account deficit. As we argue later, this is not what has happened in the United States over the last ten years. The shift in z appears to be, if not permanent, at least long lasting. Yet it has not been offset by a large depreciation but has been reflected instead in a large current account deficit. This, we shall argue, is the result of two factors, both closely linked to imperfect substitutability. The first is that, under imperfect substitutability, the initial depreciation in response to an increase in z is more limited, and, by implication, the current account deficit is larger and longer lasting. The second is that, under imperfect substitutability, asset preferences matter. An increase in foreign demand for U.S. assets, for example—an event that obviously could not be analyzed in the model with perfect substitutability we just presented—leads to an initial appreciation and a current account deficit.

And such a shift has indeed played an important role since the mid-1990s.

3.2.2 Imperfect Substitutability and Portfolio Balance

We now introduce imperfect substitutability between assets. Let W denote the wealth of U.S. investors, measured in units of U.S. goods. W is equal to the stock of U.S. assets, X, minus the net debt position of the United States, F:

$$W = X - F$$

Similarly, let W^* denote foreign wealth and X^* denote foreign assets, both in terms of foreign goods. Then the wealth of foreign investors, expressed in terms of U.S. goods, is given by

$$\frac{W^*}{E} = \frac{X^*}{E} + F$$

Let R^e be the relative expected gross real rate of return on holding U.S. assets versus foreign assets:

$$R^e \equiv \frac{1+r}{1+r^*}\frac{E^e_{+1}}{E} \tag{3.1}$$

Under perfect substitutability, the case studied above, R^e was always equal to 1; this need not be the case under imperfect substitutability[4]

U.S. investors allocate their wealth W between U.S. and foreign assets. They allocate a share α to U.S. assets and, by implication, a share $(1-\alpha)$ to foreign assets. Symmetrically, foreign investors invest a share α^* of their wealth W^* in foreign assets and a share $(1-\alpha^*)$ in U.S. assets. Assume that these shares are functions of the relative rate of return, so that

$$\alpha = \alpha(R^e, s), \;\; \alpha_{R^e} > 0, \; \alpha_s > 0 \qquad \alpha^* = \alpha^*(R^e, s), \;\; \alpha^*_{R^e} < 0 \; \alpha^*_s < 0$$

[4] One may wonder whether, even if many investors have strong asset preferences, the effects of these preferences on expected returns are not driven away by arbitrageurs, so that expected returns are equalized. The empirical work of Gourinchas and Rey (2005), which we discuss later, strongly suggests that this does not happen, and that financial assets denominated in different currencies are indeed imperfect substitutes.

A higher relative rate of return on U.S. assets leads U.S. investors to increase the share they invest in U.S. assets, and leads foreign investors to decrease the share they invest in foreign assets.

The variable s is a shift factor, standing for all the factors that shift portfolio shares for a given relative return. By convention, an increase in s leads both U.S. and foreign investors to increase the share of their portfolio in U.S. assets for a given relative rate of return.

An important parameter in the model is the degree of home bias in U.S. and foreign portfolios. We assume that there is indeed home bias, and we capture it by assuming that the sum of portfolio shares falling on own-country assets exceeds 1:

$$\alpha(R^e, s) + \alpha^*(R^e, s) > 1$$

Equilibrium in the market for U.S. assets (and, by implication, in the market for foreign assets) implies

$$X = \alpha(R^e, s)\, W + (1 - \alpha^*(R^e, s))\, \frac{W^*}{E}$$

The supply of U.S. assets must be equal to U.S. demand plus foreign demand for those assets. Given the definition of F introduced earlier, this condition can be rewritten as

$$X = \alpha(R^e, s)(X - F) \; + \; (1 - \alpha^*(R^e, s))\, (\frac{X^*}{E} + F) \tag{3.2}$$

where R^e is given in turn by equation (3.1) and depends in particular on E and E^e_{+1}. This gives us the first relation, which we refer to as the *portfolio balance* relation, between net debt, F, and the exchange rate, E.

To see its implications most clearly, consider the limiting case where the degree of substitutability is zero, so that the shares α and α^* do not depend on the relative rate of return. In this case

- The portfolio balance condition fully determines the exchange rate as a function of the world distribution of wealth, $(X - F)$ and $[(X^*/E) + F)]$. In sharp contrast to the case of perfect substitutability, news about current or future current account balances, such as a

permanent shift in z, has no effect on the current exchange rate.

- Over time, current account deficits lead to changes in F, and thus to changes in the exchange rate. The slope of the relation between the exchange rate and net debt is given by

$$\frac{dE/E}{dF} = -\frac{\alpha + \alpha^* - 1}{(1-\alpha^*)X^*/E} < 0$$

So, in the presence of home bias, an increase in net debt is associated with a lower exchange rate. The reason is that, as wealth is transferred from the United States to the rest of the world, home bias leads to a decrease in the demand for U.S. assets, which in turn requires a decrease in the exchange rate.

Outside this limiting case, the portfolio balance determines a relation between net debt and the exchange rate for a given expected rate of depreciation. The exchange rate is no longer determined myopically. But the two insights from the limiting case remain: On the one hand, the exchange rate will respond less to news about the current account than it does under perfect substitutability. On the other, it will respond to changes either in the world distribution of wealth or in portfolio preferences.

3.2.3 Imperfect Substitutability and Current Account Balance

Assume, as before, that U.S. and foreign goods are imperfect substitutes and that the U.S. trade deficit, in terms of U.S. goods, is given by

$$D = D(E, z),\ D_E > 0,\ D_z > 0$$

Turn now to the equation giving the dynamics of the U.S. net debt position. Given our assumptions, U.S. net debt is given by

$$F_{+1} = (1 - \alpha^*(R^e, s))\,\frac{W^*}{E}\,(1+r)\ -\ (1 - \alpha(R^e, s))\,W\,(1+r^*)\frac{E}{E_{+1}}\ +\ D(E_{+1}, z_{+1})$$

In words: Net debt in the next period is equal to the value of U.S. assets held by foreign

investors next period, minus the value of foreign assets held by U.S. investors next period, plus the trade deficit next period:

- The value of U.S. assets held by foreign investors next period is equal to their wealth in terms of U.S. goods this period times the share they invest in U.S. assets this period times the gross rate of return on U.S. assets in terms of U.S. goods.
- The value of foreign assets held by U.S. investors next period is equal to U.S. wealth this period times the share they invest in foreign assets this period times the realized gross rate of return on foreign assets in terms of U.S. goods.

The previous equation can be rewritten as

$$F_{+1} = (1+r)F + (1-\alpha(R^e,s))(1+r)(1-\frac{1+r^*}{1+r}\frac{E}{E_{+1}})(X-F) + D(E_{+1},z_{+1}) \qquad (3.3)$$

We shall call this the *current account balance* relation[5].

The first and last terms on the right-hand side are standard: next-period net debt is equal to this-period net debt times the gross rate of return, plus the trade deficit next period. The term in the middle reflects valuation effects, recently stressed by Pierre-Olivier Gourinchas and Hélène Rey and by Philip Lane and Gian Maria Milesi-Ferretti[6]. Consider, for example, an unexpected decrease in the price of U.S. goods - an unexpected decrease in E_{+1} relative to E. This dollar depreciation increases the dollar value of U.S. holdings of foreign assets, decreasing the U.S. net debt position.

Putting things together, a depreciation improves the U.S. net debt position in two ways:

[5] This appears to give a special role to α rather than α^*, but in fact this is not the case. A symmetrical expression can be derived with α^* appearing instead of α. Put another way, F, α and α^* are not independent. F_{+1} can be expressed in terms of any two of the three.

[6] Gourinchas and Rey (2005); Lane and Milesi-Ferretti (2004). As a matter of logic, one can have both perfect substitutability and valuation effects. (Following standard practice, we ignored valuation effects in the perfect substitutability model presented earlier by implicitly assuming that, if net debt was positive, U.S. investors did not hold foreign assets and net debt was therefore equal to the foreign holdings of dollar assets.) Under perfect substitutability, however, there is no guide as to what determines the shares, and therefore what determines the gross positions of U.S. and foreign investors.

the first, conventional way through the improvement in the trade balance, and a second way through asset revaluation. Note that

- The strength of the valuation effects depends on gross rather than net positions and so on the share of the U.S. portfolio in foreign assets $(1-\alpha)$ and on U.S. wealth $(X-F)$. It is present even if $F=0$.
- The strength of the valuation effects depends on our assumption that U.S. gross liabilities are denoted in dollars, so that their value in dollars is unaffected by a dollar depreciation. Valuation effects would obviously be very different when, as is typically the case for emerging market economies, gross positions are smaller and liabilities are denominated in foreign currency.

3.2.4 Steady State and Dynamics

Assume the stocks of assets X and X^* and the shift variables z and s to be constant. Assume also r and r^* to be constant and equal to each other. In this case the steady-state values of net debt F and E are characterized by two relations.

The first is the portfolio balance relation (equation (3.2)). Given the equality of interest rates and the constant exchange rate, $R^e=1$, the relation takes the form

$$X=\alpha(1,s)(X-F)\ +\ (1-\alpha^*(1,s))\ (\frac{X^*}{E}+F)$$

This first relation implies a negative relation between net debt and the exchange rate: As we showed earlier, in the presence of home bias, a larger U.S. net debt, which transfers wealth to foreign investors, shifts demand away from U.S. assets and thus lowers the exchange rate.

The second relation is the current account balance relation (equation (3.3)). Given the equality of interest rates, and given the constant exchange rate and net debt, the relation takes the form

$$0=rF+D(E,z)$$

This second relation also implies a negative relation between net debt and the exchange rate. The larger the net debt, the larger the trade surplus required in steady state to finance interest payments on the debt, and thus the lower the exchange rate[7]. This raises the question of the stability of the system. The system is (locally saddle point) stable if, as drawn in figure 1, the portfolio balance locus is steeper than the current account balance locus. (Appendix A characterizes the dynamics.) To understand this condition, consider an increase in U.S. net debt. This increase has two effects on the current account deficit, and thus on the change in net debt: it increases interest payments, but it also leads, through the portfolio balance relation, to a lower exchange rate and thus a decrease in the trade deficit. For stability to prevail, the net effect must be that the increase in net debt reduces the current account deficit. This condition appears to be satisfied for plausible parameter values (the next section explores this issue further), and we assume that it is satisfied here. In this case the path of adjustment—the saddle path—is downward sloping, as drawn in figure 1.

3.2.5 The Effects of a Shift toward Foreign Goods

We can now characterize the effects of shifts in preferences for goods or assets. Figure 2 shows the effect of an unexpected and permanent increase in z. One can think of this increase as coming either from an increase in U.S. activity relative to foreign activity, or from a shift in exports or imports at a given level of activity and a given exchange rate; we defer to later a discussion of the sources of the actual shift in z over the past decade in the United States.

For any given level of net debt, current account balance requires a lower exchange rate: the current account balance locus shifts down. The new steady state is at point C, associated with a lower exchange rate and a larger net debt.

Valuation effects imply that any unexpected depreciation leads to an unexpected decrease in the net debt position. If we denote by ΔE the unexpected change in the exchange rate at the time of the shift, it follows from equation (3.3) that the relation between the two at the time of the shift is given by

[7] If we had allowed r and r^* to differ, the relation would have an additional term and take the form $0 = rF + (1-\alpha)(r-r^*)(X-F) + D(E,z)$. This additional term implies that if, for example, a country pays a lower rate of return on its liabilities than it receives on its assets, it may be able to combine positive net debt with positive net income payments from abroad—the situation in which the United States remains today.

$$\Delta F = (1-\alpha)(1+r^*)(X-F)\frac{\Delta E}{E} \tag{3.4}$$

The economy jumps initially from point A to point B and then converges over time along the saddle path, from point B to point C. The shift in the trade deficit leads to an initial, unexpected depreciation, followed by further depreciation and net debt accumulation over time until the new steady state is reached.

Note that the degree of substitutability between assets does not affect the steady state; more formally, the steady state depends on $\alpha(1,s)$ and $\alpha^*(1,s)$, and so changes in α_R and α_R^* that leave $\alpha(1,s)$ and $\alpha^*(1,s)$ unchanged do not affect the steady state. In other words, the eventual depreciation is the same no matter how close substitutes U.S. and foreign assets are. But the degree of substitutability plays a central role in the dynamics of adjustment and in the relative roles of the initial unexpected depreciation and the anticipated depreciation thereafter. This is shown in figure 3, which shows the effects of three different values of α_R and α_R^* on the path of adjustment. (The three simulations are based on values for the parameters introduced in the next section. The purpose here is simply to show the qualitative properties of the paths. We return to the quantitative implications later.)

The less substitutable U.S. and foreign assets are—that is, the smaller are α_R and α_R^*—the smaller the initial depreciation and the higher the anticipated rate of depreciation thereafter. To understand why, consider the extreme case where the shares do not depend on rates of return: U.S. and foreign investors want to maintain constant shares, no matter what the relative rate of return is. In this case the portfolio balance relation (equation (3.2)) implies that there will be no response of the exchange rate to the unexpected change in z at the time it happens: any movement in the exchange rate would be inconsistent with equilibrium in the market for U.S. assets. Only over time, as the deficit leads to an increase in net debt, will the exchange rate decline.

Conversely, the more substitutable U.S. and foreign assets are, the larger will be the initial depreciation, the lower the anticipated rate of depreciation thereafter, and the longer the time taken to reach the new steady state. The limit of perfect substitutability—corresponding to the model discussed at the start—is actually degenerate: the initial depreciation is such as

to maintain current account balance, and the economy does not move from there on, never reaching the new steady state (and so the anticipated rate of depreciation is equal to zero).

To summarize: In contrast to the case of perfect substitutability between assets we saw earlier, an increase in U.S. demand for foreign goods leads to a limited depreciation initially, a potentially large and long lasting current account deficit, and a steady depreciation over time.

3.2.6 The Effects of a Shift toward U.S. Assets

Figure 4 shows the effect of an unexpected and permanent increase in s, that is, an increase in the demand for U.S. assets. Again we defer to later a discussion of the potential factors behind such an increase.

By assumption, the increase in s leads to an increase in $\alpha(1, s)$ and a decrease in $\alpha^*(1, s)$. At a given level of net debt, portfolio balance requires an increase in the exchange rate. The portfolio balance locus shifts up. The new steady state is at point C, associated with a lower exchange rate and larger net debt.

The dynamics are given by the path ABC. The initial adjustment of E and F must again satisfy the condition in equation (3.4). So the economy jumps from point A to point B and then converges over time from point B to point C. The dollar initially appreciates, triggering an increase in the trade deficit and a deterioration of the net debt position. Over time, net debt continues to increase and the dollar depreciates. In the new equilibrium the exchange rate is necessarily lower than before the shift: this reflects the need for a larger trade surplus to offset the interest payments on the now-larger U.S. net debt. In the long run the favorable portfolio shift leads to a depreciation.

Again the degree of substitutability between assets plays an important role in the adjustment. This is shown in figure 5, which plots the path of adjustment for three different values of α_R and α_R^*. The less substitutable are U.S. and foreign assets, the greater the initial appreciation and the higher the anticipated rate of depreciation thereafter. Although the depreciation is eventually the same in all cases (the steady state is invariant to the values of α_R and α_R^*), the effect of portfolio shifts is more muted but longer lasting when the degree of substitutability is high.

3.2.7 An Interpretation of the Past

Looking at the effects of shifts in preferences for goods and for assets under imperfect asset substitutability suggests three main conclusions:

- Shifts in preferences toward foreign goods lead to an initial depreciation, followed by a further anticipated depreciation. Shifts in preferences towards U.S. assets lead to an initial appreciation, followed by an anticipated depreciation.

- The empirical evidence suggests that both types of shifts have been at work in the United States in the recent past. The first shift, by itself, would have implied a steady depreciation in line with increased trade deficits, whereas instead an initial appreciation was observed. The second shift can explain why the initial appreciation has been followed by a depreciation. But it attributes the increase in the trade deficit fully to the initial appreciation, whereas the evidence is of a large adverse shift in the trade balance even after controlling for the effects of the exchange rate. (This does not do justice to an alternative, and more conventional, monetary policy explanation, in which high U.S. interest rates relative to foreign interest rates at the end of the 1990s led to an appreciation, followed since by a depreciation. The observed relative interest rate differentials seem too small, however, to explain the movement in exchange rates.)

- Both shifts lead eventually to a steady depreciation, to a lower exchange rate than before the shift. This follows from the simple condition that a larger net debt, no matter what its origin, requires larger interest payments in steady state and thus a larger trade surplus. The lower the degree of substitutability between U.S. and foreign assets, the higher the expected rate of depreciation along the path of adjustment. The United States appears to have indeed entered this depreciation phase.

3.3 How Large a Depreciation? A Look at the Numbers

The model is simple enough that one can insert some values for the parameters and draw the implications for the future. More generally, the model provides a way of looking at the data, and this is what we do in this section.

3.3.1 Parameter Values

Consider first what we know about portfolio shares: In 2003 U.S. financial wealth, W, was \$34.1 trillion, or about three times U.S. GDP of \$11 trillion[8]. Non-U.S. world financial wealth is harder to assess. For the euro area financial wealth was about €16 trillion in 2003, compared with GDP of €7.5 trillion; Japanese financial wealth was about ¥900 trillion in 2004, compared with GDP of ¥500 trillion[9]. If one extrapolates from a ratio of financial assets to GDP of about 2 for both Japan and Europe, and GDP for the non-U.S. world of approximately \$18 trillion in 2003, a reasonable estimate for W^*/E is \$36 trillion—roughly the same as for the United States.

The net U.S. debt position, F, measured at market value, was \$2.7 trillion in 2003, up from approximate balance in the early 1990s[10]. By implication, U.S. assets, X, were $W + F =$ \$36.8 trillion (\$34.1 trillion + \$2.7 trillion), and foreign assets, X^*/E, were $W^*/E - F =$ \$33.3 trillion (\$36.0 trillion - \$2.7 trillion). Put another way, the ratio of U.S. net debt to U.S. assets, F/X, was 7.3 percent (\$2.7 trillion/\$36.8 trillion); the ratio of U.S. net debt to U.S. GDP was 24.5 percent (\$2.7 trillion/\$11.0 trillion).

In 2003 gross U.S. holdings of foreign assets, at market value, were \$7.9 trillion. Together with the value for W, this implies that the share of U.S. wealth in U.S. assets, α, was 1 – (7.9/34.1), or 0.77. Gross foreign holdings of U.S. assets, at market value, were \$10.6 trillion. Together with the value of W^*/E, this implies that the share of foreign wealth in foreign assets, α^*, was equal to 1 – (10.6/36.0), or 0.71.

To get a sense of the implications of these values for α and α^*, note from equation (3.2) that a transfer of one dollar from U.S. wealth to foreign wealth implies a decrease in the demand for U.S. assets of $(\alpha + \alpha^* - 1)$ dollars, or 48 cents[11].

To summarize:

$W =$ \$34.1 trillion

[8] Financial wealth data are from the Flow of Funds Accounts of the United States 1995-2003, table L100, Board of Governors of the Federal Reserve System, December 2004.

[9] The figure for Europe is from ECB Bulletin, February 2005, table 3.1, and that for Japan from Bank of Japan, Flow of Funds (www.boj.or.jp/en/stat/sj/sj.html).

[10] The source for the numbers in this and the next paragraph is Bureau of Economic Analysis, International Transactions, Table 2, International Investment Position of the United States at Year End, 1976-2003, June 2004.

[11] Note that this conclusion depends on the assumption we make in our model that marginal and average shares are equal. This may not be the case.

W^*/E = \$36.0 trillion

X = \$36.8 trillion

X^*/E = \$33.3 trillion

F = \$2.7 trillion

$\alpha = 0.77$

$\alpha^* = 0.71$.

We would like to know not only the values of the shares, but also their dependence on the relative rate of return—the values of the derivatives α_R and α_R^*. Little is known about these values. Gourinchas and Rey provide indirect evidence of the relevance of imperfect substitutability by showing that a combination of the trade deficit and the net debt position helps predict a depreciation (we return to their results later)[12]; this would not be the case under perfect substitutability. However, it is difficult to back out estimates of α_R and α_R^* from their results. Thus, when needed below, we derive results under alternative assumptions about these derivatives.

The next important parameter in our model is θ, the effect of the exchange rate on the trade balance. The natural starting point here is the Marshall-Lerner relation:

$$\frac{dD}{Exports} = [\eta_{imp} - \eta_{\exp} - 1] \frac{dE}{E}$$

where η_{imp} and η_{exp} are, respectively, the elasticities of imports and exports with respect to the real exchange rate.

Estimates of the η_s based on estimated U.S. import and export equations range quite widely[13]. In some cases the estimates imply that the Marshall-Lerner condition (the condition that the term in brackets be positive, so that a depreciation improves the trade balance) is barely satisfied. Estimates used in macroeconometric models imply a value for the term in brackets between 0.5 and 0.9. Put another way, together with the assumption that the ratio of U.S. exports to U.S. GDP is 10 percent, they imply that a reduction of the ratio of the trade deficit to GDP by 1 percentage point requires a depreciation somewhere between 11 and 20 percent.

[12] Gourinchas and Rey (2005).

[13] See the survey by Chinn (2004).

One may believe, however, that measurement error, complex lag structures, and misspecification all bias these estimates downward. An alternative approach is to derive the elasticities from plausible specifications of utility and the pass-through behavior of firms. Using such an approach in a model with nontradable goods, domestic tradable goods, and foreign tradable goods, Obstfeld and Rogoff find that a 1-percentage-point decrease in the ratio of the trade deficit to GDP requires a decrease in the real exchange rate of somewhere between 7 and 10 percent—a smaller depreciation than implied by the macroeconometric models[14].

Which value to use is obviously crucial in assessing the scope of the required exchange rate adjustment. We choose an estimate for the term in brackets in the Marshall-Lerner equation of 0.7—toward the high range of empirical estimates but lower than the Obstfeld-Rogoff elasticities. This estimate, together with an exports-to-GDP ratio of 10 percent, implies that a reduction of the ratio of the trade deficit to GDP of 1 percentage point requires a depreciation of 15 percent.

3.3.2 A Simple Exercise

We have argued that a depreciation of the dollar has two effects: a conventional one through the trade balance, and another through valuation effects. To get a sense of their relative magnitudes, consider the effects of an unexpected depreciation in our model. More specifically, consider the effects of an unexpected 15 percent decrease in E_{+1} relative to E on net debt, F_{+1}, in equation (3.3).

The first effect of the depreciation is to improve the trade balance. Given our earlier discussion and assumptions, such a depreciation reduces the trade deficit by 1 percent of GDP (which is why we chose to look at a depreciation of 15 percent).

The second effect is to increase the dollar value of U.S. holdings of foreign assets (and to reduce the foreign currency value of foreign holdings of U.S. assets) and thus reduce the U.S. net debt position. From equation (3.3) (with both sides divided by U.S. output, Y, to make the interpretation of the magnitudes easier), this effect is given by

$$\frac{dF_{+1}}{Y} = -(1-\alpha)(1+r^*)\frac{X-F}{Y}\frac{dE}{E}$$

[14] Obstfeld and Rogoff (2004).

From the earlier discussion, $(1-\alpha)$ is equal to 0.23, and $(X-F)/Y$ to 3. Assume that r^* is equal to 4 percent. The effect of a 15 percent depreciation is then to reduce the ratio of net debt to GDP by 10 percentage points ($0.23 \times 1.04 \times 3 \times 0.15$). This implies that, after the unexpected depreciation, interest payments are lower by 4 percent times 10 percent, or 0.4 percent of GDP[15]. Putting things together, a 15 percent depreciation improves the current account balance by 1.4 percent of GDP, with roughly one-third of the improvement due to valuation effects[16].

It is tempting at this point to ask how large an unexpected depreciation would lead to a sustainable U.S. current account deficit *today*?[17] Take the actual current account deficit of about 6 percent. What the "sustainable" current account deficit is depends on the ratio of net debt to GDP that the United States is willing to sustain, and on the growth rate of GDP: if g is the growth rate of U.S. GDP, the United States can sustain a current account deficit of gF/Y. Assuming, for example, a GDP growth rate of 3 percent and a ratio of net debt to GDP of 25 percent (the ratio prevailing today, but one that has no particular claim to being the right one for this computation) implies that the United States can run a current account deficit of 0.75 percent while maintaining a constant ratio of net debt to GDP. In this case the depreciation required to shift from the actual to the sustainable current account deficit would be roughly 56 percent, or (6 percent - 0.75 percent)$\times$(15 percent/1.4 percent).

This is a large number, and despite the uncertainty attached to the underlying values of many of the parameters, it is a useful number to keep in mind. But one should be clear about the limitations of the computation:

- The United States surely does not need to shift to sustainable current account balance right away. The rest of the world is still willing to lend to it, if perhaps not at the current rate. The longer the United States waits, however, the higher the ratio of net debt to GDP becomes, and thus the larger the eventual required depreciation. In this sense our

[15] This computation assumes that all foreign assets held by U.S. investors are denominated in foreign currency. In reality, some foreign bonds held by U.S. investors are denominated in dollars. This reduces the valuation effects.

[16] Lane and Milesi-Ferretti (2004) give a similar computation for a number of countries, although not for the United States.

[17] This is also the question taken up by Obstfeld and Rogoff in this volume. Their focus, relative to ours, is on the required adjustments in both the terms of trade and the real exchange rate, starting from a micro-founded model with nontraded goods, exportables, and importables.

computation gives a lower bound on the eventual depreciation.

- The computation is based on the assumption that, at the current exchange rate, the trade deficit will remain as large as it is today. If, for example, we believed that part of the current trade deficit reflects the combined effect of recent depreciations and J-curve effects, the computation above would clearly overestimate the required depreciation.

The rest of this section deals with these issues. First, by returning to dynamics, we try to get a sense of the eventual depreciation and of the rate at which it may be achieved. Second, we look at the evidence on the origins of the shifts in z and s.

3.3.3 Returning to Dynamics

How large is the effect of a given shift in z (or in s) on the accumulation of net debt and on the eventual exchange rate? And how long does it take to get there? The natural way to answer these questions is to simulate our model using the values of the parameters we derived earlier. This is indeed what the simulations presented in figures 3 and 5 did; we look now more closely at their quantitative implications.

Both sets of simulations are based on the values of the parameters given above. Recognizing the presence of output growth (which we did not allow for in the model), and rewriting the equation for net debt as an equation for the ratio of net debt to output, we take the term in front of F in the current account balance relation (equation (3.3)) to stand for the interest rate minus the growth rate. We choose an interest rate of 4 percent and a growth rate of 3 percent, so that their difference is 1 percent. We write the portfolio shares as

$$\alpha(R^e, s) = a + bR^e + s, \qquad \alpha^*(R^e, s) = a^* - bR^e - s$$

The simulations show the results for three values (10, 1.0, and 0.1) of the parameter b. A value of 1 implies that an increase in the expected relative return on U.S. assets of 100 basis points increases the desired share in U.S. assets by 1 percentage point.

Figure 3 showed the effect of an increase in z of 1 percent of U.S. GDP. Figure 5 showed the effect of an increase in s of 5 percentage points, leading to an increase in α and a decrease

in α^* of 5 percentage points at a given relative rate of return. Time is measured in years.

Figure 3 leads to two main conclusions. First, the effect of a permanent increase in z by 1 percent is to eventually increase the ratio of net debt to GDP by 17 percent and require an eventual depreciation of 12.5 percent. (Recall that the long-run effects are independent of the degree of substitutability between assets—that is, independent of the value of b.) Second, it takes a long time to get there: the figure is truncated at fifty years, by which time the adjustment is still not complete.

Figure 5 leads to similar conclusions. The initial effect of the increase in s is an appreciation of the dollar: by 23 percent if $b = 0.1$, and by 12 percent if $b = 10$. The long-run effect of the increase in s is an increase in the ratio of U.S. net debt to GDP of 35 percent and a depreciation of 15 percent. But even after fifty years the adjustment is far from complete, and the exchange rate is still above its initial level.

What should one conclude from these exercises? We conclude that, under the following assumptions—that there are no anticipated changes in z or in α or α^*, that investors have been and will be rational (the simulations are carried out under rational expectations), and that there are no surprises—the dollar will depreciate by a large amount, but at a steady and slow rate. There are good reasons to question each of these assumptions, and this we do next.

3.3.4 A Closer Look at the Trade Deficit

To think about the likely path of z, and thus of the path of the trade deficit at a given exchange rate, it is useful to write the trade deficit as the difference between the value of imports in terms of domestic goods and exports:

$$D(E, z) \equiv E\ imp(E, Z, \tilde{z}) - \exp(E, Z^*, \tilde{z}^*)$$

We have decomposed z into two components: total U.S. spending, Z, and $\tilde{z}$, which represents shifts in the relative U.S. demand for U.S. versus foreign goods, at a given level of spending and a given exchange rate. Similarly, z^* is decomposed into Z^* and $\tilde{z}^*$, the latter measuring shifts in the relative foreign demand for U.S. versus foreign goods.

Most of the large current account fluctuations in developed countries of the last few decades

have come from relative fluctuations in activity, that is, in Z relative to Z^*[18]. It has indeed been argued that the deterioration of the U.S. trade balance has come mostly from faster growth in the United States than in its trade partners, leading imports by the United States to increase faster than U.S. exports to the rest of the world. This appears, however, to have played a limited role. Europe and Japan indeed have had slower growth than the United States (U.S. output grew a cumulative 45 percent from 1990 to 2004, compared with 29 percent for the euro area and 25 percent for Japan), but these countries account for only 35 percent of U.S. exports, and meanwhile other U.S. trade partners have grown as fast as or faster than the United States. Indeed, a study by the International Monetary Fund finds nearly identical output growth rates for the United States and its export-weighted partners since the early 1990s[19].

Some have argued that the deterioration in the trade balance reflects instead a combination of rapid growth both in the United States and abroad and a U.S. import elasticity with respect to domestic spending that is higher (1.5 or above) than the elasticity of U.S. exports with respect to foreign spending. In this view rapid U.S. growth has led to a more than proportional increase in imports and an increasing trade deficit. The debate about the correct value of the U.S. import elasticity is an old one, dating back to the estimates by Hendrik Houthakker and Stephen Magee; we tend to side with the recent conclusion by Jaime Marquez that the elasticity is close to 1[20]. For our purposes, however, this discussion is not relevant. Whether the growth in the U.S. trade deficit is the result of a high import elasticity or of shifts in the $\tilde{z}$s, there are no obvious reasons to expect either the shift to reverse or growth in the United States to drastically decrease in the future.

One way of assessing the relative roles of shifts in spending, the exchange rate, and other factors is to look at the performance of import and export equations in detailed macroeconometric models. The numbers obtained using the macroeconometric model of Global Insight (formerly the Data Resources, Inc., or DRI, model) are as follows[21]: The U.S. trade deficit in goods increased from $221 billion in the first quarter of 1998 to $674 billion in the third quarter

[18] For a review of current account deficits and adjustments for twenty-one countries over the last thirty years, and references to the literature, see Debelle and Galati (2005).

[19] International Monetary Fund, Article IV United States Consultation—Staff Report, 2004. As the case of the United States indeed reminds us, output is not the same as domestic spending, but the differences in growth rates between the two over a decade are small.

[20] Houthakker and Magee (1969); Marquez (2000).

[21] We thank Nigel Gault of Global Insight for communicating these results to us.

of 2004. Of this $453 billion increase, $126 billion was due to the increase in the value of oil imports, leaving $327 billion to be explained. When the export and import equations of the model are used, activity variables and exchange rates explain $202 billion, or about 60 percent of the increase. Unexplained time trends and residuals account for the remaining 40 percent, a substantial amount[22].

Looking to the future, whether growth rate differentials, Houthakker-Magee effects, or unexplained shifts are behind the increase in the trade deficit is probably not essential. The slower growth in Europe and Japan reflects in large part structural factors, and neither Europe nor Japan is likely to make up much of the cumulative growth difference since 1995 over the next few years. One can still ask how much a given increase in growth in Europe would reduce the U.S. trade deficit. A simple computation is as follows. Suppose that Europe and Japan made up the roughly 20-percentage-point growth gap they have accumulated since 1990 vis-à-vis the United States—an unlikely scenario in the near future—so that U.S. exports to Europe and Japan increased by 20 percent. Given that U.S. exports to these countries are currently about $350 billion, the improvement would be 0.7 percent of U.S. GDP—not negligible, but not a major increase either.

One other factor, however, may hold more hope for a reduction in the trade deficit, namely, the working out of the J-curve. Nominal depreciations increase import prices, but these decrease the volume of imports only with a lag. Thus, for a while, a depreciation can increase the value of imports and worsen the trade balance, before improving it later.

One reason to think this may be important is the "dance of the dollar" and the movements of the dollar and the current account during the 1980s. From the first quarter of 1979 to the first quarter of 1985, the real exchange rate of the United States (measured by the trade-weighted major currencies index constructed by the Federal Reserve Board) increased by 41 percent. This appreciation was then followed by a sharp depreciation, with the dollar falling by 44 percent from the first quarter of 1985 to the first quarter of 1988. The appreciation was accompanied by a steady deterioration in the current account deficit, from rough balance in the early 1980s

[22]The model has a set of export and import equations disaggregated by product type. Most of the elasticities of the different components with respect to domestic or foreign spending are close to 1, indicating that Houthakker-Magee effects play a limited role (except for imports and exports of consumption goods, where the elasticity of imports with respect to consumption is 1.5 for the United States, but the elasticity of U.S. exports with respect to foreign GDP is an even higher 2.0).

to a deficit of about 2.5 percent of GDP when the dollar reached its peak in early 1985. The current account continued to worsen, however, for more than two years, reaching a peak of 3.4 percent of GDP in 1987. The divergent paths of the exchange rate and the current account from 1985 to 1987 led a number of economists to explore the idea of hysteresis in trade[23]: the notion that, once appreciation has led to a loss of market share, an equal depreciation may not be sufficient to reestablish trade balance. Just as the idea was taking hold, however, the current account position rapidly improved, and trade was roughly in balance by the end of the decade[24].

The parallels with more recent developments are clear from figure 6, which plots the dollar exchange rate and the U.S. current account during both episodes, aligned in the figure so that the dollar peak of 1985:1 coincides with the dollar peak of 2001:2. The figure suggests two conclusions:

- If the earlier episode is a reliable guide, and the lags today are similar to those that prevailed in the 1980s, the current account deficit may start to turn around soon. Today's deficit, however, is much larger than the earlier deficit was at its peak in 1987 (6 percent of GDP versus 3.5 percent), and the depreciation so far has been more limited (23 percent from 2001:2 to 2004:4, compared with 33 percent over the equivalent period from 1985:1 to 1988:3)[25].

- Hence one can surely not conclude that the depreciation so far is enough to restore the current account deficit to sustainable levels. But it may be that, in our computation, the appropriate place to start is from a J-curve-adjusted ratio of the current account deficit to GDP of 4 or 5 percent instead of 6 percent[26]. If we choose 4 percent—a very

[23] In particular, Baldwin and Krugman (1987).

[24] These issues were discussed at length in the Brookings Papers at the time. Besides Baldwin and Krugman (1987), see, for example, Cooper (1986), Dornbusch (1987), and Sachs and Lawrence (1988), with post mortems by Lawrence (1990) and Krugman (1991). Another much-discussed issue, to which we return later, was the relative roles of fiscal deficit reduction and exchange rate adjustment in closing the deficit.

[25] On the other hand, the gross positions, and thus the scope for valuation effects from dollar depreciation, are much larger now than they were then. In 1985 gross U.S. holdings of foreign assets were $1.5 trillion, compared with $8 trillion today.

[26] Forecasts by "Macroeconomic Advisors" are for an improvement in the trade balance of $75 billion, or less than 1 percent of GDP, over the next two years. (The forecast is based on a depreciation of the dollar of 4 percent over that period.) The residuals of the import price equations of the model, however, suggest an unusually low pass-through of the dollar decline to import prices over the recent past, and the forecast assumes that the low

optimistic assumption—the remaining required depreciation is 34 percent, or (4 percent - 0.75 percent) × (15 percent/1.4 percent)[27].

3.3.5 A Closer Look at Portfolio Shares

One striking aspect of the simulations presented above is how slow the depreciation is along the adjustment path. This is in contrast with some predictions of much more abrupt falls in the dollar in the near future[28]. This raises two issues: Can the anticipated depreciation be greater than in these simulations? And are there possible surprises under which the depreciation might be much faster (or slower), and, if so, what are they? We take both questions in turn.

To answer the first question, we go back to the model. We noted earlier that the lower the degree of substitutability between assets , the higher the anticipated rate of depreciation. So, by assuming zero substitutability—that is, constant asset shares except for changes coming from shifts in s—we can derive an upper bound on the anticipated rate of depreciation. Differentiating equation 2 gives

$$\frac{dE}{E} = -\frac{(\alpha + \alpha^* - 1)X}{(1-\alpha^*)X^*/E}\ d\left(\frac{F}{X}\right) + \frac{(X-F)\ d\alpha + (X^*/E + F)\ d\alpha^*}{(1-\alpha^*)X^*/E}$$

In the absence of anticipated shifts in shares (so that the second term equals zero), the anticipated rate of depreciation depends on the change in the ratio of U.S. net debt to U.S. assets: the faster the increase in net debt, the faster the decrease in the relative demand for U.S. assets, and therefore the higher the rate of depreciation needed to maintain portfolio balance. Using the parameters we constructed earlier, this equation implies

$$\frac{dE}{E} = -1.8\ d\left(\frac{F}{X}\right) + (3.5d\alpha - 3.7d\alpha^*)$$

Suppose shares remain constant. If we take the annual increase in the ratio of net debt to U.S. GDP to be 5 percent and the ratio of U.S. GDP to U.S. assets to be one-third, this gives

pass-through continues. If the pass-through were to return to its historical average, the improvement in the trade balance would be larger.

[27] This number is surprisingly close to the 33 percent obtained by Obstfeld and Rogoff in this volume.

[28] For example, by Roubini and Setser (2005).

an anticipated annual rate of depreciation of 3 percent a year $(1.8 \times 0.05/3)$[29].

If, however, shares of U.S. assets in the portfolios of either domestic or foreign investors are expected to decline, the anticipated depreciation can clearly be much larger. If, for example, we anticipate that the share of U.S. assets in foreign portfolios will decline by 2 percent over the coming year, the anticipated depreciation is 8.7 percent (2.7 percent as calculated above, plus 3.0 times 2 percent). This is obviously an upper bound on the size of the depreciation, as it assumes that the remaining investors—those who do not anticipate selling—are willing to keep a constant share of their wealth in U.S. assetsbecause of the assumption of zero substitutability) in the face of high negative expected rates of return over the coming year . Still, it implies that, under imperfect substitutability, and under the assumption that desired shares in U.S. assets will decrease, it is logically acceptable to predict a substantial depreciation of the dollar in the near future.

Are there good reasons to expect these desired shares to decrease in the near future? This is the subject of a contentious debate. Some argue that the United States can continue to finance its current account deficits at today's level for a long time to come at the same exchange rate. They argue that the poor development of financial markets in Asia and elsewhere, together with the need for Asian countries to accumulate international collateral, implies a steadily increasing relative demand for U.S. assets. They point to the latent demand for U.S. assets on the part of Chinese private investors, currently limited by capital controls. In short, they argue that foreign investors will be willing to further increase their holdings of U.S. assets for many years to come[30].

Following this argument, we can ask what increase in shares—say, what increase in $(1-\alpha^*)$, the share of U.S. assets in foreign portfolios—would be needed to absorb the current increase in net debt at a given exchange rate. From the relation derived above, setting dE/E and $d\alpha$ equal to zero gives

[29] Although comparison is difficult, this rate appears lower than that implied by the estimates of Gourinchas and Rey (2005). Their results imply that a combination of net debt and trade deficits 2 standard deviations from the mean—a situation that would appear to characterize well the United States today—implies an anticipated annual rate of depreciation of about 5 percent over the following two years.

[30] See, for example, Dooley, Folkerts-Landau, and Garber (2004) and Caballero, Farhi, and Hammour (2004).

$$d\alpha^* = -\frac{(\alpha^* + \alpha - 1)X}{X^*/E + F}\, d\,(\frac{F}{X})$$

For the parameters we have constructed, a change of five percentage points in F/Y requires an increase in the share of U.S. assets in foreign portfolios of about 0.8 percentage points a year (0.47×5 percent$/3$)[31].

We find more plausible the argument that the relative demand for U.S. assets may actually decrease rather than increase in the future. This argument is based, in particular, on the fact that much of the recent accumulation of U.S. assets has taken the form of accumulation of reserves by the Japanese and the Chinese central banks. Many worry that this will not last, that the pegging of the renminbi will come to an end, or that both central banks will want to change the composition of their reserves away from U.S. assets, leading to further depreciation of the dollar. Our model provides a simple way of discussing the issue and thinking about the numbers.

Consider pegging first: the foreign central bank buys or sells dollar assets so as to keep $E = \bar{E}$[32]. Let B denote the reserves (U.S. assets) held by the foreign central bank, so that

$$X = B + \alpha(1)(X - F) + (1 - \alpha^*(1))(\frac{X^*}{E} + F)$$

Figure 7 illustrates the resulting dynamics. Suppose that, in the absence of pegging, the steady state is given by point A and that the foreign central bank pegs the exchange rate at $\bar{E}$. At that level the U.S. current account is in deficit, and so F increases over time. Wealth gets steadily transferred to the foreign country, and so the private demand for U.S. assets steadily decreases. To keep E unchanged, B must increase further over time. Pegging by the

[31] A related argument is that, to the extent that the rest of the world is growing faster than the United States, an increase in the ratio of net debt to GDP in the United States is consistent with a constant share of U.S. assets in foreign portfolios. This argument falls quantitatively short: although some Asian countries are growing fast, their weight and their financial wealth are still far too small to absorb the U.S. current account deficit while maintaining constant shares of U.S. assets in their portfolios.

[32] Our two-country model has only one foreign central bank, and so we cannot discuss what happens if one foreign bank pegs its currency and the others do not. The issue is, however, relevant in thinking about the paths of the dollar-euro and the dollar-yen exchange rates. We discuss this further in the next section.

foreign central bank is thus equivalent to a continuous outward shift in the portfolio balance schedule: in effect, the foreign central bank is keeping world demand for U.S. assets unchanged by offsetting the fall in private demand. Pegging leads to a steady increase in U.S. net debt and a steady increase in the foreign central bank's reserves, offsetting the steady decrease in private demand for U.S. assets (represented by the path DC in figure 7). What happens when the foreign central bank unexpectedly stops pegging? From point C just before the peg is abandoned, the economy jumps to point G (recall that valuation effects lead to a decrease in net debt, and therefore a capital loss for the foreign central bank, when there is an unexpected depreciation) and then adjusts along the saddle-point path $GA\prime$. The longer the peg lasts, the larger the initial and the eventual depreciation.

In other words, an early end to the Chinese peg would obviously lead to a depreciation of the dollar (an appreciation of the renminbi). But the sooner it takes place, the smaller the required depreciation, both initially and in the long run. Put another way, the longer the Chinese wait to abandon the peg, the larger the eventual appreciation of the renminbi.

The conclusions are very similar with respect to changes in the composition of reserves. We can think of such changes as changes in portfolio preferences, this time not by private investors but by central banks, and so we can apply our earlier analysis directly. A shift away from U.S. assets will lead to an initial depreciation, leading in turn to a lower current account deficit, a smaller increase in net debt, and thus to a smaller depreciation in the long run.

How large might these shifts be? Chinese reserves currently equal $610 billion, and Japanese reserves are $840 billion. Assuming that these reserves are now held mostly in dollars, if the People's Bank of China and the Bank of Japan reduced their dollar holdings to half of their portfolio, this would represent a decrease in the share of U.S. assets in total foreign (private and central bank) portfolios, $(1 - \alpha^*)$, from 30 percent to 28 percent. The computations we presented earlier suggest that this would be a substantial shift, leading to a decrease in the dollar exchange rate possibly as large as 8.7 percent.

To summarize: Avoiding a depreciation of the dollar would require a steady and substantial increase in shares of U.S. assets in U.S. or foreign portfolios at a given exchange rate. This seems unlikely to hold for very long. A more likely scenario is the opposite, a decrease in shares, due in particular to diversification of reserves by central banks. If and when this happens, the

dollar will depreciate. Note, however, that the larger the adverse shift, the larger the initial depreciation but the smaller the accumulation of debt thereafter, and therefore the smaller the eventual depreciation. "Bad news" on the dollar now may well be good news in the long run (and vice versa).

3.3.6 The Path of Interest Rates

Our model takes interest rates as given, and the discussion thus far has taken them as constant[33]. Yield curves in the United States, Europe, and Japan indeed indicate little expected change in interest rates over the near and the medium term. However, it is easy to think of scenarios where changes in interest rates play an important role, and this leads us to discuss the role of budget deficit reduction in the adjustment process.

First, however, we briefly show the effects of an increase in the U.S. interest rate in our model. Figure 8 shows the effects of an unexpected permanent increase in r over r^*. (In contrast to the case of perfect substitutability, it is possible for the two interest rates to differ even in the steady state.) The portfolio balance locus shifts upward: At a given level of net debt, U.S. assets are more attractive, and so the exchange rate increases. The current account balance locus shifts down: the higher interest rate implies larger payments on foreign holdings of U.S. assets and thus requires a larger trade surplus, and in turn a lower exchange rate. The adjustment path is given by ABC. In response to the increase in r, the economy jumps from point A to point B and then moves over time from point B to point C. As drawn, there is an appreciation initially, but, in general, the initial effect on the exchange rate is ambiguous. If gross liabilities are large, for example, the effect of higher interest payments on the current account balance may dominate the more conventional "overshooting" effects of increased attractiveness and lead to an initial depreciation rather than an appreciation. In either case the steady-state effect is greater net debt accumulation, and thus a larger depreciation than if r had not increased.

Thus, under the assumption that an increase in interest rates leads initially to an appreciation, an increase in U.S. interest rates beyond what is already implicit in the yield curve would

[33] Remember that, when financial assets are imperfect substitutes, the interest rate differential no longer directly reflects expected exchange rate changes. It is thus perfectly rational for the level of long-term interest rates in the United States and in other countries to be very similar, even as the market anticipates a depreciation of the dollar. Therefore, if we consider that financial assets denominated in different currencies can be imperfect substitutes, there is no "interest rate puzzle," contrary to what is sometimes claimed in the financial press.

delay the depreciation of the dollar, at the cost of greater net debt accumulation and a larger eventual depreciation.

Interest rate changes however do not take place in a vacuum. It is more interesting to think about what may happen to interest rates as the dollar depreciates, either slowly along the saddle path or more sharply, in response, for example, to adverse portfolio shifts. As the dollar depreciates, relative demand shifts toward U.S. goods, reducing the trade deficit but also increasing total demand for U.S. goods. Suppose also that output is initially at its natural level (the level associated with the natural rate of unemployment), which appears to be a good description of the United States today. Three outcomes are possible:

- Interest rates and fiscal policy remain unchanged. The increase in demand leads to an increase in output but also an increase in imports, which partly offsets the effect of the depreciation on the trade balance. (In terms of our model, it leads to an increase in domestic spending, Z, and thus to a shift in z.)

- Interest rates remain unchanged, but fiscal policy is adjusted to offset the increase in demand and leave output at its natural level; in other words the budget deficit is reduced so as to maintain internal balance.

- Fiscal policy remains unchanged, but the Federal Reserve increases interest rates so as to maintain output at its natural level. In this case, higher U.S. interest rates limit the extent of the depreciation and mitigate the current account deficit reduction. In doing so, however, they lead to larger net debt accumulation and to a larger eventual depreciation.

In short, an orderly reduction of the current account deficit—that is, one that occurs while maintaining internal balance—requires both a decrease in the exchange rate and a reduction in the budget deficit[34]. The two are not substitutes: the exchange rate depreciation is needed to

[34] Many of the discussions at Brookings in the late 1980s were about the relative roles of budget deficit reduction and exchange rate adjustment. For example, Sachs (1988) argued that "the budget deficit is the most important source of the trade deficit. Reducing the budget deficit would help reduce the trade deficit... [while] an attempt to reduce the trade deficit by a depreciating exchange rate induced by easier monetary policy would produce inflation with little benefit on the current account," a view consistent with the third scenario above. Cooper (1986), in a discussion of the policy package best suited to eliminate the U.S. imbalances, stated, "The drop in the dollar is an essential part of the policy package. The dollar's decline will help offset the fiscal contraction through expansion of net exports and help maintain overall U.S. economic activity at a satisfactory level," a view consistent with the second scenario.

achieve current account balance, and the budget deficit reduction is needed to maintain internal balance at the natural level of output[35]. (The frequently heard statement that deficit reduction would reduce the need for dollar depreciation leaves us puzzled.) If the decrease in the budget deficit is not accompanied by a depreciation, the result is likely to be lower demand and a recession. Although the recession would reduce the current account deficit, this is hardly a desirable outcome. If the depreciation is not accompanied by a reduction in the budget deficit, one of two things can happen: demand will increase, and with it the risk that the economy will overheat, or, more likely, interest rates will increase so as to maintain internal balance. This increase would either limit or delay the depreciation of the dollar, but, as we have made clear, this would be a mixed blessing. Such a delay implies less depreciation in the short run but more net debt accumulation and more depreciation in the long run.

3.4 The Euro, the Yen, and the Renminbi

The depreciation of the dollar since the peak of 2002 has been very unevenly distributed: as of April 2005 the dollar had fallen 45 percent against the euro, 25 percent against the yen, and not at all against the renminbi. In this section we return to the questions asked in the introduction: if substantially more depreciation is indeed to come, against which currencies will the dollar fall? If China abandons its peg, or if Asian central banks diversify their reserves, how will the euro and the yen be affected?

The basic answer is simple. Along the adjustment path, what matters—because of home bias in asset preferences—is the reallocation of wealth across countries, and thus the bilateral current account balances of the United States with its partners. Wealth transfers modify countries' relative demands for assets, thus requiring corresponding exchange rate movements. Other things equal, countries with larger trade surpluses with the United States will see a larger appreciation of their currency.

Other things may not be equal, however. Depending on portfolio preferences, a transfer of wealth from the United States to Japan, for example, may change the relative demand for euro assets and thus the euro exchange rate. In that context one can think of central banks as

[35] Obstfeld and Rogoff (2004) emphasize a similar point.

investors with different asset preferences. For example, a central bank that holds most of its reserves in dollars can be thought of as an investor with strong dollar preferences. Any increase in its reserves is likely to lead to an increase in the relative demand for dollar assets and thus an appreciation of the dollar. Any diversification of its reserves is likely to lead to a depreciation of the dollar.

It is beyond the scope of this paper to construct and simulate a realistic multicountry portfolio model. But we can make some progress in thinking about mechanisms and magnitudes. The first step is to extend our model to allow for more countries.

3.4.1 Extending the Portfolio Model to Four Regions

In 2004 the U.S. trade deficit in goods (the only component of the current account for which a decomposition of the deficit by country is available) was $665 billion. Of this, $162 billion was with China, $77 billion with Japan, $85 billion with the euro area, and the remainder, $341 billion, with the rest of the world. We ignore the rest of the world here and think of the world as composed of four countries or regions: the United States, Europe, Japan, and China (indexed 1 through 4, respectively). We shall therefore think of China as accounting for roughly half the U.S. current account deficit, and Europe and Japan as accounting each for roughly one-fourth.

We extend our portfolio model as follows. We assume that the share of asset j in the portfolio of country i is given by

$$\alpha_{ij}(.) = a_{ij} + \sum_k \beta_{ijk} \, R_k^e$$

where R_k^e is the expected gross real rate of return, in dollars, from holding assets of country k (so that R_k^e denotes a rate of return, not a relative rate of return as in our two-country model).

We assume further that $b_{ijk} = b_{jk}$, so that the effect of the return on asset k on demand for asset j is the same for all investors, independent of the country of origin. This implies that differences in portfolio preferences across countries show up only as different constant terms, and derivatives with respect to rates of return are the same across countries.

The following restrictions apply: From the budget constraint (the condition that the shares sum to 1, for any set of expected rates of return), it follows that $\sum_j a_{ij} = 1$ for all i, and

$\sum_j \beta_{jk} = 0$ for all k. The home bias assumption takes the form: $\sum_i a_{ii} > 1$. The demand functions are assumed to be homogenous of degree zero in expected gross rates of return, so $\sum_k \beta_{jk} = 0$ for all j.

Domestic interest rates, in domestic currency, are assumed to be constant and all equal to r. Exchange rates, E_k, are defined as the price of U.S. goods in terms of foreign goods (so $E_1 = 1$, and an increase in E_2 for example indicates an appreciation of the dollar vis a vis the euro--equivalently, a depreciation of the euro vis a vis the dollar.) It follows that the expected gross real rate of return, in dollars, from holding assets of country k is given by $R_k^e = (1+r)E_k/E_{k,+1}^e$. In steady state, $R_k^e = (1+r)$, so that $\sum_k \beta_{jk} R_k^e = 0$ and we can concentrate on the a_{ij} s. The portfolio balance conditions, absent central bank intervention, are given by:

$$\frac{X_j}{E_j} = \sum_i a_{ij} \left(\frac{X_i}{E_i} - F_i\right)$$

where F_i denotes the net foreign debt position of country i, so $\sum_i F_i = 0$.

So far we have treated all four countries symmetrically. China, however, is special in two respects: it enforces strict capital controls, and it pegs the renminbi to the dollar. We capture these two features as follows:

- We formalize capital controls as the assumption that $a_{4i} = a_{i4} = 0$ for all $i \neq 4$; that is, capital controls prevent Chinese residents from investing in foreign assets but also prevent investors outside China from acquiring Chinese assets[36].

- We assume that, to peg the renminbi-dollar exchange rate ($E_4 = 1$), the People's Bank of China passively acquires all dollars flowing into China: the wealth transfer from the United States to the euro area and Japan is thus the U.S. current account minus the fraction that is financed by the Chinese central bank: $dF_1 + dF_4 = -dF_2 - dF_3$.

[36] This ignores inflows of foreign direct investment into China, but since we are considering the financing of the U.S. current account deficit, this assumption is inconsequential for our analysis.

3.4.2 Some Simple Computations

Consider now an increase in U.S. net debt equal to dF_1. Assume that a share γ of the U.S. net debt is held by China. Assume that a fraction x of the remaining portion is held by the euro area and a fraction $(1-x)$ by Japan, so that the changes in net debt are given by

$$dF_2 = -x(1-\gamma)dF_1,\ dF_3 = -(1-x)(1-\gamma)dF_1,\ dF_4 = -\gamma dF_1$$

Assume further that China imposes capital controls and pegs the renminbi, that the remaining three economies are all of the same size, and that the matrix of a_{ij} 's is symmetric in the following way: $a_{ii} = a$ and $a_{ij} = c = (1-a)/2 < a$ for $i \neq j$[37]. In other words, investors want to put more than one-third of their portfolio into domestic assets (the conditions above imply $a > 1/3$) and allocate the rest of their portfolio equally among foreign assets.

Under these assumptions, $dE_4 = 0$ (because of pegging) and dE_2 and dE_3 are given by:

$$\begin{aligned}
\frac{dE_2}{dF_1} &= -\frac{(a-c)(1-\gamma)[x(1-a)+c(1-x)]}{(1-a)^2-c^2} + \frac{c\gamma}{1-a-c} \\
\frac{dE_3}{dF_1} &= -\frac{(a-c)(1-\gamma)[xc+(1-a)(1-x)]}{(1-a)^2-c^2} + \frac{c\gamma}{1-a-c}
\end{aligned}$$

Consider first the effects of γ, the share of U.S. net debt held by China.

- For $\gamma = 0$, dE_2/dF_1 and dE_3/dF_1 are both negative. Not surprisingly, an increase in U.S. net debt leads to a depreciation of the dollar vis a vis both the euro and the yen.
- As γ increases, the depreciation of the dollar vis a vis the euro and the yen becomes smaller. This, too, is not surprising. What may be more surprising however is that for high values of γ, the depreciation turns into an appreciation. For $\gamma = 1$ for example, the dollar appreciates vis a vis both the euro and the yen. The explanation is straightforward,

[37] The assumption of countries of equal size allows us to specify the matrix in a simple and transparent way. Allowing countries to differ in size, as they obviously do, would lead to a more complex, size-adjusted matrix; but the results would be unaffected.

and is found in portfolio preferences. The transfer of wealth from the United States to China is a transfer of wealth from U.S. investors, who are willing to hold dollar, euro and yen assets, to the PBC, who only holds dollars. This transfer to an investor with extreme dollar preferences leads to a relative increase in the demand for dollars, an appreciation of the dollar against both the euro and the yen.

Consider now the effects of x, the share of the U.S. net debt held by Europe, excluding the net debt held by China (for simplicity, we set γ equal to zero):

- Consider first the case where $x = 0$, so that the accumulation of net debt is entirely vis-à-vis Japan. In this case, it follows that $dE_3/dF_1 = 2\, dE_2/dF_1$. Both the yen and the euro appreciate against the dollar, with the yen appreciating twice as much as the euro. This result might again be surprising: why should a transfer of wealth from the United States to Japan lead to a change in the relative demand for euros? The answer is that it does not. The euro appreciates against the dollar, but depreciates against the yen. The real effective exchange rate of the euro remains unchanged.

- If $x = 1/2$ (which seems to correspond roughly to the ratio of trade deficits and thus to the relative accumulation of U.S. net debt today), then obviously the euro and the yen appreciate in the same proportion against the dollar.

This simple framework also allows us to think about what would happen if China stopped pegging, or diversified its reserves away from dollars, or relaxed capital controls on Chinese and foreign investors, or any combination of these. Suppose China stopped pegging but maintained capital controls. Because the end of the peg, together with the assumption of maintained capital controls, implies a zero Chinese surplus, the renminbi would have to appreciate against the dollar. From then on, reserves of the Chinese central bank would remain constant. So, as the United States continued to accumulate net debt vis-à-vis Japan and Europe, relative net debt vis-à-vis China would decrease. In terms of our model, γ, the proportion of U.S. net debt held by China, would decrease[38]. Building on our results, this would lead to a decrease in the

[38] Marginal γ, the proportion of the increase in U.S. net debt absorbed by China, would equal zero.

role of an investor with extreme dollar preferences, the People's Bank of China, and would lead to an appreciation of the euro and the yen.

Suppose instead that China diversified its reserves away from dollars. Then, again, the demand for euros and for yen would increase, leading to an appreciation of both currencies against the dollar.

To summarize: The trade deficits of the United States with Japan and the euro area imply an appreciation of both the yen and the euro against the dollar. For the time being, this effect is partly offset by the Chinese policies of pegging and keeping most of its reserves in dollars. If China were to give up its peg or to diversify its reserves, the euro and the yen would appreciate further against the dollar. This last argument is at odds with the often-heard statement that the Chinese peg has "increased the pressure on the euro," and that therefore the abandonment of the peg would remove some of the pressure, leading to a depreciation of the euro. We do not understand the logic behind that statement.

3.4.3 Two Simulations and a Look at Portfolios

We have looked so far at equilibrium for a given distribution of Fs. This distribution is endogenous, however, in our model, determined by trade deficits and portfolio preferences. We now report the results of two simulations of our extended model.

In the first simulation we keep the symmetric portfolio assumptions introduced above. We take the three economies to be of the same size, and we use the values for the portfolio parameters introduced above of 0.70 for a and 0.15 for c. We consider a shift in the U.S. trade deficit, with half of the change in the deficit falling on China, one-fourth on Japan, and one-fourth on the euro area. We assume that each country trades only with the United States, so that we can focus on the bilateral balances with the United States.

We perform this simulation under two alternative assumptions about Chinese policy. In both we assume capital controls, but in the first case we assume that China continues to peg the renminbi, and in the second we assume that the renminbi floats; together with the assumption of capital controls, this implies, as indicated above, a zero Chinese trade surplus.

The top panel of figure 9 presents the results. Because of symmetry, the responses of the euro and the yen are identical and thus represented by the same line. The lower line shows the

depreciation of the dollar against the euro and the yen when the renminbi floats. The higher locus shows the more limited depreciation of the dollar (and more limited appreciation of the euro and the yen) when the renminbi is pegged and the Chinese central bank accumulates dollars.

One may wonder whether the preferences of private investors are really symmetric however. Constructing portfolio shares for Japanese, European, and U.S. investors requires rather heroic assumptions. We have nevertheless given it a try, and the results are reported in table 1. Appendix B presents details of the construction.

Note in table 1 the much larger share of dollar assets in European than in Japanese portfolios. Note also the small share of Japanese assets held by euro-area investors relative to the share of euro-area assets held by Japanese investors (the difference is much larger than the difference in relative size of the two economies). Portfolio preferences appear indeed to be asymmetric.

To see what difference this asymmetry makes, the bottom panel of figure 9 presents results of a second simulation. This simulation is identical to that in the top panel but now takes into account the relative size of the three economies (the Xs) and uses the shares reported in table 1.

The main conclusion we draw from the bottom panel is that it looks very similar to the top, except that the dollar depreciates initially a bit more against the yen than against the euro. This difference is due to the larger share of dollar assets in European than in Japanese portfolios: a dollar transferred from the United States to Europe leads to a smaller decrease in the demand for U.S. assets than does a dollar transferred from the United States to Japan.

3.5 Summary and Conclusions

We have argued that there have been two main forces behind the large U.S. current account deficits of the past ten years: an increase in the U.S. demand for foreign goods, and an increase in the foreign demand for U.S. assets. The path of the dollar since the late 1990s can be explained as the reaction to these forces.

The shift in portfolio preferences toward U.S. assets manifested itself first, in the late 1990s, in the form of high private demand for U.S. equities, and more recently in the form of high

central bank demand for U.S. bonds. The shift in demand away from U.S. goods is often attributed to more rapid growth in the United States than in its trading partners. This appears, however, to have played only a limited role: the performance of import and export equations in macroeconometric models shows that activity variables and exchange rates explain only about 60 percent of the increase in the U.S. trade deficit, with unexplained time trends and residuals accounting for the rest. We interpret this as evidence of a shift in the U.S. trade balance relation.

Either shift could have induced the observed paths of the dollar and the U.S. current account only in a world where financial assets are imperfect substitutes. A shift in asset preferences could not account for these paths, because it would be meaningless in a world where assets are perfect substitutes. Nor can the shift in preferences for goods explain these paths, because with perfect substitutability such a shift—provided it were perceived as long lasting—would have induced a quicker and sharper depreciation of the exchange rate and a smaller increase in the current account than we have observed.

As a way of organizing our thoughts about the U.S. current account deficit and the dollar, we have studied a simple model characterized by imperfect substitutability both among goods and among assets. The model allows for valuation effects, whose relevance has recently been emphasized in a number of papers. The explicit integration of valuation effects in a model of imperfect substitutability is, we believe, novel.

We find that the degree of substitutability between assets does not affect the steady state. In other words, the eventual dollar depreciation induced by either shift is the same no matter how closely U.S. and foreign assets substitute for each other. But the degree of substitutability does play a central role in the dynamics of adjustment.

In contrast to the case of perfect substitutability between assets, an increase in U.S. demand for foreign goods leads to a limited depreciation initially, a potentially large and long-lasting current account deficit, and a slow and steady depreciation over time. An increase in foreign demand for U.S. assets leads to an initial appreciation, followed by a slow and steady depreciation.

The slow rate of dollar depreciation implied by imperfect substitutability contrasts with predictions by others of much more abrupt falls in the dollar in the near future. We show that,

in the absence of anticipated portfolio shifts, the anticipated rate of depreciation depends on the change in the ratio of U.S. net debt to U.S. assets: the faster the increase in net debt, the faster the decrease in the relative demand for U.S. assets, and therefore the higher the rate of depreciation needed to maintain portfolio balance. If we take the annual increase in the ratio of net debt to U.S. GDP to be 5 percent, we derive an upper bound on the anticipated annual rate of depreciation of 2.7 percent a year.

If, however, shares in U.S. assets in the portfolios of either U.S. or foreign investors are expected to decline, the anticipated depreciation can be much larger. If, for example, we anticipate that central banks will diversify their reserves away from dollars and, as a result, that the share of U.S. assets in foreign portfolios declines by 2 percent over the coming year, then the anticipated depreciation may be as large as 8.7 percent. This is obviously an upper bound, derived by assuming that private investors are willing to keep a constant share of their wealth in U.S. assets despite a high expected negative rate of return. Still, it implies that, under imperfect substitutability, and under the assumption that desired shares in U.S. assets will decrease, it is logically acceptable to predict a substantial depreciation of the dollar in the near future.

On the other hand, a further shift in investors' preferences toward dollar assets would slow down, or even reverse, the path of dollar depreciation. The relief, however, would only be temporary. It would lead to an initial appreciation, but the accompanying loss of competitiveness would speed up the accumulation of foreign debt. The long-run value of the dollar would be even lower. The argument that the United States, thanks to the attractiveness of its assets, can keep running large current account deficits with no effect on the dollar appears to overlook the long-run consequences of a large accumulation of external liabilities.

For basically the same reason, an increase in interest rates would be self-defeating. It might temporarily strengthen the dollar, but the depreciation eventually needed to restore equilibrium in the current account would be even larger—because (as in the case of a shift in portfolio preferences) the accumulation of foreign liabilities would accelerate, and eventually the United States would need to finance a larger flow of interest payments abroad. A better mix would be a decrease in interest rates and a reduction in budget deficits to avoid overheating. (To state the obvious: tighter fiscal policy is needed to reduce the current account deficit, but it is not a

substitute for the dollar depreciation. Both are needed.)

The same will happen so long as China keeps pegging the exchange rate. One should think of the People's Bank of China as a special investor whose presence has the effect of raising the portfolio share of the world outside the United States invested in dollar assets. The longer the Chinese central bank intervenes, the larger this share. Sooner or later, however—as in the case of Korea in the late 1980s—the People's Bank of China will find it increasingly difficult to sterilize the accumulation of reserves. Eventually, when the peg is abandoned, the depreciation of the dollar will be larger, the longer the peg will have lasted, because in the process the United States will have accumulated larger quantities of foreign liabilities. Thus, if China is worried about a loss of competitiveness, pegging may be a myopic choice.

What would abandonment of the Chinese peg imply for the euro and the yen? Contrary to a commonly heard argument, if the renminbi were allowed to float, both currencies would be likely to appreciate further against the dollar. The reason is that, when the People's Bank of China stops intervening, the market effectively loses an investor with extreme dollar preferences, to be replaced by private investors with less extreme preferences. A similar argument holds if the People's Bank of China diversifies its reserves away from dollar assets. For Europe and Japan, however, what matter are effective exchange rates, and their currencies may well depreciate in effective terms even if they appreciate relative to the dollar in bilateral terms.

We end with one more general remark. A large fall in the dollar would not by itself be a catastrophe for the United States. It would lead to higher demand for U.S. goods and higher output, and it would offer the opportunity to reduce budget deficits without triggering a recession. The danger is more serious for Japan and Europe, which suffers from low growth already, and has little room for use of expansionary fiscal or monetary policy at this stage.

Figure 1. Determination of Exchange Rate and Net Debt in Steady State

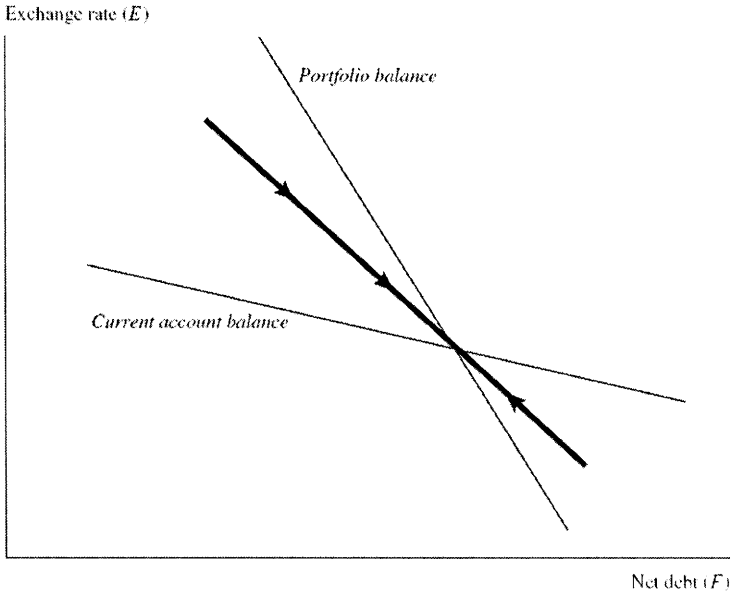

Source: Authors' model described in the text.

Figure 2. Adjustment of Exchange Rate and Net Debt to an Increase in z

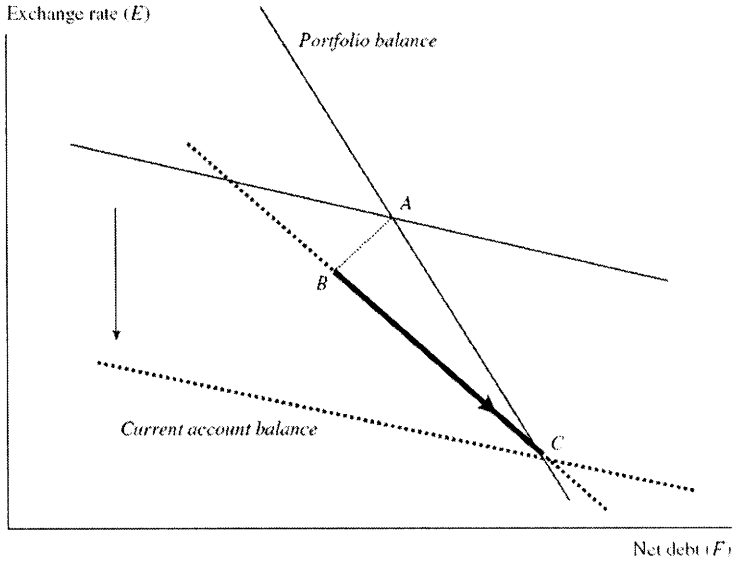

Source: Authors' model described in the text.

Figure 3. Responses of the Exchange Rate and Net Debt to a Shift in z

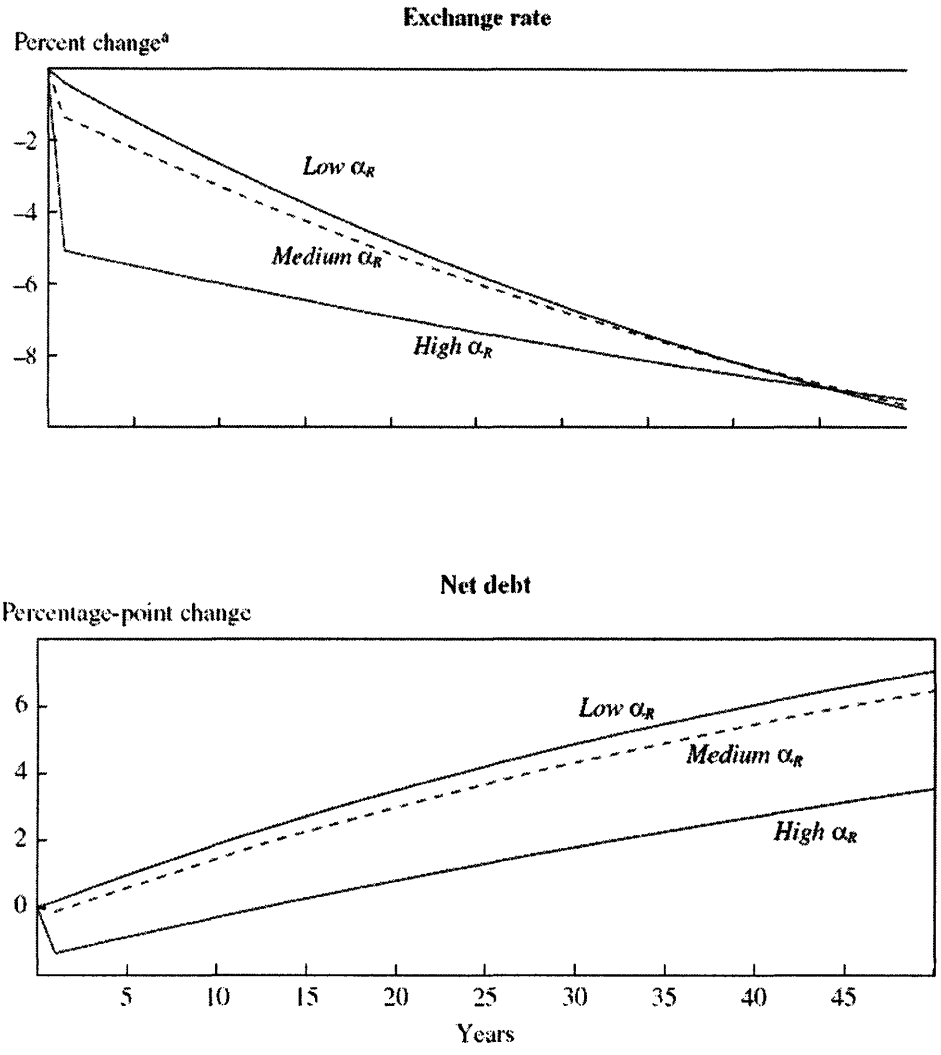

Source: Authors' calculations.
a. All simulations are for a shift in z of 1 percent of U.S. GDP.

Figure 4. Adjustment of Exchange Rate and Net Debt to an Increase in s

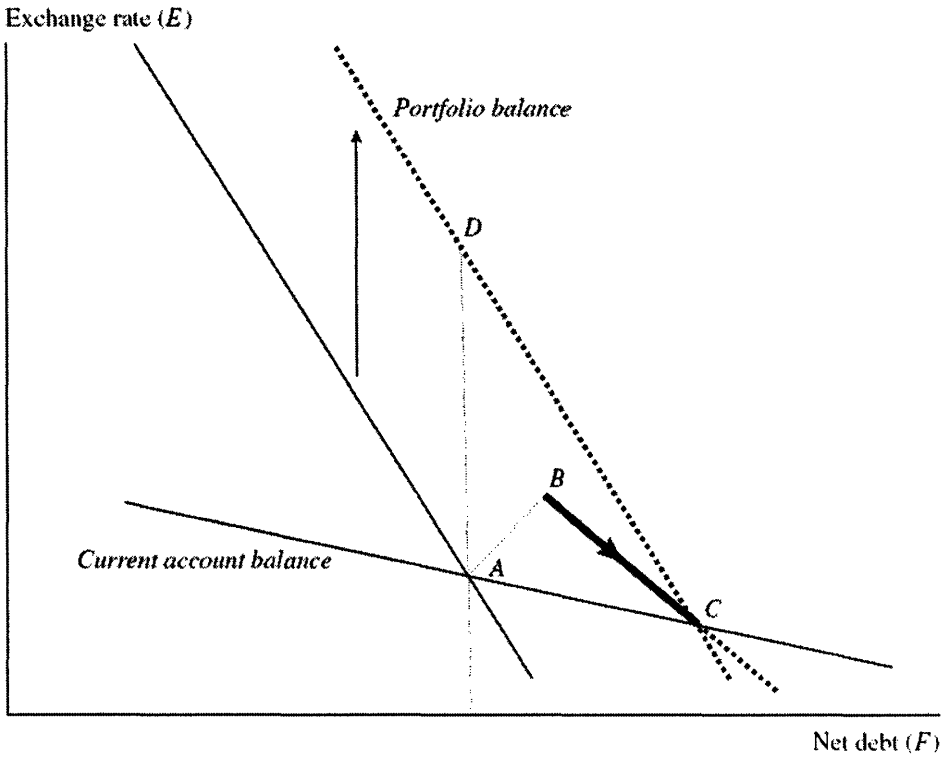

Source: Authors' model described in the text.

Figure 5. Responses of the Exchange Rates to a Shift in *s*

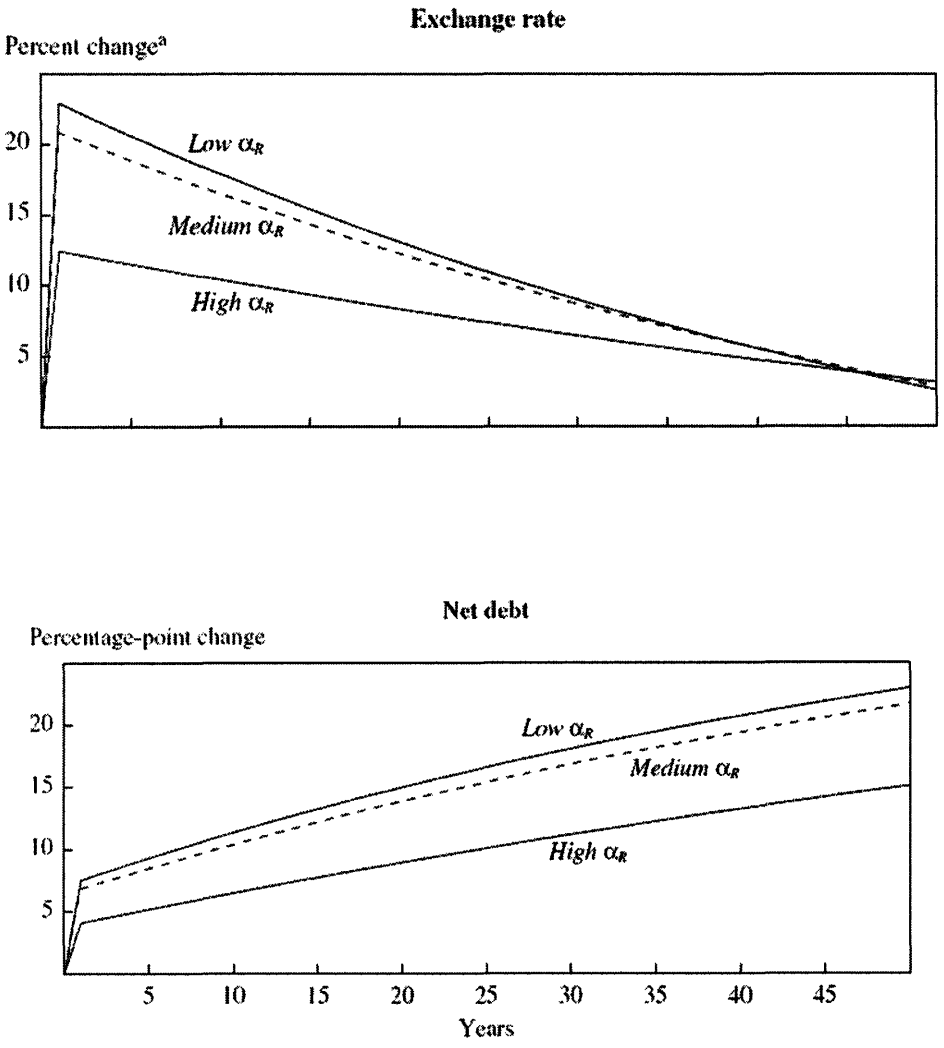

Source: Authors' calculations.
a. All specifications are for a 5-percentage-point shift in *s*.

Figure 6. Current Account Deficit and Effective Real Exchange Rate, 1978–93 and 1995–2004

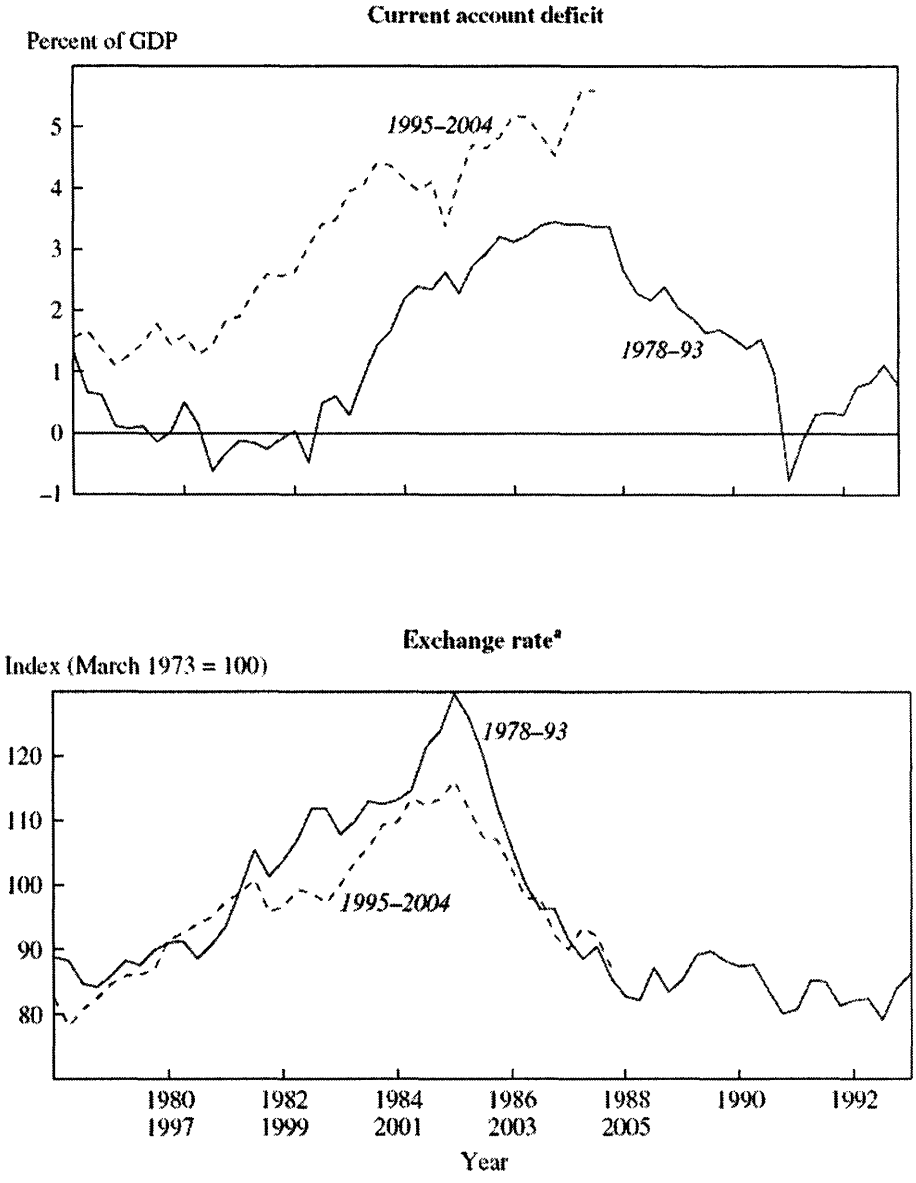

Source: Bureau of Economic Analysis, Table 1, U.S. International Transactions; Federal Reserve data.
a. Price-adjusted Major Currencies index.

Figure 7. Adjustment of Exchange Rate and Net Debt to Abandonment of Foreign Peg

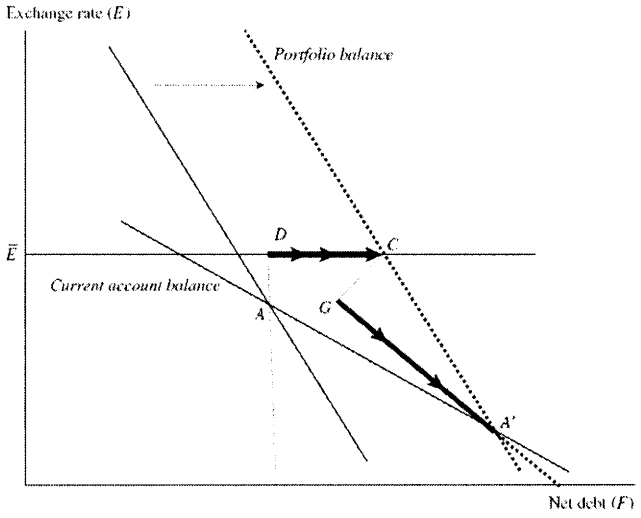

Source: Authors' model described in the text.

Figure 8. Adjustment of Exchange Rate and Net Debt to an Increase in the Domestic Interest Rate

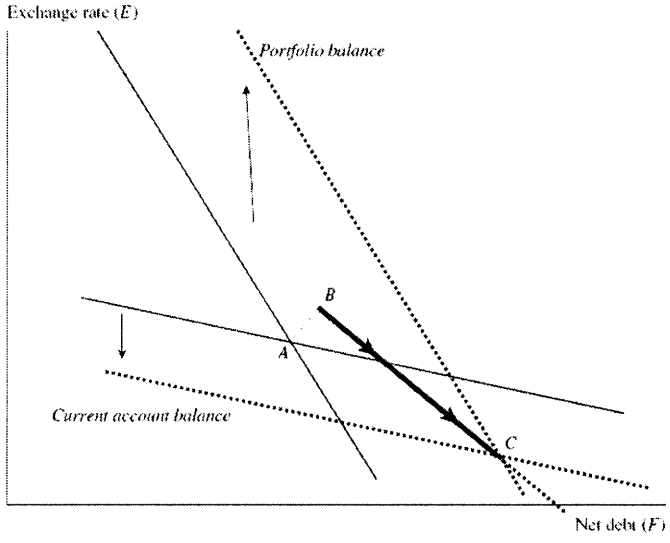

Source: Authors' model described in the text.

Figure 9. Effects of a Shift in the U.S. Trade Deficit on Euro-Dollar and Yen-Dollar Exchange Rates, with and without Chinese Peg[a]

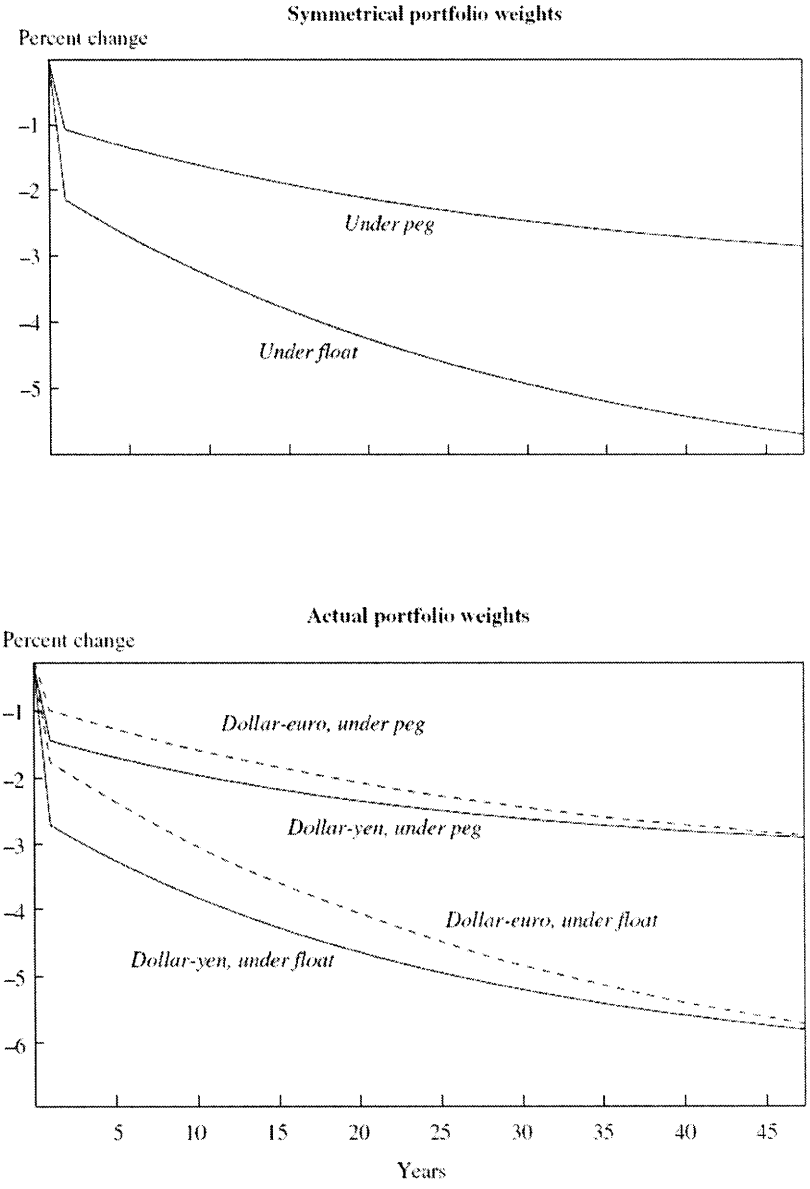

Source: Authors' calculations.
a. All simulations assume that China maintains capital controls.

Figure A-1. Adjustment of Exchange Rate and Net Debt in Continuous Time

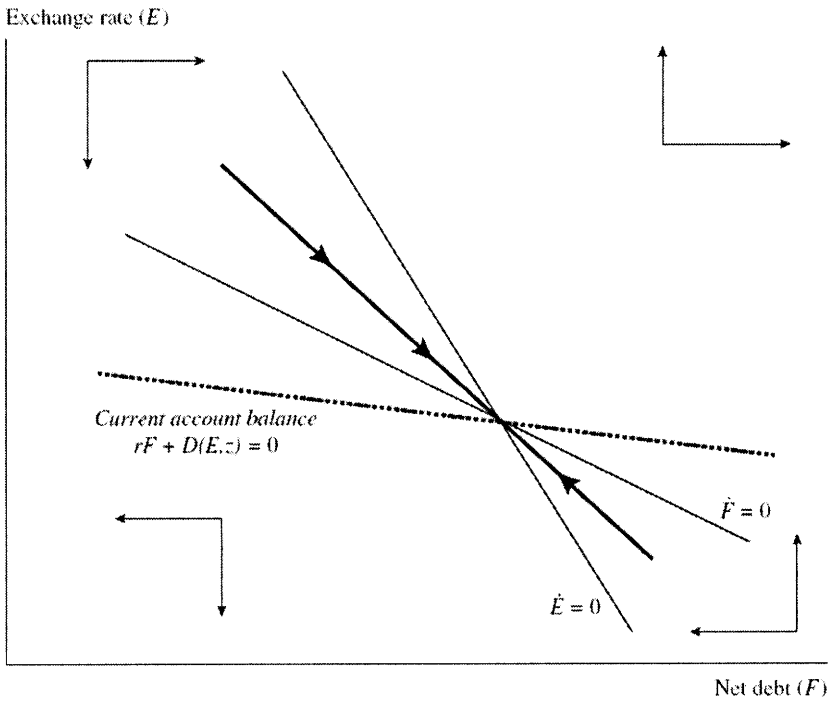

Source: Authors' model described in the text.

Table B-1. Calculated Portfolio Shares by Investment Destination[a]

	Investing country		
Destination	*United States*	*Euro area*	*Japan*
United States	0.77	0.19	0.17
Euro area	0.08	0.53	0.12
Japan	0.04	0.02	0.63
Rest of the world	0.11	0.27	0.08

Sources: Authors' calculations using data from the International Monetary Fund, the Organization for Economic Cooperation and Development, and national central banks.

a. Investment includes both portfolio investment and foreign direct investment. Shares may not sum to 1.00 because of rounding.

Appendix A: Dynamics of the Model

The dynamics of the system composed of equations (3.2) and (3.3) are more easily characterized by taking the continuous time limit. In continuous time the portfolio and current account balance equations become, respectively,

$$X = \alpha(1 + r - r^* + \frac{\dot{E}^e}{E}, s)\ (X - F) + (1 - \alpha^*(1 + r - r^* + \frac{\dot{E}^e}{E}), s)\ (\frac{X^*}{E} + F)$$

$$\dot{F} = rF + (1 - \alpha(1 + r + r^* + \frac{\dot{E}^e}{E}), s)\ \frac{\dot{E}}{E}\ (X - F) + D(E, z)$$

Note the presence of both expected and actual depreciation in the current account balance equation. Expected appreciation determines the share of the U.S. portfolio invested in foreign assets; actual appreciation determines the change in the value of that portfolio, and in turn the change in the U.S. net debt position.

We limit ourselves to a characterization of the equilibrium and local dynamics, using a phase diagram. (The global dynamics are more complex. The nonlinearities imbedded in the equations imply that the economy is likely to have two equilibria, only one of which is potentially saddle-point stable. This is the equilibrium we focus on.) We do so here under the additional assumption that $r = r^*$. The extension to differences in interest rates, which we used to construct figure 8, is straightforward.

The locus ($\dot{E} = \dot{E}^e = 0$) is obtained from the portfolio balance equation, and is downward sloping. In the presence of home bias, an increase in net debt shifts wealth abroad, decreasing the demand for U.S. assets, and requiring a depreciation.

The locus ($\dot{F} = 0$) is obtained by assuming ($\dot{E}^e = \dot{E}$) in the current account balance relation and replacing ($\dot{E}^e$) by its implied value from the portfolio balance equation. This locus is also downward sloping: a depreciation leads to a smaller trade deficit and thus allows for a larger net debt position consistent with current account balance.

Note that the locus ($\dot{F} = 0$) is not the same as the current account balance locus in figure 1; that locus is derived under the assumption that both $\dot{F}$ and $\dot{E}$ are zero. Using that locus makes for a simple graphical characterization of the equilibrium but is not appropriate for studying

stability or dynamics.

The derivatives α_R and α_R^* do not affect the slope of the locus $\dot{E} = 0$ but do affect that of the locus $\dot{F} = 0$. The smaller these derivatives are (that is, the lower the degree of substitutability between assets), the closer the locus ($\dot{F} = 0$) is to the locus ($\dot{E} = 0$). In the limit, if the degree of substitutability between U.S. and foreign assets is zero, the two loci coincide. The larger these derivatives are (that is, the higher the degree of substitutability between assets), the closer the ($\dot{F} = 0$) locus is to the current account balance locus: $0 = rF + D(E)$.

The condition for the equilibrium to be saddle-point stable is that the locus ($\dot{E} = 0$) be steeper than the locus ($\dot{F} = 0$); this turns out to be the same as the condition given in the text, that the portfolio balance locus be steeper than the current account balance locus. For this to hold, the following condition must be satisfied:

$$\frac{r}{ED_E} < \frac{\alpha + \alpha^* - 1}{(1 - \alpha^*)X^*/E}$$

The interpretation of this condition was given in the text. It is more likely to be satisfied, the lower the interest rate, the larger the home bias, and the larger the response of the trade balance to the exchange rate. If the condition is satisfied, the dynamics are as shown in figure A-1. The saddle path is downward sloping, implying that the adjustment to the steady state from below is associated with an expected depreciation, and the adjustment from above with an expected appreciation. Valuation effects imply that unexpected shifts in z or s are associated with initial changes in F, according to

$$\Delta F = (1 - \alpha)(1 + r^*)(X - F)\frac{\Delta E}{E}$$

The effect of the degree of substitutability on the dynamics is as follows. The smaller are α_R and α_R^*, the closer the locus ($\dot{F} = 0$) is to the locus ($\dot{E} = 0$), and so the closer the saddle point path is to the locus ($\dot{E} = 0$). In the limit, if the degree of substitutability between U.S. and foreign assets is zero, the two loci and the saddle point path coincide, and the economy remains on and adjusts along the ($\dot{E} = 0$) locus, the portfolio balance relation.

The larger α_R and α_R^*, the closer the ($\dot{F} = 0$) locus is to the locus given by $0 = rF + D(E)$, and the closer the saddle point path is to that locus as well. Also, the larger α_R and α_R^*, the

slower the adjustment of F and E over time. The slow adjustment of F comes from the fact that the current account is close to balance. The slow adjustment of E comes from the fact that, the larger the elasticities, the smaller is $\dot{E}$ for a given distance from the $\dot{E} = 0$ locus.

The limiting case of perfect substitutability is degenerate. The rate of adjustment to an unexpected, permanent shifts in z goes to zero. The economy is always on the locus $0 = rF + D(E)$. For any level of net debt, the exchange rate adjusts so net debt remains constant, and, in the absence of shocks, the economy stays at that point. There is no unique steady state, and where the economy is depends on history.

Appendix B: Construction of Portfolio Shares

Data on the country allocation of gross portfolio investment are from the International Monetary Fund's Coordinated Portfolio Survey for 2002. Data for the country allocation of direct investment are from the Organization for Economic Cooperation and Development and likewise refer to 2002. Financial wealth for the United States, the euro area, and Japan, which we need to compute the home bias of portfolios, are from the Flow of Funds data of the Federal Reserve Board[39].

From these data we construct a_{ij} in two steps. First, we compute the geographical allocation of net foreign investment positions by weighting the shares of portfolio assets and foreign direct investment allocated to country j by the relative importance of portfolio (pf) and direct investment (fdi) in country i's total investment abroad. We then scale these shares by the share of total foreign investment $(1 - a_{ii})$, so that

$$a_{ij} = [(pf_i/(pf_i + fdi_i))\ a_{ij,p} + (fdi_i/(pf_i + fdi_i))a_{ij,fdi}] * (1 - a_{ii})$$

Table B-1 presents the results.

To perform the simulation described in the text, we then allocate the shares invested in the "rest of the world" to foreign holdings so as to keep the relative shares in the remaining foreign assets the same. For the United States, for example, we increase the foreign shares in euro and yen assets to approximately 0.15 and 0.08, respectively. This gives us the numbers reported in table 1.

The simulation presented in figure 9 uses these values, together with asset levels of $36.8 trillion for the United States, $23.0 trillion for the euro area, and $8.0 trillion for Japan. Trade is assumed to be bilateral between the United States and each of the other regions, with elasticities of the trade balance all being equal to the elasticity used in our earlier two-country model.

[39] For the United States, see footnote 8. The source for Japan is the Bank of Japan flow of funds data (www.boj.or.jp/en/stat/sj/stat_f.htm), and that for the euro area is the ECB Economic Bulletin (released February, 2005 and available at www.ecb.int/pub/html/index.en.html).

Bibliography

[1] Baldwin, R., and P. Krugman (1987), 'The Persistence of the U.S. Trade Deficit', *BPEA*, No. 1, pp. 1–43.

[2] Bernanke, B. (2005), 'The Global Saving Glut and the U.S. Current Account Deficit', Homer Jones Lecture, Federal Reserve Bank of St. Louis, April 14. Available on the Internet at www.federalreserve.gov/boarddocs/speeches/2005/20050414/default.htm.

[3] Bernanke, B., V. Reinhart, and B. Sack (2004), 'Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment', *BPEA*, No. 2, pp. 1–78.

[4] Branson, W., and D. Henderson (1985), 'The Specification and Influence of Asset Markets', in *Handbook of International Economics II*, edited by R. Jones and P. Kenen, Amsterdam: Elsevier Science Publishers.

[5] Caballero, R., E. Farhi, and M. Hammour (2004), 'Speculative Growth: Hints from the U.S. Economy', Massachusetts Institute of Technology.

[6] Chinn, M. (2004), 'Incomes, Exchange Rates and the U.S. Trade Deficit, Once Again', *International Finance*, Vol. 7, No. 3, pp. 451–69.

[7] Cooper, R. (1986), 'Dealing with the Trade Deficit in a Floating Rate System', *BPEA*, No. 1, pp. 195–207.

[8] Debelle, G. and G. Galati (2005), 'Current Account Adjustment and Capital Flows', BIS Working Paper 169, Basel: Monetary and Economic Department, Bank for International Settlements.

[9] Dooley, M., D. Folkerts-Landau, and P. Garber (2004), 'The U.S. Current Account Deficit and Economic Development: Collateral for a Total Return Swap', National Bureau of Economic Research Working Paper 10727, Cambridge, Mass..

[10] Dornbusch, R. (1976), 'Expectations and Exchange Rate Dynamics', *Journal of Political Economy*, Vol. 84, pp. 1161–76.

[11] Dornbusch, R. (1987), 'External Balance Correction: Depreciation or Protection?', *BPEA*, No. 1, pp. 249–69.

[12] Gourinchas, P., and H. Rey (2005), 'International Financial Adjustment', National Bureau of Economic Research Working Paper 11155. Cambridge, Mass..

[13] Henderson, D., and K. Rogoff (1982), 'Negative Net Foreign Asset Positions and Stability in a World Portfolio Balance Model', *Journal of International Economics*, Vol. 13, pp. 85–104.

[14] Houthakker, H., and S. Magee (1969), 'Income and Price Elasticities in World Trade', *Review of Economics and Statistics*, Vol. 51, No. 2, pp. 111–25.

[15] Kouri, P. (1976), 'Capital Flows and the Dynamics of the Exchange Rate', Seminar Paper 67, Stockholm: Institute for International Economic Studies.

[16] Kouri, P. (1983), 'Balance of Payments and the Foreign Exchange Market: A Dynamic Partial Equilibrium Model', in *Economic Interdependence and Flexible Exchange Rates*, edited by J. Bhandari and B. Putnam, MIT Press.

[17] Krugman, P. (1991), 'Introduction', in *International Adjustment and Financing*, edited by C. Fred Bergsten and P. Krugman, Washington: Institute for International Economics.

[18] Lane, P., and G. Milesi-Ferretti (2002), 'Long-Term Capital Movements', in *NBER Macroeconomics Annual 2001*, edited by B. Bernanke and K. Rogoff, MIT Press.

[19] Lane, P., and G. Milesi-Ferretti (2004), 'Financial Globalization and Exchange Rates', CEPR Discussion Paper 4745, London: Centre for Economic Policy Research.

[20] Lawrence, R. (1990), 'U.S. Current Account Adjustment: An Appraisal', *BPEA*, No. 2, pp. 343–82.

[21] Marquez, J. (2000), 'The Puzzling Income Elasticity of U.S. Imports', Washington: Federal Reserve Board.

[22] Masson, P. (1981), 'Dynamic Stability of Portfolio Balance Models of the Exchange Rate', *Journal of International Economics*, Vol. 11, pp. 467–77.

[23] Obstfeld, M. (2004), 'External Adjustment', *Review of World Economics*, Vol. 140, No. 4, pp. 541–68.

[24] Obstfeld, M., and K. Rogoff (2004), 'The Unsustainable U.S. Current Account Position Revisited', National Bureau of Economic Research Working Paper 10869, Cambridge, Mass..

[25] Roubini, N., and B. Setser (2005), 'Will the Bretton Woods 2 Regime Unravel Soon? The Risk of a Hard Landing in 2005–2006', paper presented at a symposium on the 'Revived Bretton Woods System: A New Paradigm for Asian Development?', organized by the Federal Reserve Bank of San Francisco and the University of California, Berkeley, San Francisco, February 4, 2005.

[26] Sachs, J. (1988), 'Global Adjustments to a Shrinking U.S. Trade Deficit', *BPEA*, No. 2, pp. 639–67.

[27] Tille, C. (2003), 'The Impact of Exchange Rate Movements on U.S. Foreign Debt', *Issues in Economics and Finance*, Vol. 9, No. 1, pp. 1–7.