

NBER WORKING PAPER SERIES

LEARNING THE WEALTH OF NATIONS

Francisco J. Buera
Alexander Monge-Naranjo
Giorgio E. Primiceri

Working Paper 14595
<http://www.nber.org/papers/w14595>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2008

We would like to thank Daron Acemoglu, Graziella Bertocchi, V.V. Chari, Chad Jones, Ramon Marimon, Joel Mokyr, Juan Pablo Nicolini, Monika Piazzesi, Martin Schneider, Ruben Segura, Tao Zha and participants in many seminars and conferences. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by Francisco J. Buera, Alexander Monge-Naranjo, and Giorgio E. Primiceri. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Learning the Wealth of Nations
Francisco J. Buera, Alexander Monge-Naranjo, and Giorgio E. Primiceri
NBER Working Paper No. 14595
December 2008
JEL No. O4,O43,O50

ABSTRACT

We study the evolution of market-oriented policies over time and across countries. We consider a model in which own and neighbors' past experiences influence policy choices, through their effect on policymakers' beliefs. We estimate the model using a large panel of countries. We find that there is a strong geographical component to learning, which is crucial to explain the slow adoption of liberal policies during the postwar period. Our model also predicts that there would be a substantial reversal to state intervention if nowadays the world was hit by a shock of the size of the Great Depression.

Francisco J. Buera
Department of Economics
University of California, Los Angeles
8283 Bunche Hall Office 8357
Mail Stop: 147703
Los Angeles, CA 90095
and NBER
fjbuera@econ.ucla.edu

Giorgio E. Primiceri
Department of Economics
Northwestern University
3218 Andersen Hall
2001 Sheridan Road
Evanston, IL 60208-2600
and NBER
g-primiceri@northwestern.edu

Alexander Monge-Naranjo
Department of Economics
Pennsylvania State University
502 Kern Graduate Building
University Park, PA 16802-3306
aum26@psu.edu

LEARNING THE WEALTH OF NATIONS

FRANCISCO J. BUERA, ALEXANDER MONGE-NARANJO, AND GIORGIO E. PRIMICERI

ABSTRACT. We study the evolution of market-oriented policies over time and across countries. We consider a model in which own and neighbors' past experiences influence policy choices, through their effect on policymakers' beliefs. We estimate the model using a large panel of countries. We find that there is a strong geographical component to learning, which is crucial to explain the slow adoption of liberal policies during the postwar period. Our model also predicts that there would be a substantial reversal to state intervention if nowadays the world was hit by a shock of the size of the Great Depression.

1. INTRODUCTION

This paper studies the dynamics of economic policies across countries and over time. We focus on one important aspect of a government's economic policy decision, namely whether to be interventionist or instead favor a market-based allocation of resources. These issues are obviously crucial to the debate on growth and economic development. At the same time, there exist widely different views around the world about the relative benefits of the markets versus the state. Moreover, history shows that countries undergo cycles of liberalizations and reversals to state intervention, and viceversa. Our goal is understanding these transitions across regimes.

We explore one specific mechanism driving countries' transitions between regimes of *state intervention* and *market orientation*. This mechanism relies on the evolution of policymakers' beliefs about the desirability of free markets. We argue that these beliefs change because policymakers learn from their own past experience, as well as the experience of other countries.

Date: First version: September 2007. This version: November 2008. We would like to thank Daron Acemoglu, Graziella Bertocchi, V.V. Chari, Chad Jones, Ramon Marimon, Joel Mokyr, Juan Pablo Nicolini, Monika Piazzesi, Martin Schneider, Ruben Segura, Tao Zha and participants in many seminars and conferences.

The notion that past experience shapes beliefs and policy decisions is well rooted in policy circles. For instance, Stanley Fischer suggests the rise of socialist ideas in the 20th century as a clear case of interaction of outcomes, ideas and policies:

“It is not hard to see why views on the role of the state changed between 1914 and 1945 [...] A clear-headed look at the evidence of the last few decades at that point should have led most people to view the market model with suspicion, and a large role for the state with approbation- and it did.” [Fischer (1995), p. 102]

In a similar vein, the dismantlement of these ideas was a consequence of the change in views that “resulted from a combination and interaction of research and experience with development and development policy.” (Krueger (1997)) An important witness of this era also attests to this fact:

“I remember the foreign minister and the finance minister from another country saying to me: ‘You’re the first prime minister who is ever tried to roll back the frontiers of socialism. We want to know what’s going to happen. Because if you succeed others will follow.’” [Margaret Thatcher in Cran (2002)]

Despite this crucial connection between ideas, policies and economic development, there are very few formal treatments of these mechanisms in the literature. This paper fills this gap by proposing and estimating a model of the evolution of beliefs and policies where these interactions are at play.

In our model, countries start off with some prior beliefs about the effects of state intervention and market orientation on economic growth. Policymakers choose a regime of market orientation if they think that it is growth enhancing, and if the political costs entailed by this regime are not too large. However, the arrival of new information forces policymakers to revise their beliefs over time, which might lead to changes in policies. In updating their beliefs, policymakers pay also attention to the experience of other countries. However, they are allowed to discount the information coming from countries that are geographically, culturally or economically distant from the home country. In our model, it is optimal to do so if the effects of alternative policies are perceived to be country specific.

We estimate the parameters of the model (prior beliefs, structure of learning and the distribution of political costs) adopting data on the market orientation and GDP growth for a panel of 133 countries, over the years 1950-2001. We discipline our estimation procedure by imposing that countries' prior beliefs must be consistent with the data available for the 1900-1950 period.

We find that our simple learning mechanism rationalizes the evolution of market-oriented policies, correctly predicting 93% of the policy choices observed in the data and 60% of the policy changes within an interval of ± 3 years. In particular, the model captures well the early reversals towards state intervention typical of regions like Latin America, and the heterogeneous timing of liberalizations. To build intuition, it is useful to explain why, according to the model, many countries became market-oriented in the 1980s and 1990s. This is due to the growth slow-down of late 1970s and 1980s. At that time, most countries were state interventionists, and hence expectations about the benefits of state intervention must have been downgraded due to the low –and even negative– observed growth. Moreover, market-oriented economies performed better during those years, which must have further tilted beliefs in favor of market orientation.

We use the estimated model to answer several questions of interest. What are the main features of the learning model that are needed to explain the data? For example, how geographically localized is learning? How do cross-country information spillovers affect the diffusion of market-oriented policies? How dispersed are beliefs about the benefits of market-oriented policies by the end of our sample? How likely are policy reversals to state intervention after large negative global shocks to growth?

Our estimation results indicate that there is a strong geographical component to learning. The weight assigned to the experience of other countries declines by approximately 40% for every 1000Km (621.37 miles) of physical distance between countries. This turns out to be quite important since the median country has only four countries within a radius of 1000Km. Economic and cultural distance also matter in determining the relevant information neighborhood. More generally, according to our estimates, the localized nature of learning is essential to understand the slow adoption of free market doctrines around the world.

The localization of learning is also responsible for the substantial dispersion of countries' beliefs by the end of the sample period. In fact, even if the available global evidence favors

market orientation as a strategy for development, there still exists a considerable number of countries with negative views about the merits of market economies. Moreover, many of those with a favorable view are either quite uncertain or believe that the gains are quantitatively small. This motivates our concern that large growth shocks might induce a sequence of policy reversals to interventionism. Our concern is warranted: our model predicts that between 10% and 15% of the countries would revert to state intervention within 5 years, following a world-wide shock of the size of the Great Depression. Our exercises indicate that the evolution of beliefs is a central ingredient for the dynamics of policies.

Our paper relates to a large literature that studies the determinants of policy decisions. This literature mainly explores political economy dimensions, like redistributional issues, interest-group politics, the role of multilateral institutions (see, for example, Grossman and Helpman (1995), Krusell and Rios-Rull (1996) or Acemoglu, Johnson, and Robinson (2005)). Our paper complements this literature by studying the formation and evolution of beliefs about the benefits of different policies, as advocated by Kremer, Onatski, and Stock (2001). In this respect, our work is more closely related to Piketty (1995), Mukand and Rodrik (2005) and Strulovici (2008), where policy choices are related to the behavior of rational agents learning from past experience. Our focus, however, is on a formal exploration of the quantitative role played by the evolution of beliefs for the policy outcomes observed in the data. To this end, we purposely abstract from the political economy aspects, concentrating on tractable models of beliefs formation.

Models of policymakers as learning agents have been successfully applied to explain the rise and fall of US inflation (see, for instance, Sargent (1999), Cogley and Sargent (2005), Primiceri (2006) and Sargent, Williams, and Zha (2006)). Differently from this literature, policymakers in our model do not face a complex trade-off between alternative policy objectives. Our interest is instead in the role of learning spillovers among countries. From this point of view, our paper is also related to the literature on social learning and information spillovers in technology adoption and diffusion (see, for instance, Besley and Case (1994), Foster and Rosenzweig (1995) or Conley and Udry (2005)).

Finally, this paper draws from the empirical literature studying the behavior of countries income and growth (see, for instance, Jones (1997)) and their connection with trade

and market orientation policies (see Dollar (1992), Sachs and Warner (1995), Edwards (1998) or Rodriguez and Rodrik (2000)). However, most of this literature investigates the impact of policies on economic development, while our focus is exactly on the converse, i.e. understanding the determinants of market-oriented policies. In this sense, our objective is closer to Blattman, Clemens, and Williamson (2002), Clemens and Williamson (2004) and, to a lesser extent, Sachs and Warner (1995), although none of these papers stresses the importance of past growth performances for current policy choices.

The rest of the paper is organized as follows. Section 2 presents our theoretical model. Section 3 examines the dynamics and geography of market orientation and economic growth during the postwar period. Section 4 describes the estimation methodology. The estimation results and counterfactual exercises are discussed in sections 5 and 6 respectively. Section 7 contains some extensions and robustness checks and section 8 concludes.

2. MODEL

In this section we present a simple model in which the arrival of new information continuously reshapes policymakers' beliefs and affects policy decisions in each of the N countries of the world economy. At each point in time, a country adopts market-oriented policies if their perceived impact on growth overtakes their political and social cost.

2.1. The environment. In light of our available data, we simplify the economic policy problem to a dichotomic choice, in which countries choose to either rely on markets or state intervention. Policies are chosen period by period and we let $\theta_{i,t} \in \{0,1\}$ be an indicator variable that equals 1 if country i choose to be a market-oriented economy in period t and 0 otherwise.

Let $Y_{i,t}$ denote per-capita GDP in country i at time t and $y_{i,t} \equiv \log Y_{i,t} - \log Y_{i,t-1}$ its growth rate. Policymakers choose the sequence $\{\theta_{i,s}\}_{s=t}^{\infty}$ to maximize:

$$(2.1) \quad \max_{\{\theta_{i,s}\}_{s=t}^{\infty}} E_{i,t} \sum_{s=t}^{\infty} \delta^{s-t} [(1-\delta) \log Y_{i,s} - \theta_{i,s} K_{i,s}]$$

$$(2.2) \quad \text{s.t.} \quad y_{i,s} = \beta_i^S (1 - \theta_{i,s}) + \beta_i^M \theta_{i,s} + \varepsilon_{i,s}$$

$$(2.3) \quad \varepsilon_{i,s} \stackrel{i.i.d.}{\sim} N(0, \sigma_i^2)$$

$$(2.4) \quad K_{i,t} \stackrel{i.i.d.}{\sim} N(0, \sigma_{k,i}^2), \quad \text{all } s > t.$$

From the perspective of policymakers (but not necessarily from ours), the relationship between policy choices and GDP growth is given by equation (2.2), where β_i^S and β_i^M represent the average growth rates of country- i GDP under policies of state intervention and market orientation respectively, and $\varepsilon_{i,s}$ is a country specific shock.

Policy choices are also affected by a random utility term $K_{i,t}$. The variable $K_{i,t}$ is an exogenous random variable that captures the present value of political and social costs of market-oriented policies at time t . The constant $\sigma_{k,i}^2$ denotes the variance of these costs in country i . The random political cost $K_{i,t}$ is meant to capture variations in policies that are unexplained by our learning mechanism. In our estimation procedure, we treat $\left\{ \sigma_{k,i}^2 \right\}_{i=1}^N$ as a set of unknown parameters. Small estimates of $\sigma_{k,i}^2$ indicate that large variations in exogenous political costs are not needed to rationalize the observed dynamics of policies over time. In other words, we can interpret the estimates of $\sigma_{k,i}^2$ as a metric for the fit of the model for country i .¹

The information structure is as follows: policymakers do not know the value of β_i^S and β_i^M . However, for simplicity, we assume that they have perfect knowledge of all the remaining model parameters, including the variance of the growth shock $\varepsilon_{i,s}$.

The timing of events is as follows: at the end of time $t - 1$, policymakers in country i observe data on policy choices and GDP growth of all N countries, and update their beliefs about β_i^S and β_i^M . At the beginning of time t , they observe the realization of $K_{i,t}$ and decide what policy to adopt.

2.2. Optimal policy. We follow Sargent (1999) and assume that policymakers do not intentionally experiment. In other words, once they update their beliefs, policymakers solve the programming problem given by equations (2.1)-(2.4) as if these beliefs will not change in the future.² Given this assumption, optimal policy at time t is given by

$$\theta_{i,t} = 1 \left\{ E_{i,t-1}(\beta_i^M) - E_{i,t-1}(\beta_i^S) > K_{i,t} \right\},$$

¹ Section 7 presents the results of an alternative model in which policies determine countries' relative income and the speed of convergence to the technological frontier (Barro and Sala-I-Martin (1995)). Section 7 also analyzes the case in which the level of the political cost depends on a country's political environment (measured by the Polity2 variable) and inherits its autocorrelation.

² Therefore, expectations in (2.1) are taken with respect to the probability measure implied by this assumption.

where $1\{\cdot\}$ is the indicator function. Notice that the optimal policy decision only depends on the expected average growth rates and not the entire distribution of beliefs: policymakers choose to pursue market-oriented policies if their expectation of the difference between average growth under this regime and a regime of state intervention outweighs the political cost.

By abstracting from voluntary experimentation we rule out the possibility of a country opting for a policy regime perceived as detrimental for growth, but for which the country does not have much information, with the sole purpose of learning about it. We rule out experimentation for a number of reasons. First, while from a normative perspective experimenting might be beneficial, ours is a positive study of the behavior of governments. The extent to which governments conduct social experimentation on the grand scale remains an open question. Economists like Lucas (1981) or Blinder (1998) have argued against such behavior in policymaking. Moreover, standard voting schemes seem to reduce the occurrence of collective experimentation (see Strulovici (2008)). Second, the gains from experimentation are reduced in our setting because of two forces: (i) equations (2.1)-(2.4) constitute a sequential statistical decision problem (see El-Gamal and Sundaram (1993)), in which the presence of additional state variables (the political cost $K_{i,t}$) circumvents the incomplete learning results typical of the literature on multiarmed bandit problems (Berry and Fristedt (1985)); (ii) in our multicountry learning setting, policymakers benefit from the variation in the political costs of other countries, which should lower the incentives to experiment. Third and most important, the joint decision problem of N countries is extremely hard to solve, because it involves strategic experimentation motives.³

2.3. Learning. We assume that, in period $t = 0$, policymakers of country i start off with a Gaussian prior density on the unknown coefficients of (2.2), i.e. $\beta_i \equiv [\beta_i^S, \beta_i^M]'$. More precisely,

$$(2.5) \quad \beta_i \sim N\left(\hat{\beta}_{i,0} \ ; \ \sigma_i^2 \cdot P_{i,0}^{-1}\right),$$

³ Much simplified versions of this problem have been solved by Bolton and Harris (1999) and Keller, Rady, and Cripps (2005).

where N denotes the Normal distribution, while $\hat{\beta}_{i,0}$ and $\sigma_i^2 \cdot P_{i,0}^{-1}$ represent its expected value and variance-covariance matrix respectively. We choose the following parameterization for the inverse of the precision matrix $P_{i,0}$:

$$(2.6) \quad P_{i,0}^{-1} = \nu_i^2 \cdot \begin{bmatrix} 1 & \rho_i \\ \rho_i & 1 \end{bmatrix}.$$

Notice that we are making the simplifying assumption that the diagonal elements of the covariance matrix are the same. This means that policymakers start with a similar degree of uncertainty about the effects of government controls and market friendliness on economic growth. The coefficient ν_i^2 parameterizes this uncertainty.⁴

Priors are recursively updated with every new vintage of data. In updating their beliefs, policymakers of country i might use data from other countries, depending on how useful such data are perceived to be to learn about β_i . For example, if Argentinian policymakers believe that the effect of market-oriented policies in Argentina is fundamentally different from the rest of the world, they will update their beliefs using only Argentinian data. In the opposite extreme, if they believe that the growth effect of market orientation is approximately the same in the whole world, the data from every country will carry the same weight as Argentina's own data to update their beliefs. More plausibly, Argentinian policymakers might view the experience of nearby countries such as Brazil, Chile and Uruguay as informative, while discarding the data of more distant countries such as Japan, Russia and Zimbabwe.

In order to capture this idea in a flexible and convenient way, we assume that policymakers of country i believe that the relationship between policy choices and growth in other countries is described by the following equations:

$$(2.7) \quad y_{j,s} = \beta_{j|i,s}^S (1 - \theta_{j,s}) + \beta_{j|i,s}^M \theta_{j,s} + \varepsilon_{j|i,s}$$

$$(2.8) \quad \beta_{j|i,s}^S = \beta_i^S + \sigma_j \sqrt{q_{j|i}} \eta_{j,s}^S$$

$$(2.9) \quad \beta_{j|i,s}^M = \beta_i^M + \sigma_j \sqrt{q_{j|i}} \eta_{j,s}^M$$

$$\eta_{j,s}^S \sim N(0, 1)$$

$$(2.10) \quad \eta_{j,s}^M \sim N(0, 1), \quad j = 1, \dots, N$$

⁴ Observe that we do not force countries to start off with similar priors, in contrast with the literature on common priors in games (see Morris (1995) for a review). In our time series environment with country-specific signals, the common prior assumption is not a natural one.

where the subscript $j|i$ denotes country- i view about country- j variables.

Under this formulation, policymakers of country i believe that a relationship between policy decisions and growth similar to (2.2) is in place also in country j .⁵ However, they also believe that the average effects of policies on growth in other countries might differ from those in the home country. The variable $q_{j|i}$ parameterizes these discrepancies, by scaling the random differences $\eta_{j,s}^S$ and $\eta_{j,s}^M$.⁶ In this way, $q_{j|i}$ determines how useful data of country j are for country i . For instance, if $q_{j|i} = 0$, the consequences of market-oriented policies in country i and j are the same. Policymakers of country i would then use both sources of data symmetrically to update beliefs. As $q_{j|i}$ increases, data from country j become less and less informative about the growth effect of policy choices in country i .

We assume that $q_{j|i}$ is a parametric function of a vector of covariates z_{ji} :

$$q_{j|i} = \exp[-2 \cdot z_{ji}'\gamma] - 1.$$

We borrow this formulation from the literature on geographically weighted regressions (Fotheringham, Brunsdon, and Charlton (2002)). The vector z_{ji} may include various measures of geographic, cultural or economic distance between country i and j . This specification captures the idea that policymakers might attach more weight to countries that are closer geographically, culturally or economically to the home country.

Bayes law induces simple updating rules for $\hat{\beta}_{i,t} \equiv E_{i,t}([\beta_i^S, \beta_i^M]')$, i.e. the expectation of policymakers' beliefs in country i about the average effect of state intervention and market orientation on growth. Appendix A derives these updating rules and shows that this formulation of the problem is equivalent to a weighted least squares estimation problem, in which policymakers of country i assign a weight

$$(2.11) \quad w_{j|i} = \frac{\sigma_i}{\sigma_j} \exp[z_{ji}'\gamma].$$

to data coming from country j .

Finally, notice that in our model the weight $w_{j|i}$ is perfectly known to policymakers. An interesting alternative would be to assume that policymakers do not have perfect knowledge of the coefficients $\{\sigma_i\}_{i=1}^N$ and γ , and learn about the weight over time. Under these assumptions, however, the policymakers' estimation problem would resemble a system of

⁵ Notice that country i presumes that growth shocks are uncorrelated across countries. We make this assumption in our baseline model only for simplicity, but will relax it in section 7.

⁶ As we will see below, the random effect assumption greatly simplifies policymakers' inference because it implies tractable updating rules for their beliefs.

regressions with unequal variances, for which the posterior distribution is not available in closed form and can only be evaluated numerically (see, for example, Gelman, Carlin, Stern, and Rubin (2004)). For reasons of tractability, we have therefore opted for our simpler specification.

3. POSTWAR DYNAMICS AND GEOGRAPHY OF MARKET ORIENTATION

In the rest of the paper, we explore the ability of our model of beliefs formation to explain the observed choice of countries to follow market-oriented or state-interventionist policies. We use data from the large panel of liberalizations (133 countries, 1950-2001) originally constructed by Sachs and Warner (1995) and revised and extended by Wacziarg and Welch (2008).

In this section, we describe the data, their behavior over time and across regions, and use them to present evidence of the connection between policy choices and countries' past growth performance under alternative policy regimes. The findings of this section motivate the structural estimation exercise that we will undertake in sections 4 and 5.

3.1. Measuring market orientation. In order to investigate the dynamics of market orientation, we need a measure of policy that is comparable across countries and over time. As a proxy, we use the Sachs and Warner's (1995, hereafter SW) indicator of liberalizations. This indicator was originally constructed as a measure of openness to international trade, but it is better interpreted as a broader, albeit stark, measure of market orientation. Indeed, SW have been widely criticized because their indicator captures policy interventions that go much beyond restrictions on international trade (see, for example, Rodriguez and Rodrik (2000)).

SW argue that a government disposes of a variety of mechanisms to intervene in the economy and restrict international trade. The most direct mechanism, of course, is to impose tariffs and other barriers on imports. Other mechanisms include taxing, restricting or monopolizing exports and limiting or blocking the convertibility of the country's currency. Finally, a socialist government is likely to have significant distortions on international trade. Following this logic, SW require the following five criteria to classify a country as "open": (i) The average tariff rate on imports is below 40%; (ii) Non-tariff barriers cover less than 40% of imports; (iii) The country is not a socialist economy (according to the definition of Kornai (2002)); (iv) The state does not hold a monopoly of

the major exports; (v) The black market premium is below 20%. The resulting indicator is a dichotomic variable. If in a given year a country satisfies *all* of these five criteria, SW call it open and set the indicator to 1. Otherwise, the indicator takes the value of 0.⁷

Criteria (iii)-(v) capture forms of government intervention that go clearly beyond trade policies. This is why, following the interpretation of Rodriguez and Rodrik (2000), we dub a country as “market-oriented” if the SW indicator is equal to 1, and “state interventionist” if it is equal to 0. Moreover, notice that the broadness of the indicator is an advantage for us, not a concern. In fact, we are not interested in exploring the impact of trade policies on economic growth, but rather explaining countries’ policy decisions as a function of past growth performances. Therefore, the broader the policy measure we use the better.

In our study, we use the revision of the SW indicator carried out by Wacziarg and Welch (2008) who expanded the sample of countries and years and revised the information used to classify the countries.⁸ The data are an unbalanced panel, partly because separations and unifications have led to start or stop reporting some countries as independent entities.

Finally, it is important to recognize that the dichotomic nature of the SW indicator is an important limitation. Countries with very different degrees of state intervention end up being classified equally. Moreover, the indicator fails to capture reforms if they do not simultaneously move countries in all five criteria, e.g. China in later years. Unfortunately, richer indicators, such as those produced by the World Bank or the Heritage Foundation, are only available for a reduced sample of countries and a handful of recent years. The large coverage of countries and years in the SW indicator is essential to study changes in the orientation to markets, which, as we now show, happen only so often.

3.2. The dynamics of market orientation and growth. Figure 1 plots the evolution of the world fraction of market-oriented countries according to the SW indicator. In 1950, only about 30% of the countries qualify as market oriented. The spike in the late 1950s reflects the creation of the European Economic Community, followed by the inclusion in the sample of a number of state-interventionist developing countries. In the years from 1963

⁷ Since a complete panel dataset with all these criteria is not available for the entire postwar period, SW and Wacziarg and Welch (2008) perform a complementary country by country analysis of large reform episodes, in order to find out whether large changes in policies led a country not to meet (or start meeting) some of these criteria. This extended dataset is referred to as the “panel of liberalization dates.”

⁸ For example, SW apply criterion (iv) only for African countries while Wacziarg and Welch extended it to other regions. In general, for the same years as SW, the revision of Wacziarg and Welch lead more countries to be classified as market oriented relative to the original SW measure.

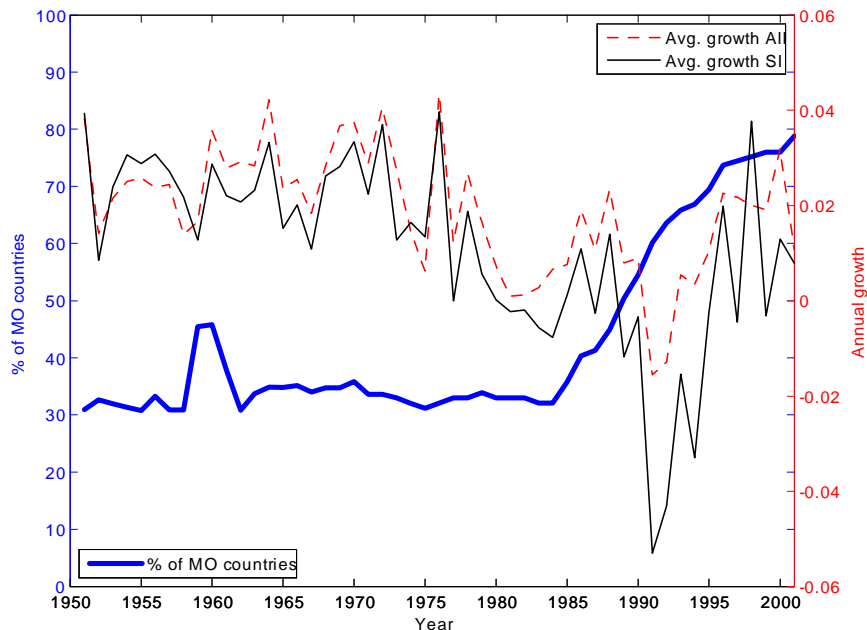


FIGURE 1. World average of market orientation and growth.

to 1984, the share of market-oriented economies increases only very marginally. However, in 1985 we observe the beginning of a global movement towards market orientation that continues until the end of the sample. By 2001, the fraction of market economies has increased to almost 80%.

The world-wide average hides interesting regional differences. Figure 2 displays the fraction of market economies (with scale on the right-hand-side axis) within eight regions of the world. Some of these regions started the sample as state interventionists and began a transition to market economies only later on. Earlier transitions towards market orientation include Western Europe, with the creation of the European Economic Community, and the Asian/Pacific region. On the contrary, North and Central America started relatively market oriented. However, after the establishment of the Central American Common Market and the diffusion of economic policies based on import substitutions, essentially only the U.S. and Canada remained market friendly between 1960 and 1985.

In order to relate policy choices and past growth performances, figure 1 also plots the (unweighted) average growth rate of per-capita GDP.⁹ Interestingly, many countries

⁹ GDP data are obtained from Penn World Table 6.2.

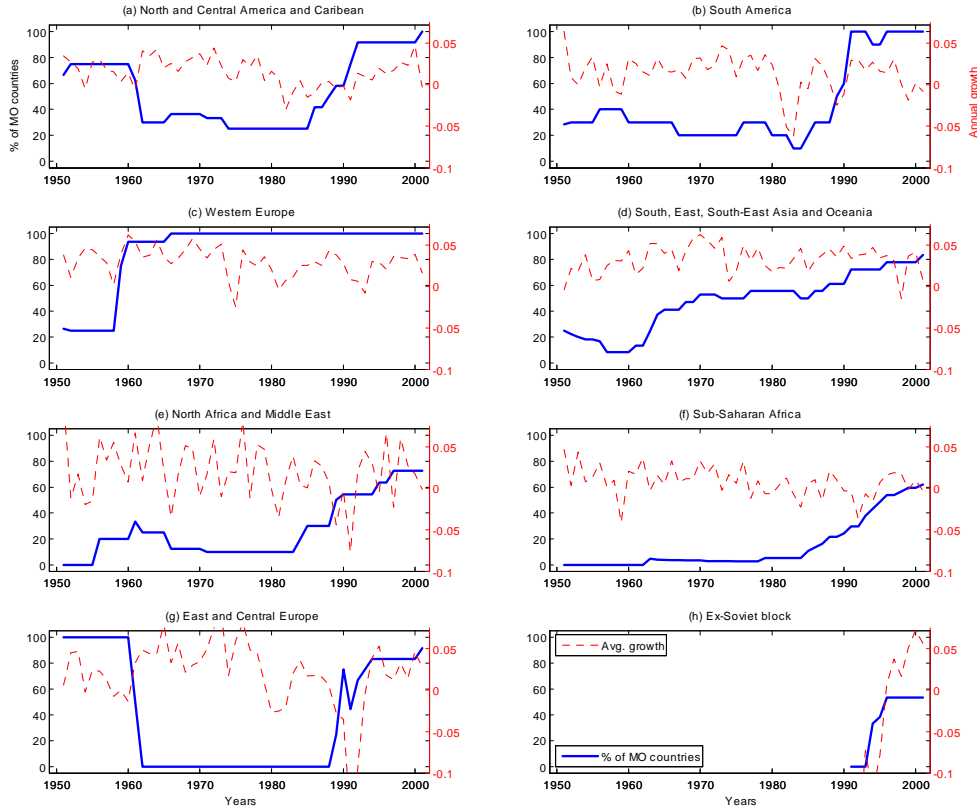


FIGURE 2. Regional patterns of market orientation and growth

started opening to the market only after the growth collapse of late 1970s and early 1980s, which was even more severe for state-interventionist economies. A similar pattern can be observed at the regional level (figure 2): most regions began a transition to market orientation after the deterioration of their growth performances. Our model of belief formation interprets those low growth episodes as the driving forces of the policy changes.

3.3. Reduced-form regressions. Before turning to structural estimation, we run reduced-form regressions to examine further the relationship between policy choices and observed growth rates.

Specifically, we consider the following linear probability model:¹⁰

$$(3.1) \quad E(\theta_{i,t}|\dots) = \phi_i + \phi_t + \phi_1\theta_{i,t-1} + \phi_2\hat{E}_{i,t-1}(y|\theta = 1) + \phi_3\hat{E}_{i,t-1}(y|\theta = 0) + \phi_4\bar{\theta}_{i,t-1}.$$

Here, the policy decision of country i in period t ($\theta_{i,t}$) is a function of its own past policy ($\theta_{i,t-1}$), a distance-weighted measure of other countries' policies ($\bar{\theta}_{i,t-1}$) and the distance-weighted average growth rate over the previous 3 years of other countries under the two policy regimes ($\hat{E}_{i,t-1}(y|\theta = 1)$ and $\hat{E}_{i,t-1}(y|\theta = 0)$).¹¹

Our theory predicts policies to be persistent ($\phi_1 > 0$), due to the persistence of beliefs implied by Bayesian updating. It also predicts that countries are more likely to pursue market-oriented policies in periods in which market-oriented neighbors grow faster ($\phi_2 > 0$), as a consequence of information spillovers. Similarly, in periods of faster growth of state-interventionist neighbors, the probability of being market oriented should decline ($\phi_3 < 0$). Finally, the precision of the information about the success of market economies depends on the number of countries following those policies. Therefore, our theory predicts a positive correlation among policies ($\phi_4 > 0$), provided that on average countries update their beliefs about the benefits of market orientation upwards (which will be the case according to the estimation results of our structural model). It is important to recognize, however, that finding $\phi_4 > 0$ can also be symptomatic of a force of “fads and fashion” in policy choices, which may also be in play in the data (see Banerjee (1992) or Bikhchandani, Hirshleifer, and Welch (1992)).

Table 1 presents the estimation results of equation 3.1 in our data. In the first column, we estimate a simplified version that excludes $\bar{\theta}_{i,t-1}$ from the regressors and sets $\delta_1 = 2,500$, which fixes the effective neighborhood of the median country, defined as $\sum_{j \neq i} e^{-d_{ij}/\delta_1}$, to be 20 countries. In columns 2-5 we progressively add various features to the model, by including country fixed and time effects (columns 2-5); optimizing over δ_1 ($=\delta_2$), estimating the effective country neighborhoods when measuring past growth performance and policy choices (columns 3-5); allowing for $\delta_1 \neq \delta_2$ in the estimation, which introduces differences in neighborhoods with regards to policies and growth performances

¹⁰ We have opted for a linear probability model because a probit with fixed effects is inconsistent (Chamberlain (1980)).

¹¹ Formally, we define: $\hat{E}_{i,t-1}(y|\theta) = \sum_{s=1}^3 \sum_{j:\theta_{j,t-s}=\theta} e^{-d_{ij}/\delta_1} y_{j,t-s} / \sum_{s=1}^3 \sum_{j:\theta_{j,t-s}=\theta} e^{-d_{ij}/\delta_1}$ $\theta = 0, 1$, and $\bar{\theta}_{i,t-1} = \sum_{j \neq i} e^{-d_{ij}/\delta_2} \theta_{j,t-1} / \sum_{j \neq i} e^{-d_{ij}/\delta_2}$.

(columns 4-5); and, finally, controlling for the political environment, as measured by the Polity2 variable (column 5).

Table 1 reveals three robust features of the data, which are consistent with the predictions of our structural model. First, policies are very persistent (second row). The probability that a country who was market oriented in period $t - 1$ keeps the same policy in period t is 85 to 95 percent larger than that of countries that were interventionist in period $t - 1$. Second, policy choices are highly correlated to past performance of policy regimes (third and fourth rows). For each additional point of per-capita GDP growth of market-oriented (state interventionist) countries in the neighborhood of country i , the long-run probability that country i is market oriented increases (decreases) by approximately 6 (9) percent in our specification with fixed and time effects. Finally, countries are more likely to be market oriented when their neighbors followed the same policies (fifth row).

	(1)	(2)	(3)	(4)	(5)
	OLS	FE-TE	FE-TE	FE-TE	FE-TE
constant	0.037** (0.005)				
$\theta_{i,t-1}$	0.958** (0.004)	0.88** (0.01)	0.86** (0.01)	0.85** (0.01)	0.84** (0.01)
$\hat{E}_{i,t-1}[y \theta = 1]$	0.041 (0.15)	0.40* (0.21)	0.50* (0.29)	1.00* (0.53)	1.03* (0.62)
$\hat{E}_{i,t-1}[y \theta = 0]$	-0.56** (0.11)	-0.62** (0.18)	-0.70** (0.26)	-1.50** (0.40)	-1.48* (0.60)
$\bar{\theta}_{i,t-1}$			0.14** (0.03)	0.08** (0.02)	0.10** (0.02)
polity IV					0.0014** (0.0005)
δ_1	2500	2500	3447.8** (502.3)	5948.3** (1515.3)	5042* (2148.2)
δ_2	δ_1	δ_1	δ_1	1139.3* (174.8)	1169.1* (411.1)
N	4755	4755	4755	4755	4441

TABLE 1: *Estimation results of reduced-form linear probability model. * (**)
significant at the 5% (1%) level.*

We conclude this section by noting that these results must be interpreted with caution, in light of the “reflection” problem emphasized by Manski (1993). Specifically, the identification of endogenous interactions effects (ϕ_4) and contextual effects (ϕ_2 and ϕ_3) is problematic due to the collinearity of the two effects. In our model, however, this problem is much less severe because neighborhoods are individual specific, which breaks the symmetry that causes the collinearity problem. On the other hand, the fact that we estimate the structure of groups adds a non-trivial element of complexity relative to more standard linear-in-mean models. Finally, the structural learning model that we study in the rest of the paper belongs to a class of binary choice models where identification is less problematic (Brock and Durlauf, 2001a and b).

We now turn to the estimation of this model.

4. INFERENCE

Like the agents of our model, we (the econometricians) are also Bayesian and wish to construct the posterior distribution for the unknown parameters of the model. These unknown coefficients are:

$$\begin{aligned}
 \left\{ \hat{\beta}_{j,0}^S \right\}_{j=0}^N & : \text{ expectations of initial beliefs about the effect of state intervention} \\
 \left\{ \hat{\beta}_{j,0}^M \right\}_{j=0}^N & : \text{ expectations of initial beliefs about the effect of market orientation} \\
 \left\{ \nu_j \right\}_{j=0}^N & : \text{ standard deviation of initial beliefs about the effect of SI and MO} \\
 \left\{ \rho_j \right\}_{j=0}^N & : \text{ correlation of initial beliefs about the effect of SI and MO} \\
 \left\{ \sigma_{j,k} \right\}_{j=1}^N & : \text{ standard deviation of the political cost} \\
 \gamma & : \text{ coefficients of the weighting function}
 \end{aligned}$$

If we collect the set of unknown coefficients in the vector α and denote by D the entire set of available data on policies and growth, standard application of Bayes rule delivers:

$$p(\alpha|D) \propto \mathcal{L}(D|\alpha) \cdot \pi(\alpha),$$

where $p(\cdot)$, $\mathcal{L}(\cdot)$ and $\pi(\cdot)$ represent the posterior, sampling and prior densities respectively, and \propto denote the proportionality relation. We now turn to the description of the priors and the construction of the likelihood function.

4.1. Priors. Our model is quite heavily parameterized. Thus, the use of somewhat informative priors helps to prevent overfitting problems. For instance, we would like to avoid cases in which we fit the data well, but only due to estimates of policymakers' initial beliefs which are clearly implausible. As an example, consider the literature on macroeconomic forecasting: highly parameterized models do well in-sample, but perform poorly out-of-sample. The use of priors considerably improves the forecasting performance of these models (see, for instance, Doan, Litterman, and Sims (1984), Litterman (1986) or, more recently, Banbura, Giannone, and Reichlin (2007)). The role of priors is similar in our context, as our ultimate goal is using the model to conduct a set of counterfactual experiments.

To begin with, in order to reduce the number of estimated parameters, we calibrate the correlation of initial beliefs about average growth under state intervention and market orientation ($\{\rho_j\}_{j=0}^N$) to zero. For the remaining parameters, we assume the following prior densities:

$$\begin{aligned}\pi(\hat{\beta}_{i,0}^S) &= N(\bar{\beta}_0^S, \omega^2), & i = 1, \dots, N \\ \pi(\hat{\beta}_{i,0}^M) &= N(\bar{\beta}_0^M, \omega^2), & i = 1, \dots, N \\ \pi(\nu_i) &= \text{IG}(s_\nu, d_\nu), & i = 1, \dots, N \\ \pi(\sigma_{i,k}) &= \text{IG}(s_\sigma, d_\sigma), & i = 1, \dots, N \\ \pi(\gamma) &= \text{Uniform.}\end{aligned}$$

These prior densities are parameterized as follows:

- We set $\bar{\beta}_0^S = 0.0275$ and $\bar{\beta}_0^M = 0.0125$. We have chosen these numbers using the Maddison data (Maddison (2006)). First of all, the average annual growth rate of per-capita GDP of countries in the Maddison dataset in 1901-1950 (excluding the years corresponding to the two wars) is approximately 2%. Then, we split the countries in the Maddison dataset in two groups, according to the value of the SW indicator in 1950. We find that, between 1946 and 1950, the state-interventionist countries grew on average 1.5% faster than the market-oriented. The values of $\bar{\beta}_0^S$ and $\bar{\beta}_0^M$ are then chosen so that $\bar{\beta}_0^S + \bar{\beta}_0^M = 2\%$ and $\bar{\beta}_0^S - \bar{\beta}_0^M = 1.5\%$. Notice that starting with a prior that most countries believed that state intervention fostered growth is consistent with the fact that only about 30% of the countries

were market oriented in 1950. These priors are also consistent with the evidence in Clemens and Williamson (2004). The value of ω is set to 0.025, which implies a quite agnostic view about the mean of initial beliefs.

- We select s_ν and d_ν so that ν_i has an a-priori mean and a standard deviation equal to 0.264. The prior on ν_i is potentially important because it affects the speed of learning, especially for those countries for which fewer data are available. In calibrating this prior we first observe that $\sigma_i^2 \cdot \nu_i^2$ should be approximately equal to $\text{var}(\bar{y}_i)$, the variance of the average growth rate of GDP.¹² We obtain an estimate of $\text{var}(\bar{y}_i) = 0.0175^2$ as the variance of the average growth rates of the countries present in the Maddison dataset between 1901 and 1950 (excluding the wars).¹³ To obtain an estimate of σ_i^2 based on pre-sample observations, we use again the Maddison data and run a regression of GDP growth on time and fixed effects. We then compute the variance of the residuals for each country and calculate the mean of these variances (which equals 0.0044). Therefore, we set the mean of the prior for ν_i to $\left[\frac{(0.0175)^2}{0.0044}\right]^{1/2} = 0.264$.
- We select s_σ and d_σ so that $\sigma_{i,k}$ has an a-priori mean and standard deviation equal to 0.01. The idea here is to discourage the model from fitting the data using very large variances of the exogenous political cost K_{it} . This prior distribution implies that, if policymakers believe that growth under market orientation is 1% higher than under state intervention, they will adopt market-oriented policies with probability 87% on average (standard deviation 10%).
- As the coefficients γ are common to all countries, we use a flat prior for γ .

4.2. The likelihood function. In order to derive the posterior distribution of the unknown coefficients, we update these priors with the likelihood information. Notice that we have not postulated a *true* data generating process for GDP growth. We can do so under the assumption that GDP growth depends only on actual policies and is not directly affected by the policymakers' beliefs that led to those policies. This is a natural assumption which greatly simplifies and robustifies the inference about our parameters of

¹² This can be seen by combining (2.5) and (2.6) and noticing that we cannot distinguish between market-oriented and state interventionist countries in the pre 1950 data.

¹³ There is a huge outlier in the distribution of the average growth rates across countries. Therefore, this variance is estimated with a robust method (squared average distance from the median of the 16 and 84th percentiles).

interest (α), because it allows us to be completely agnostic about the details of the way policy decisions affect growth outcomes.

If we denote by D^s the available data up to a generic time s , the likelihood function can then be written as a product of conditional densities:

$$\mathcal{L}(D^T|\alpha) \propto \prod_{i=2}^N \left[\mathcal{L}(\theta_{i,1}|\alpha) \cdot \prod_{t=2}^T \mathcal{L}(\theta_{i,t}|D^{t-1}, \alpha) \right].$$

In turn, the conditional density $\mathcal{L}(\theta_{i,t}|D^{t-1}, \alpha)$ can be written as

$$\mathcal{L}(\theta_{i,t}|D^{t-1}, \alpha) = \Phi \left(\frac{\hat{\beta}_{i,t-1}^M - \hat{\beta}_{i,t-1}^S}{\sigma_{i,k}} \right)^{1(\theta_{i,t}=1)} \cdot \left(1 - \Phi \left(\frac{\hat{\beta}_{i,t-1}^M - \hat{\beta}_{i,t-1}^S}{\sigma_{i,k}} \right) \right)^{1-1(\theta_{i,t}=1)},$$

where $\Phi(\cdot)$ denotes the cdf of a standard Gaussian density. These results are derived in appendix B.

5. RESULTS

This section presents various measures of fit and the estimation results for our baseline specification of the model. In this specification, the weight that country i assigns to the data of country j ($w_{j|i}$) is a function of two variables: d_{ji} , physical distance (in thousands of Km) between the capital of country j and country i , and ℓ_{ji} , a dummy variable equal to 1 if countries i and j have the same official language. In other words, $z_{ji} = [d_{ji}, \ell_{ji}]$ in expression (2.11).

Extensions of the baseline model and robustness checks are presented in section 7.

5.1. The model's fit. We begin by assessing how well our model fits the data. Using different criteria, we will argue that the model explains quite well the observed dynamics of market-oriented policies over time and across countries.

First of all, the model correctly predicts 93% of the policy choices that we observe in the data. Moreover, it accounts for 13%, 30%, 46% and 60% of the switches in policies, within $\pm 0, 1, 2$ and 3 -year windows, respectively. Explaining changes in policies is a very challenging test, because we only observe 101 policy switches in our dataset (2% of the total policies). We are therefore quite satisfied with the performance of the model.

Figure 3 reports the actual and predicted fractions of market-oriented economies present in our sample. The model-predicted series corresponds to the sequence of one-step-ahead predictions of the model, absent any political cost shock. For example, the value of the

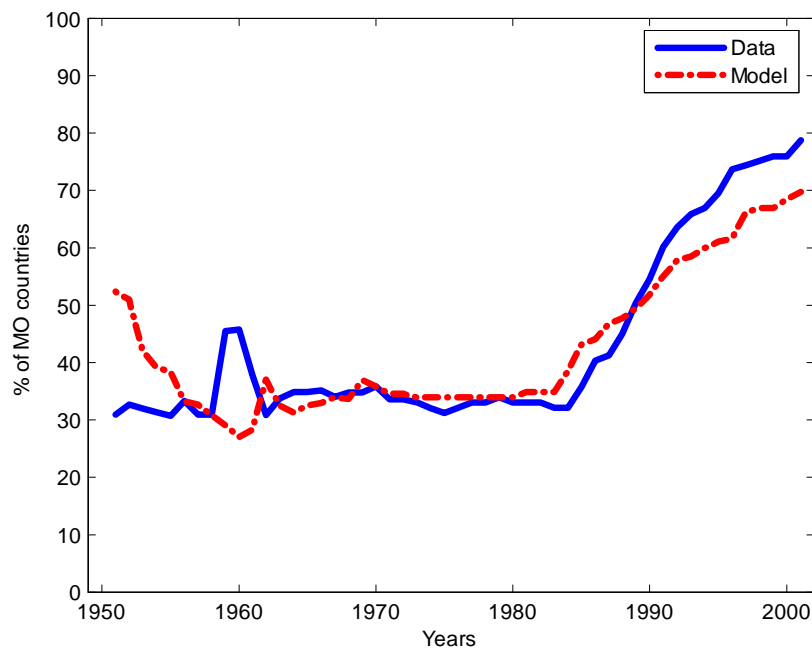


FIGURE 3. Actual and model-predicted fraction of market-oriented countries

red dashed line in 1990 represents what the model predicts for the world share of market-oriented countries, given the information on GDP growth and policies up to 1989, and assuming that in 1990 the shock to political costs is zero for each country. The figure makes clear that the model captures fairly well the high fraction of state-interventionist economies in the first part of the sample, and the run-up towards market orientation starting in the mid 1980s.

The success of the model along these dimensions is due to two main features of the data. First, the ability of the model to match the initial low share of market-oriented economies follows from the information in the pre-1950 data, which led to a calibration of initial priors consistent with most countries believing that the state outperforms the market. The first ten years of our estimation sample are also consistent with this idea. Second, the wave of liberalizations started in the mid-1980s are partially explained by the low growth observed in the 1970s and early 1980s. Through the lens of the model, past growth experiences are important determinants of beliefs and policy dynamics.

Figure 4 provides a more disaggregated picture of the fit of the model, by comparing the actual and predicted series for eight regions of the world economy. Notice that the

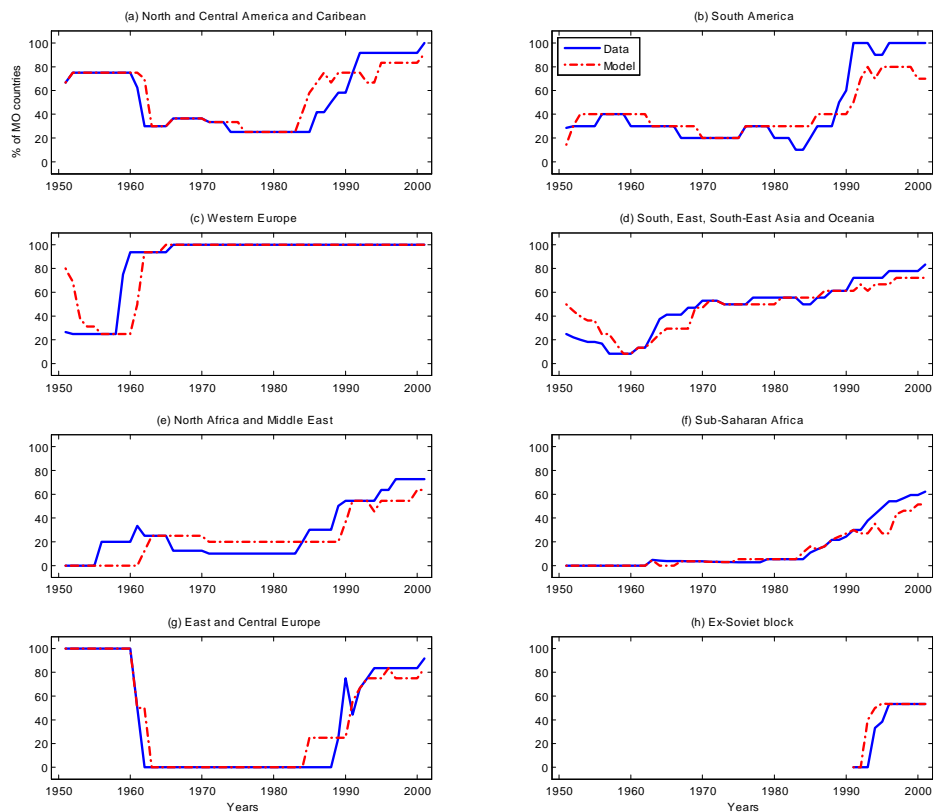


FIGURE 4. Actual and model-predicted fraction of market-oriented countries across regions

model captures well episodes of policy reversals (e.g. Central and South America) and the heterogeneous timing of liberalizations (e.g. the earlier liberalizations in Asia and the more protracted reforms in Africa).

It is also informative to analyze the cases where the model has some difficulties in fitting the data. For example, the early experience of Western Europe is difficult to rationalize based on a simple learning story. This is presumably due to the fact that the model abstracts from the postwar geopolitical forces that led to the integration of Western European countries and the creation of the European Community. Similarly, the fact that we slightly underpredict the number of liberalizations in the final part of the sample suggests a potential role for complementarity in policies, a force absent in our model.

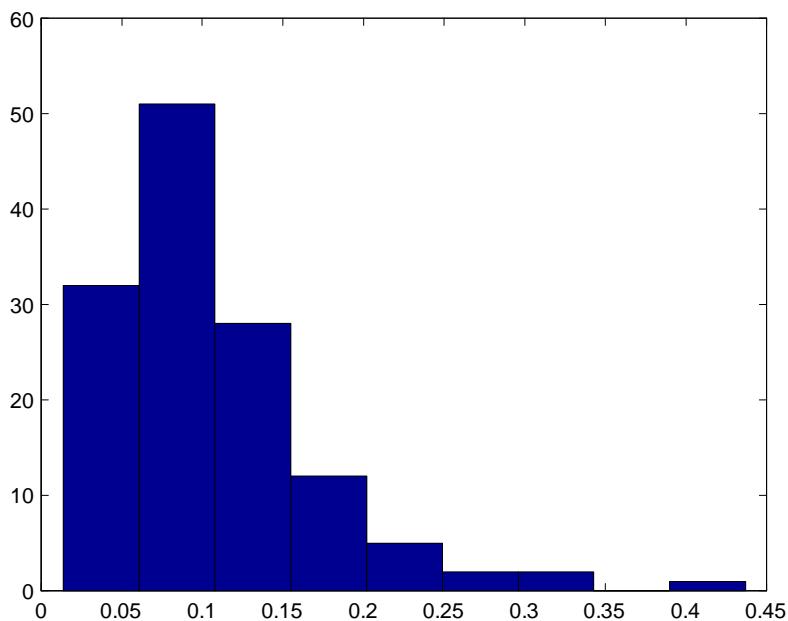


FIGURE 5. Histogram of $\left\{ \frac{\sigma_{i,k}^*}{\sqrt{\text{Var}(y_{i,t})}} \right\}_{i=1}^N$, i.e. the posterior modes of the standard deviations of political cost relative to the standard deviation of GDP growth in each country.

Another way to evaluate the fit of the model is to look at the magnitude of the variability of political costs needed to match the data. Figure 5 presents an histogram of $\{\sigma_{i,k}^*/\sqrt{\text{Var}(y_{i,t})}\}_{i=1}^N$, i.e. the ratio between the posterior mode of the standard deviation of the political cost ($\sigma_{i,k}^*$) and the standard deviation of the growth rate of GDP in each country. Figure 5 makes the point that in essentially every case the variability of the exogenous political cost necessary to explain the dynamics of policies is substantially lower (on average 10 times lower) than the typical variability of GDP growth in the same country.

5.2. Estimation results: the weighting function. Table 2 reports the estimates of the vector of coefficients γ that determines the effect of physical distance and language on the weight that country i puts on information from country j :

	Posterior mode
	(Posterior std)
d_{ij}	-0.4167 (0.015)
ℓ_{ij}	0.3163 (0.051)

TABLE 2: *Estimates of the coefficients of the weighting function in the baseline model.*

Our estimates indicate that countries put significantly more weight on data from countries nearby. Figure 6 plots the weight that country i puts on country j ($w_{j|i}$) as a function of distance (d_{ji}). For instance, everything else equal, the weight that a country puts on data from another country that is 2000Km away is almost 60% lower than the weight on its own data. The other conclusion that can be drawn from table 2 is that cultural distance matters: speaking the same language increases the weight that countries assign to each other by approximately 35%.

With this information in hand we ask: *Is learning globalized across the countries of the world?* The answer is *no*. Learning appears to be quite localized instead. Figure 7 plots the histogram of $\frac{\{\sum_{j \neq i} w_{j|i}\}_{i=1}^N}{w_{i|i}}$, i.e. the total weight assigned by each country to the rest of the world, relative to the weight assigned to their own data. Recall from equation (2.11) that this relative weight is a function of the distance measures and the relative volatility of growth rates across countries.

There are essentially no countries that weight equally the data from the rest of world. However, learning is not completely isolated either, as figure 7 shows that the weight put on data from other countries is substantial. Indeed, with the exception of two countries, all countries end up putting more weight on the information from the rest of the world than on their own data.

We conclude this subsection by comparing the fit of our baseline model to two alternative specifications. First, we assume that countries only learn from their own past data (\mathbf{M}_{own}). Second, we assume that countries learn globally, i.e., they do not discount information

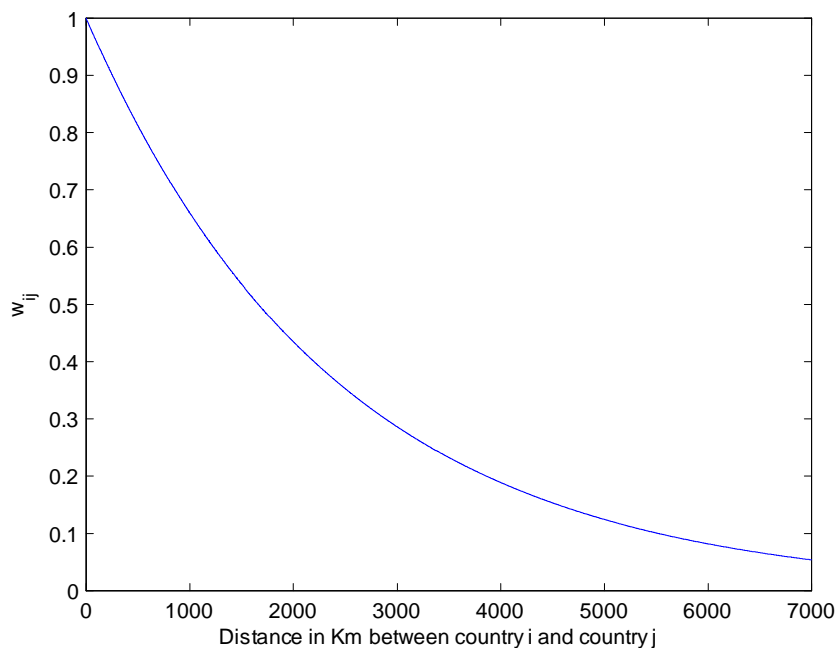


FIGURE 6. Weights on data from other countries (relative to own country) as a function of distance in Km.

on the basis of geographic and cultural distance (\mathbf{M}_{all}). Table 3 compares our baseline specification to these two alternative models. There are two things to notice: First, our model dominates these alternatives, especially in terms of capturing the policy switches in the data.¹⁴ Second, between the alternatives, the fit of the autarkic-learning model is substantially better than the global-learning extreme. This last finding also shows that a model that abstracts from a rich geographical learning structure is unable to fit the data, beside the complex structure of prior beliefs.

¹⁴ Compared to the alternatives, our model has two additional parameters. Nevertheless, both a likelihood ratio test or the Bayesian Information Criterion (BIC) would easily reject the restricted models. Of course, the formal way of performing a model comparison exercise in a Bayesian framework would be based on the comparison of the marginal data densities. However, the computation of the marginal data density is computationally very costly for these models.

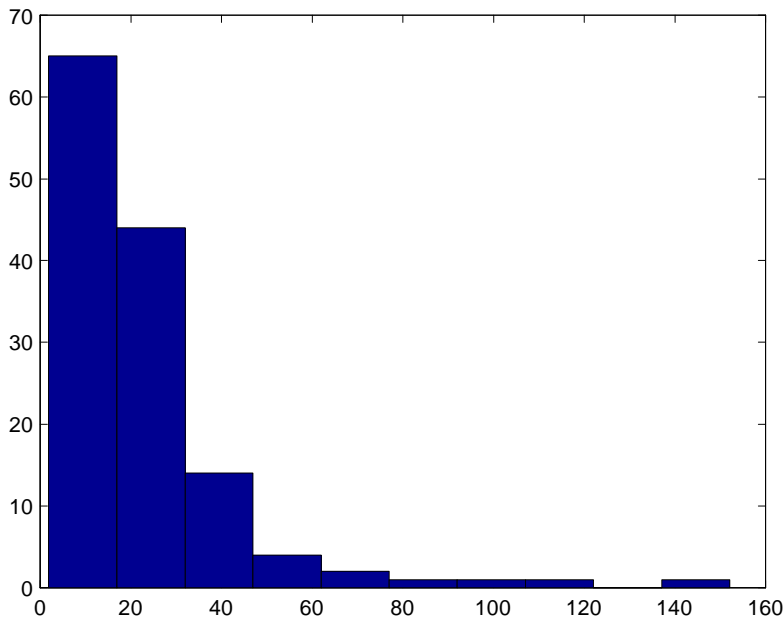


FIGURE 7. Histogram of total weight put on information from other countries relative to own country.

	Baseline model	\mathbf{M}_{own}	\mathbf{M}_{all}
log-likelihood	-865.59	-1,241.28	-1,928.18
share of correct predictions	93.0%	86.8%	69.4%
policy switches (± 0 -y window)	12.9%	8.9%	3.0%
policy switches (± 1 -y window)	29.7%	20.8%	9.9%
policy switches (± 2 -y window)	45.5%	28.7%	15.8%
policy switches (± 3 -y window)	60.4%	35.6%	30.7%

TABLE 3: *Measures of fit for the baseline model, a model in which countries learn only from their own past data (\mathbf{M}_{own}) and a model in which countries do not discount information on the basis of distance (\mathbf{M}_{all}).*

In sum, the weighting function plays an important role in our analysis. In the next sections we will analyze its contribution to the evolution of policymakers' beliefs about the effect of state intervention and market orientation on growth. Moreover, in section 7,

we will check the robustness of our results to alternative specifications of this weighting function.

5.3. Estimation results: evolution of beliefs. Figure 8 presents a summary of the estimated evolution of beliefs over time and across countries. Figure 8a plots a histogram of $\left\{ \hat{\beta}_{i,0}^{M*} - \hat{\beta}_{i,0}^{G*} \right\}_{i=1}^N$, i.e. the difference between the posterior mode of the mean of policymakers' prior beliefs about the effect of market-oriented and state-interventionist policies on growth. In estimating the model, we assume that the beliefs of countries entering the sample after 1950 reflect their 1950 priors and subsequent data. Notice that initial beliefs about the success of market economies were quite negative and characterized by considerable dispersion across countries. According to our estimates, in 1950 about 65% of the countries believed that market-oriented policies were inferior to state intervention in terms of economic growth.

Figure 8b shows that, with the information accumulated over fifty years, the beliefs of countries have shifted considerably. By 2001, the histogram of implied $\left\{ \hat{\beta}_{i,T}^{M*} - \hat{\beta}_{i,T}^{S*} \right\}_{i=1}^N$ has moved to the right, with a perceptible majority of countries believing that market-oriented policies are growth enhancing. However, quite interestingly, the dispersion of beliefs across countries has declined, but certainly not disappeared.

Figure 8c provides a time series perspective on the evolution of beliefs, by plotting expected growth under the two policies in the median country. As anticipated, the median country started off with beliefs biased towards state intervention and has slowly shifted towards favoring market-oriented regimes.

Figure 9 summarizes the precision of these beliefs across countries and over time. Figure 9a plots a histogram of $\left\{ \sigma_i \cdot \nu_i^* \right\}_{i=1}^N$, i.e. the posterior mode of the standard deviation of policymakers' prior beliefs about the growth effect of market-oriented and state-interventionist policies. For most countries, the initial uncertainty about the effect of these policies is substantial. By 2001, countries have sharpened their views considerably, although significant doubts remain (figure 9b). Finally, figure 9c plots the evolution of this uncertainty over time. At the end of the sample, the median country attaches a standard deviation of approximately 0.4% to its estimates of average growth under state intervention and market orientation.

We now have enough information to ask the important question: *Has the world reached a point in which the vast majority of policymakers are convinced that a market-oriented*

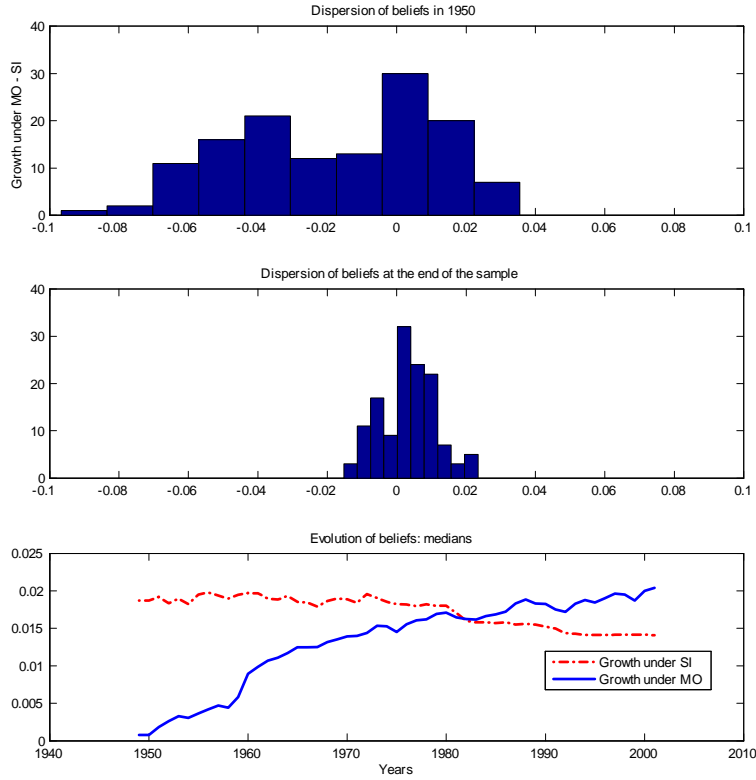


FIGURE 8. Estimates of policymakers' beliefs: (a) histogram of the posterior mode of the difference of expected growth under market orientation and state intervention in 1950; (b) histogram of the posterior mode of the difference of expected growth under market orientation and state intervention in 2001; (c) evolution of the posterior mode of expected growth under market orientation and state intervention for the median country.

regime is beneficial for economic growth? Figure 8 and 9 tell us that this is certainly not the case. There still exists a considerable amount of negative views about the merits of market economies. Moreover, many of those with a favorable view are either quite uncertain or believe that the gains are quantitatively minimal. More specifically, in our dataset only about 35 countries estimate average growth under market orientation to be statistically significantly (at the 95% confidence level) higher than average growth under state intervention, in 2001. As one would expect, the US, Western Europe and countries like Honk Kong or Singapore belong to this small group.

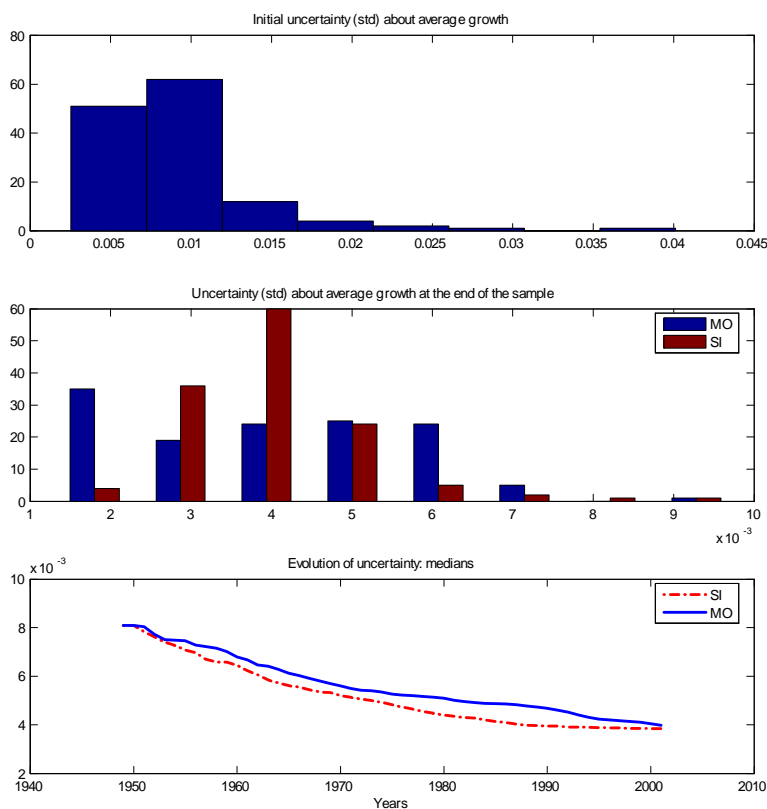


FIGURE 9. Estimates of policymakers' beliefs: (a) histogram of the posterior mode of initial uncertainty about growth under market orientation and state intervention in 1950; (b) histogram of the posterior mode of initial uncertainty about growth under market orientation and state intervention in 2001; (c) evolution of the posterior mode of initial uncertainty about growth under market orientation and state intervention for the median country.

6. COUNTERFACTUALS

In this section we address two questions:

- (1) *Do spillovers of information matter for the diffusion of market-oriented policies?*
- (2) *Would many countries revert to state intervention if the world was hit by another Great Depression?*

We will argue that the answer to both of these questions is *yes*.

6.1. Autarkic and global learning. Our multicountry learning model allows for a rather flexible specification of the weighting function (2.11), according to which policymakers can discount information from other countries based on measures of physical and cultural distance. In this subsection, we perform a counterfactual exercise in order to understand how beliefs and policies would have evolved under some alternative scenarios. In particular, we analyze two restricted versions of the model: a model in which countries learn only from their own past experience (\mathbf{M}_{own}), and another one in which countries do not discount information from places that are far away physically or culturally (\mathbf{M}_{all}). Model \mathbf{M}_{own} corresponds to driving the elements of γ to $-\infty$, while model \mathbf{M}_{all} is obtained by setting $\gamma = 0$. The other model coefficients are kept fixed at the posterior mode of the baseline estimation.

Figure 10 plots the evolution of median beliefs on the expected growth rates under state intervention and market-oriented policies in the \mathbf{M}_{own} model. Compared to the evolution of beliefs estimated in our baseline model (see figure 8c), countries learn very slowly since they do not gain any knowledge from the experience of other countries. The absence of cross-country spillovers would have slowed down the diffusion of market-oriented policies to the point that the world would still be, to a large extent, a collection of state-interventionist economies.

On the contrary, if countries had weighted foreign data as much as their own, learning would have been very fast. This point is nicely illustrated in figure 11, which plots the evolution of median beliefs about the effect of state intervention and market-oriented policies in the \mathbf{M}_{all} model. With all the information available in the world uniformly used across countries, the sample evidence would have quickly overtaken the initial unfavorable views on markets, and market-based regimes would have prevailed as early as in the 1960s. We conclude that the localization of learning might have severely slowed down the global diffusion of market-oriented policies.

6.2. Policy reversals: another Great Depression. We now consider how large negative world-wide shocks to GDP growth impact policy decisions. To accomplish this task we conduct another counterfactual simulation exercise. Differently from the previous subsection, we keep the estimated weighting scheme, but generate artificial paths of GDP growth, in order to induce a global depression.

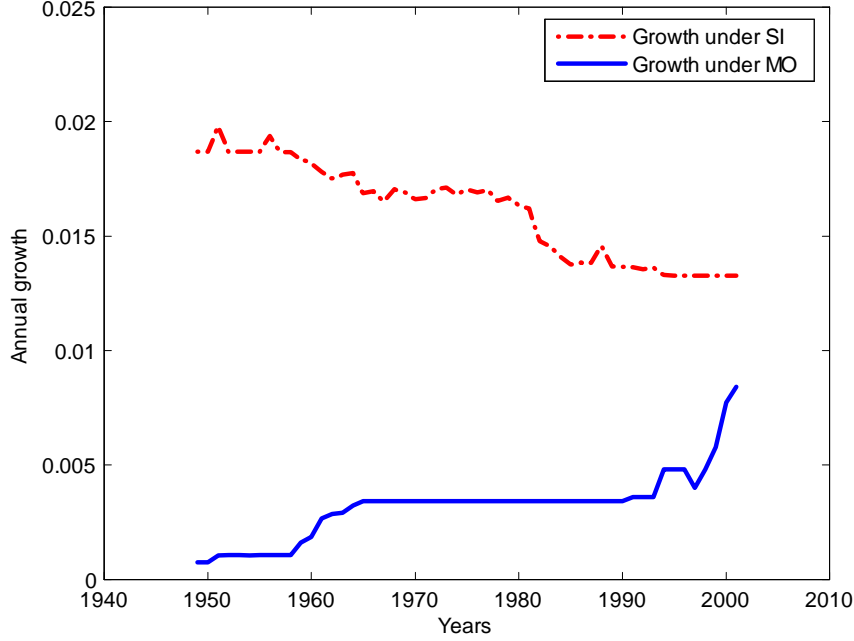


FIGURE 10. Evolution of beliefs (expected growth under market orientation and state intervention in the median country) under the counterfactual scenario in which countries learn only from their own past data.

More specifically, these simulations are constructed as follows. Suppose that time τ is the starting point of the counterfactuals. Based on time- τ beliefs and the realization of the exogenous political cost at time $\tau + 1$, policymakers choose the value of the policy variable; this policy choice contributes to the realization of GDP growth in period $\tau + 1$; a new vintage of data is now available and policymakers form time- $\tau + 1$ beliefs by updating their priors with the new information; and so on.

Contrary to the rest of the paper, in this subsection we need to postulate a true data generating process for GDP growth. To keep things simple, we simulate the realization of GDP growth in every period using the following stochastic process:

$$(6.1) \quad \begin{aligned} y_{i,s} &= \beta^S (1 - \theta_{i,s}) + \beta^M \theta_{i,s} + f_{i,s} + e_{i,s}, \quad s = \tau + 1, \dots, \tau + H \\ e_{i,s} &\sim N(0, \varpi_i^2), \end{aligned}$$

where H denotes the length of the simulation and $f_{i,s}$ is an exogenous forcing variable. We obtain specific values for β^S , β^M and $\{\varpi_i^2\}_{i=1}^N$ by running a simple panel regression

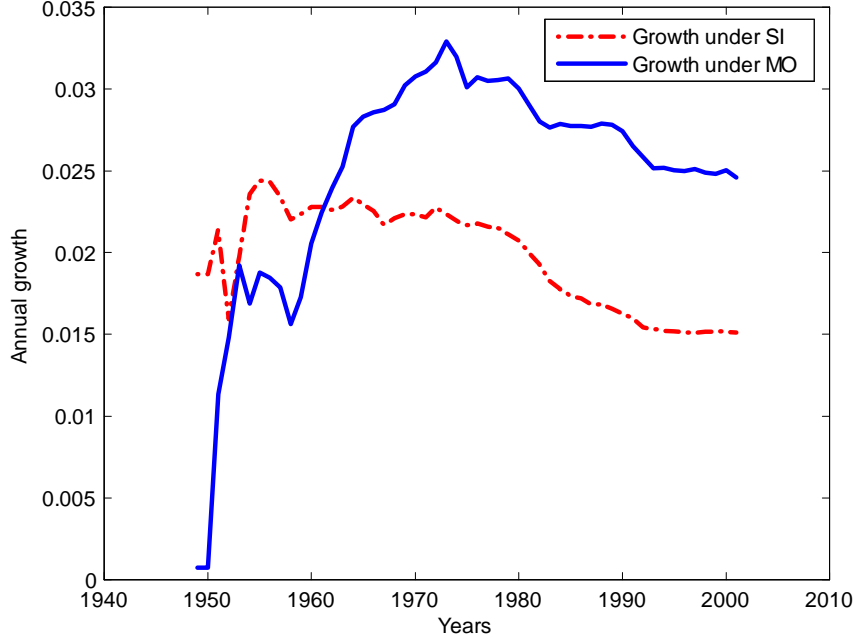


FIGURE 11. Evolution of beliefs (expected growth under market orientation and state intervention in the median country) under the counterfactual scenario in which countries do not discount information based on distance.

of GDP growth on the SW indicator, using the entire sample. This procedure delivers $\beta^S = 0.0107$, $\beta^M = 0.0259$ and $\{\varpi_i^2\}_{i=1}^N$ equal to the variances of the residuals. We calibrate $\{f_{i,s}\}_{s=1}^H$ to match the size of the Great Depression in the 1930s. Using the Maddison dataset we compute

$$\{f_{i,s}\}_{s=1}^H = [-0.0047, -0.0517, -0.0867, -0.0687, 0.0023, 0, \dots, 0], \quad i = 1, \dots, N.$$

which corresponds to the average deviation across countries of the growth rate in 1929-1933 relative to the average growth between 1919 and 1928. Finally, the other model parameters (initial beliefs, volatility of the political costs and coefficients of the weighting function) are set to the estimated posterior mode.

Before turning to the description of the results, we want to stress that the results of this subsection are conditional on (6.1), i.e. the particular data generating process that we have chosen for GDP growth. The process in (6.1) has two main advantages. First, it is transparent and easy to cast into our model. Second, it closely resembles the growth regressions that we assume our policymakers estimate to update their beliefs. Therefore,

if the forcing variable is set to $f_{i,s} = 0$, our learning model eventually converges to a self-confirming equilibrium in which everybody knows the truth. On the other hand, in using (6.1), we are making the strong assumption that the residuals of the GDP growth equation ($e_{i,s}$) are uncorrelated with the policy variable ($\theta_{i,s}$). While this is the case in our model,¹⁵ it is possible to imagine situations in which this condition fails. The most obvious of these situations is, for instance, a case in which the shocks to GDP growth ($e_{i,s}$) are not independent from the political cost ($K_{i,s}$).

Figure 12 plots the impulse response of the fraction of market-oriented countries to the global shock. The two lines correspond to different starting points for the experiment (different τ). For instance, had such global depression hit the world in 1961, when the world was mostly characterized by state intervention, it would have spawn the diffusion of market-oriented policies. A deep recession would have cast serious doubts on the growth perspectives of economies heavily based on government controls. The opposite would happen with a global recession in 2002. Given that by 2001 most countries are market friendly, policymakers around the world would receive very precise negative signals about the growth prospects of market economies. As a consequence, they would become pessimistic about the success of market-oriented policies much faster than about the effects of government intervention.

In sum, the recession would have large and persistent effects. It would persuade about 13% of the countries to revert to state intervention and it would take almost forty years for this effect to wash out.

Of course things would be even worse if, for some reason, the recession had a relatively more severe effect on market economies. For instance, figure 13 shows that the fraction of countries reverting to interventionism after five years would be 20% if the recession affected market-oriented economies twice as much as state-interventionist ones (keeping fixed the average effect).

Summing up, would we observe policy reversals towards state intervention if the world was hit by a severe recession? The answer of our estimated model is yes.

¹⁵ Policy decisions are predetermined in our model.

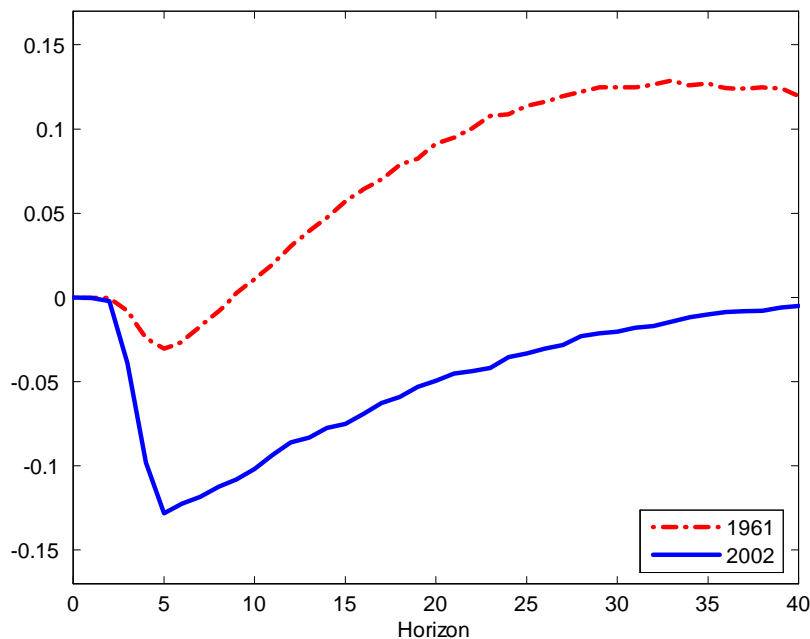


FIGURE 12. Fraction of countries becoming market oriented as a consequence of a counterfactual severe recession hitting at different possible points in time.

7. Sensitivity Analysis and Extensions

Our benchmark specification assumes that policymakers base their policy decisions on a relatively stylized model. In this section, however, we demonstrate that our results are robust to various alternative specifications and extensions of the baseline framework.

7.1. Alternative weighting schemes. We first examine the robustness of our results to changes in the specification of the weighting function. For example, our baseline model abstracts from the possibility that colonizers receive a disproportionately large or small weight from former colonies. In fact, anecdotal evidence suggests that former colonies might either follow closely the lead of a former metropolis or, alternatively, might be biased against its policies. To explore the importance of these considerations, we augment the weighting function with a dummy variable c_{ji} , which equals 1 if country j has been a colonizer of country i and 0 otherwise. The second column of table 4 reports the results of this specification. The weight depends negatively on having been a colonizer, but the estimated coefficient is far from being statistically different from zero. Moreover, the

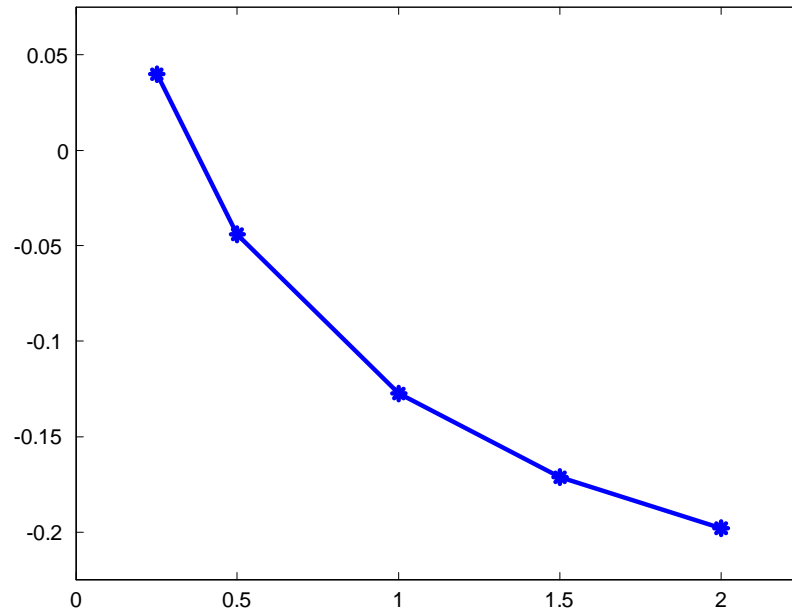


FIGURE 13. Fraction of countries becoming market oriented as a consequence of a severe world-wide recession in 2002, as a function of the differential impact of the recession on market-oriented and state-interventionist countries.

improvement in the value of the log-likelihood is negligible and model's ability to predict policy changes is reduced.

	baseline (\mathbf{M}_1)	\mathbf{M}_2	\mathbf{M}_3	\mathbf{M}_4	\mathbf{M}_5	\mathbf{M}_6
d_{ji}						
geographic	-0.417	-0.403	-0.079	-0.438	-0.462	-0.494
distance	(0.015)	(0.017)	(0.006)	(0.03)	(0.019)	(0.027)
ℓ_{ji}						
language	0.317	0.346	0.308	0.207	0.253	0.285
	(0.051)	(0.067)	(0.077)	(0.078)	(0.065)	(0.079)
c_{ji}		-5.727				
colony		(43.805)				
$\Delta_{ji,t}$						
log-difference			-1.748			
GDP per-capita			(0.279)			
$\Pi_{i,t}$				0.001		
Polity2				(0.000)		
log-likelihood	-865.59	-863.26	-824.03	-703.73 ¹⁶	-863.93	-971.62
% correct predictions	93	93	93.6	93.4	92.9	91.4
% correctly predicted						
switches (± 3 -y window)	60.4	57.4	55.5	65.6	61.4	52.5

TABLE 4: *Coefficient estimates in the baseline model and various alternative specifications (described in section 7).*

¹⁶ The dataset used for estimating this model has more missing values than the baseline dataset, due to the fact that we also need the Polity2 variable. As a consequence, the log-likelihood of this model is not comparable to the others.

We next examine the role of economic distance, as measured by differences in the level of economic development. In addition to geography and cultural proximity, policymakers might discount information from countries that are in a very different stage of the economic development process. In order to investigate this possibility, we include an additional variable $\Delta_{ji,t}$, defined as the absolute value of the log-difference in GDP per-capita between country i and country j at time t . The third column of table 4 shows a negative and statistically significant coefficient on $\Delta_{ji,t}$, suggesting that, indeed, policymakers tend to attach more weight to the experience of countries characterized by a similar level of development. For example, information coming from a country in which GDP per-capita is twice as much as in the home country receives 70% less weight. Also, notice that the inclusion of $\Delta_{ji,t}$ in the weighting function reduces the coefficient on geographical distance, This is not surprising, since development has an important geographical component and d_{ji} and $\Delta_{ji,t}$ are highly correlated. Interpreting a weighting structure that incorporates per-capita GDP differences is promising and deserves further investigation. This being said, it is important to point out that the quantitative implications of this specification of the model are almost the same as in the baseline.

To assess the economic differences of alternative models, figure 14 plots the evolution of the difference between median expected growth under market orientation and state intervention. All models (\mathbf{M}_4 , \mathbf{M}_5 and \mathbf{M}_6 are explained below) tell essentially the same story about the evolution of beliefs.

As a final comparison, figure 15 reports our counterfactual exercise of simulating the policy response to another Great Depression starting in 2002 (see section 6.2). All models have very similar quantitative implications.

7.2. Controlling for the political environment. Another important dimension of our model is given by the specification of the political cost. So far, we have assumed that $K_{i,t}$ has mean zero and is uncorrelated over time. We now examine the implications of allowing this political cost to depend on variables measuring a country’s political environment.

For this extension, we need a measure of political environments that is comparable across countries and over time. The “Polity IV” dataset is a multidimensional panel that measures various aspects of the governance of countries and is widely employed in research on political science and international relations. The main variable of the dataset is the Polity2 score, which aggregates democratic and autocratic qualities of institutions.

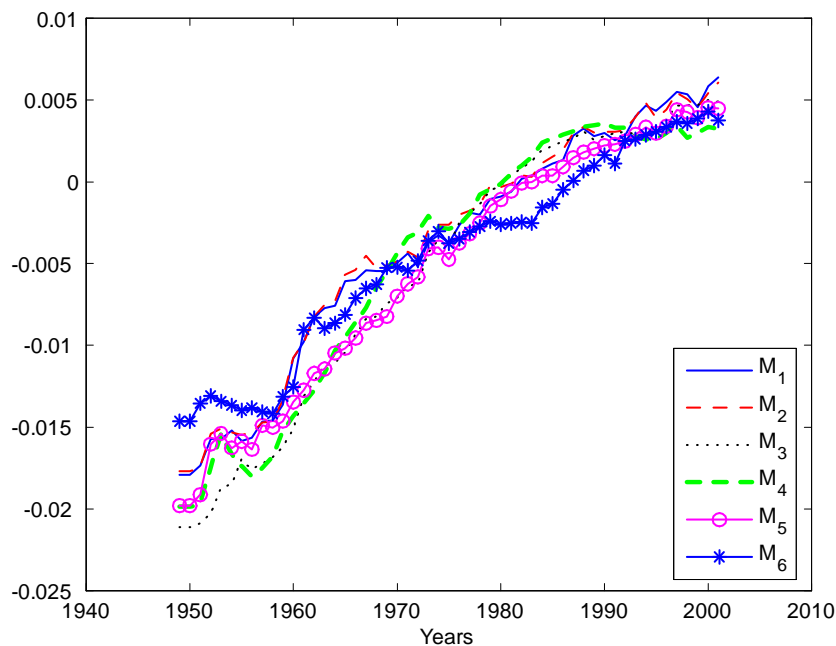


FIGURE 14. Evolution of beliefs (difference between expected growth under market orientation and state intervention in the median country) implied by baseline model and alternative specifications.

It classifies countries in a spectrum on a 21-point scale, ranging from -10 to $+10$ (see Marshall and Jaggers (2005)).

To explore the possibility that the political cost of market-oriented policies depends on the political environment, we replace (2.4) with the following specification for $K_{i,t}$:

$$K_{i,t} = \xi \Pi_{i,t} + k_{i,t}$$

$$k_{i,t} \stackrel{i.i.d.}{\sim} N(0, \sigma_{k,i}^2),$$

where $\Pi_{i,t}$ denotes the deviation of the Polity2 score for country i at time t from its pooled mean. Notice that this specification allows average political costs to differ across countries because of the cross-country differences in the average Polity2 score. Moreover, the political cost inherits some of the autocorrelation of $\Pi_{i,t}$, which is usually highly persistent.

The fourth column of table 4 reports the estimation results for this extension of the model. First, notice that this model is the best in terms of explaining policy switches.

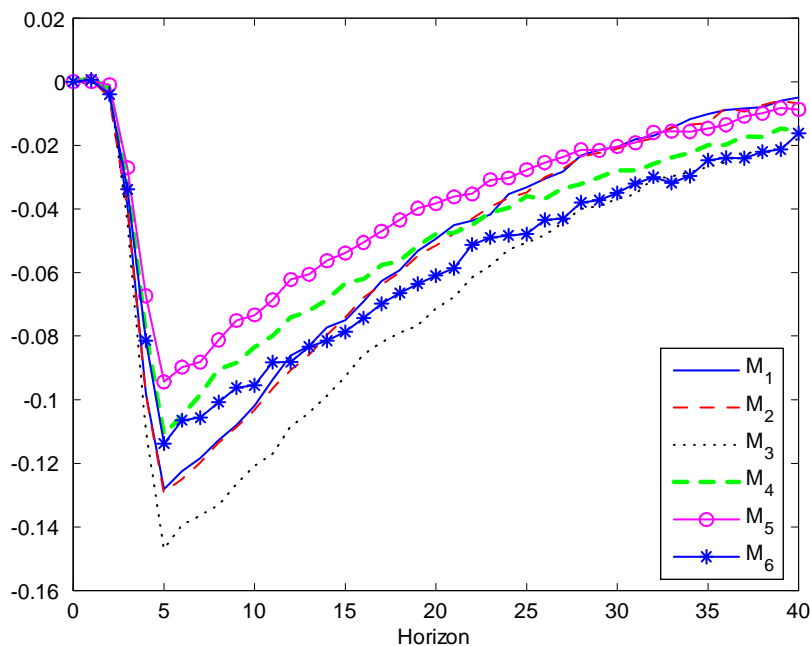


FIGURE 15. Fraction of countries becoming market oriented predicted by baseline model and alternative specifications, as a consequence of a counterfactual severe recession hitting in 2002.

Second, the coefficients of the weighting function are very similar to the baseline. Third, the coefficient ξ has a negative and significant sign, implying that more democratic countries have a lower political cost of being market-oriented. To evaluate the magnitude of ξ , figure 16 plots the model implied probability of being market oriented as a function of the Polity2 score. The figure is constructed for the case in which countries hold similar beliefs about the effects of market orientation and state intervention, and assuming that $\sigma_{k,i} = 0.53\%$.¹⁷ Everything else equal, a country with low Polity2 score is substantially less likely to be market oriented relative to more democratic countries.

These results are sensible and indicate that the interaction of political variables with the countries' learning mechanism that we introduce in this paper deserves further attention. In terms of quantitative implications, however, the model's implied evolution of beliefs and response in the aftermath of a global Great Depression (recall figure 15 and 14) are very similar to those of our baseline model.

¹⁷ This number corresponds to the cross-country mean of the estimated $\sigma_{i,k}$.

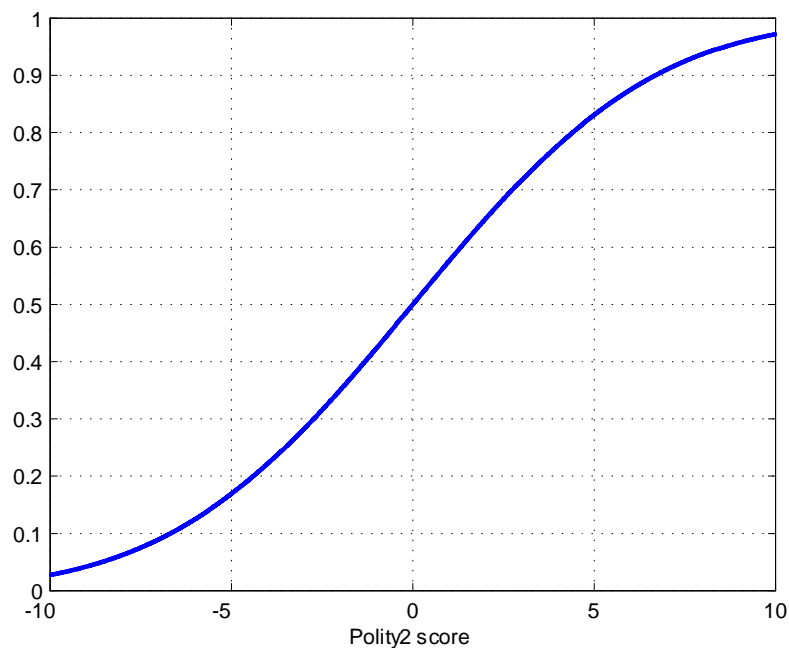


FIGURE 16. Probability of becoming market oriented as a function of the Polity2 score.

7.3. Spatially correlated growth shocks. In our benchmark model, policymakers treat growth shocks as uncorrelated across countries. If in reality these shocks exhibit substantial spacial correlation, this assumption might bias our estimates of the geographic localization of learning. Therefore, in this subsection, we analyze a more general model in which policymakers take into account the spatial correlation structure of growth shocks.

In order to do so, we use all available data to run a panel regression of GDP growth on the SW indicator and extract the residuals. We then estimate their covariance matrix by assuming that the correlation coefficient between country i and country j is non-negative and has the following form:

$$\text{corr}(\varepsilon_{i,t}, \varepsilon_{j,t}) = cr^{d_{ji}},$$

where $0 \leq c \leq 1$ and $0 \leq r \leq 1$ are two coefficients that we estimate by maximum likelihood, jointly with $\{\sigma_i^2\}_{i=1}^N$.¹⁸ Finally, we endow policymakers with the knowledge of this correlation structure.

Our estimates of c and r indicate that growth shocks are weakly spatially correlated. For instance, this correlation is equal to approximately 0.2 if two countries are 1000 km apart. The fifth column of table 4 shows that the estimates of the weighting function when we incorporate this correlation structure are very similar to the baseline model. The implications of this model for the behavior of policymakers' beliefs and the consequences of a counterfactual severe recession are also quite similar to the baseline (figure 14 and 15). In other words, empirically plausible levels of spacial correlation among growth shocks do not substantially affect our results.

7.4. Conditional convergence. In our final extension, we modify the specification of the empirical model that policymakers use to update their beliefs. More concretely, we replace (2.2) with

$$(7.1) \quad y_{i,s} = \beta_i^S (1 - \theta_{i,s}) + \beta_i^M \theta_{i,s} + \zeta (y_{i,s-1} - y_{s-1}^*) + \chi \iota_{i,s-1} + \varepsilon_{i,s},$$

This specification allows for a country's GDP growth rate to depend not only on policy choices, but also on the logarithm of its investment-to-GDP ratio (ι_i) and the gap between the country's GDP and the US (y^*). Everything else equal, and assuming that $\zeta < 0$, under this specification poorer countries grow faster, as in the convergence regressions of Barro and Sala-I-Martin (1995).

Ideally, we would allow policymakers to learn about all of the unknown coefficients of (7.1). In practice, however, this is computationally extremely costly. Therefore, we set $\zeta = -0.0025$ and $\chi = 0.0149$ and assume that policymakers know these values (that we have chosen by running a panel regression using the entire set of available data).¹⁹

The last column of table 4 reports the estimates of the coefficients of the weighting function for this specification of the model. Notice that these values are in line with our

¹⁸ We cannot estimate this correlations by simply using the sample covariances because of the unbalanced nature of our panel. Therefore, the estimate of the covariance matrix obtained by combining pairwise sample covariances need not be positive definite. Alternatively, the covariance matrix estimated using only overlapping observations for all countries would be very imprecise, since **for some country pairs** we only have five of these observations.

¹⁹ We obtain a lower value for ζ than Barro and Sala-I-Martin (1995), partly because we use a larger sample of countries and years, and because we exploit variation of growth rates across countries and over time.

baseline estimates. The overall fit of the model is slightly worse than the benchmark but the evolution of beliefs (figure 14) and the implied response to a Great Depression (figure 15) are very similar, reinforcing our conclusion that that our main results are quite robust.

8. CONCLUDING REMARKS

We have explored the ability of a learning model to explain the observed transition of countries between regimes of state intervention and market orientation. In our model, the crucial determinants of policy choices are policymakers' beliefs about the relative merits of the markets versus the state. These beliefs, in turn, are influenced by past experience.

The estimated model fits well the data of 133 countries for the postwar period. Our results indicate that policymakers update their beliefs relatively slowly because they attach little weight to information coming from countries that are geographically, culturally or economically distant from their home country. This is important to explain the slow adoption of market-oriented policies in the postwar period. Moreover, our findings suggest that by the end of the sample many countries remain agnostic about the growth prospects of market economies. Indeed, according to the model, a global GDP growth shock of the size of the Great Depression would induce almost 15% of countries in the world to revert to state intervention.

In this paper, we have focused on the evolution of beliefs as the driving mechanism of policy changes and have abstracted from other forces that are undoubtedly important. For example, we have not considered the higher incentives that a country might have to be market oriented when other countries are also market oriented, as well as all other reasons that may lead countries to herd behind the policy choices of other countries. It would be interesting to allow for these policy complementarities and strategic interaction among policymakers of different countries. Similarly, it seems important to enrich our model with explicit political economy dimensions and examine how they interact with the mechanism of beliefs formation that we have presented in this paper. These are our priorities for future research.

APPENDIX A. UPDATING BELIEFS

This appendix shows how to use the assumptions in section 2.3 to derive the updating formulas for $\hat{\beta}_{i,t} \equiv E_{i,t}([\beta_i^S, \beta_i^M]')$ resulting from the application of Bayes rule. The calculations are conducted from the perspective of policymakers of country i .

First, define the vector of regressors $x_{j,t} \equiv [1 - \theta_{j,t}, \theta_{j,t}]'$ and coefficients $\beta_{j|i,t} \equiv [\beta_{j|i,t}^S, \beta_{j|i,t}^M]'$. Rewrite equation (2.7) as

$$(A.1) \quad y_{j,t} = x'_{j,t} \beta_{j|i,t} + \varepsilon_{j|i,t}, \quad j = 1, \dots, N,$$

We can now substitute equations (2.8) and (2.9) into (A.1). We obtain

$$(A.2) \quad y_{j,t} = x'_{j,t} \beta_i + \tilde{\varepsilon}_{j|i,t}, \quad j = 1, \dots, N,$$

where

$$\text{var}(\tilde{\varepsilon}_{j|i,t}) = \sigma_j^2 (1 + q_{j|i}), \quad j = 1, \dots, N.$$

Finally, rewrite (A.2) as

$$(A.3) \quad y_{j,t} = x'_{j,t} \beta_i + \frac{\sigma_i}{w_{j|i}} \varepsilon_{j|i,t}^*, \quad j = 1, \dots, N$$

where $\varepsilon_{j|i,t}^* \equiv \frac{w_{ij}}{\sigma_i} \tilde{\varepsilon}_{j|i,t}$ and $w_{j|i} \equiv \frac{\sigma_i}{\sigma_j} \frac{1}{\sqrt{1+q_{j|i}}}$. As $w_{i|i} = 1$, equation (A.3) holds for any j . Moreover, notice that $\text{var}(\varepsilon_{j|i,t}^*) = 1$. The estimation of equation (A.3) corresponds to a weighted least square estimation problem. Given our assumption that policymakers know the variance of the shocks to growth, $\{\sigma_j\}_{j=1}^N$, the weights $\{w_{j|i}\}_{j=1}^N$ are known to them as well. In practice, we run a panel regression of GDP growth on the SW indicator, compute the variance of the residuals and endow policymakers with the knowledge of these coefficients.

It is then easy to show that the optimal updating formulas for the expectation of policymakers' beliefs in country i are:

$$\begin{aligned} P_{i,t} &= P_{i,t-1} + X'_t W^2 X_t \\ \hat{\beta}_{i,t} &= P_{i,t}^{-1} \left(P_{i,t-1} \hat{\beta}_{i,t-1} + X'_t W^2 y_t \right), \end{aligned}$$

where $y_t \equiv [y_{1,t}, \dots, y_{N,t}]'$, $X_t = [x_{1,t}, \dots, x_{N,t}]'$ and $W = \text{diag}([w_{1|i}, \dots, w_{N|i}])$. The recursion is initialized at $\hat{\beta}_{i,0}$ and $P_{i,0}$ which denote the prior mean and precision matrix respectively.

APPENDIX B. THE LIKELIHOOD FUNCTION

Let $\alpha \equiv \left[\left\{ \hat{\beta}_{j,0}^S \right\}_{j=0}^N, \left\{ \hat{\beta}_{j,0}^M \right\}_{j=0}^N, \left\{ \nu_j \right\}_{j=0}^N, \left\{ \rho_j \right\}_{j=0}^N, \left\{ \sigma_{j,k} \right\}_{j=1}^N, \gamma \right]$ and let ψ denote the set of coefficients that parameterize the true data generating process for GDP growth, which we have left unrestricted. The likelihood function can be written as a product of conditional densities:

$$(B.1) \quad \mathcal{L}(D^T | \alpha, \psi) = \mathcal{L}(D_1 | \alpha, \psi) \prod_{t=2}^T \mathcal{L}(D_t | D^{t-1}, \alpha, \psi),$$

with a slight abuse of notation in which \mathcal{L} is used generically to denote an arbitrary density function. Under the assumption that the distribution of the vector $y_t \equiv [y_{1,t}, \dots, y_{N,t}]'$ depends on α only through the vector $\theta_t \equiv [\theta_{1,t}, \dots, \theta_{N,t}]'$, it follows that

$$(B.2) \quad \mathcal{L}(D_t | D^{t-1}, \alpha, \psi) = \mathcal{L}(y_t | \theta_t, D^{t-1}, \psi) \mathcal{L}(\theta_t | D^{t-1}, \alpha)$$

$$(B.3) \quad = \mathcal{L}(y_t | \theta_t, D^{t-1}, \psi) \prod_{i=1}^N \mathcal{L}(\theta_{i,t} | D^{t-1}, \alpha)$$

This assumption is natural in our framework. It implies that the realization of GDP growth is affected by past and current policy decisions, but does not directly depend on the beliefs that have led to those policy choices. This assumption allows us to evaluate the likelihood function of α without taking a stand on ψ and the true data generating process of GDP growth. It is important to note that, even in those cases in which this assumption is undesirable, our methodology can still be interpreted as a limited-information estimation strategy. Combining (B.1) and (B.3), we obtain the following result:

$$\mathcal{L}(D^T | \alpha) \propto \prod_{i=2}^N \left[\mathcal{L}(\theta_{i,1} | \alpha) \cdot \prod_{t=2}^T \mathcal{L}(\theta_{i,t} | D^{t-1}, \alpha) \right].$$

Since the policy decision is given by

$$\theta_{i,t} = 1 \{ E_i(\beta_i^M | D^{t-1}) - E_i(\beta_i^S | D^{t-1}) > K_{i,t} \}, \quad i = 1, \dots, N,$$

it follows that

$$\begin{aligned} \Pr(\theta_{i,t} = 1 | D^{t-1}, \alpha) &= \Pr(K_{i,t} < E_i(\beta_i^M | D^{t-1}) - E_i(\beta_i^S | D^{t-1})) = \\ &= \Phi \left(\frac{\hat{\beta}_{i,t-1}^M - \hat{\beta}_{i,t-1}^S}{\sigma_{i,k}} \right). \end{aligned}$$

This implies

$$\mathcal{L}(\theta_{i,t}|D^{t-1}, \alpha) = \Phi\left(\frac{\hat{\beta}_{i,t-1}^M - \hat{\beta}_{i,t-1}^S}{\sigma_{i,k}}\right)^{1(\theta_{i,t}=1)} \cdot \left(1 - \Phi\left(\frac{\hat{\beta}_{i,t-1}^M - \hat{\beta}_{i,t-1}^S}{\sigma_{i,k}}\right)\right)^{1-1(\theta_{i,t}=1)}.$$

REFERENCES

- ACEMOGLU, D., S. JOHNSON, AND J. A. ROBINSON (2005): “Institutions as the Fundamental Cause of Long-Run Growth,” in *Handbook of Economic Growth*, ed. by P. Aghion, and S. Durlauf, chap. 6, pp. 385–472. North Holland.
- BANBURA, M., D. GIANNONE, AND L. REICHLIN (2007): “Bayesian VARs with Large Panels,” CEPR Working Paper N. 6326.
- BANERJEE, A. V. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107(3), 797–817.
- BARRO, R. J., AND X. SALA-I-MARTIN (1995): *Economic Growth*. McGraw Hill, New York.
- BERRY, D. A., AND B. FRISTEDT (1985): *Bandit Problems*. Chapman and Hall, New York.
- BESLEY, T. J., AND A. CASE (1994): “Diffusion as a Learning Process: Evidence from HYV Cotton,” Woodrow Wilson School, Research Program in Development Studies Discussion Paper No. 174.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 100(5), 992–1026.
- BLATTMAN, C., M. A. CLEMENS, AND J. G. WILLIAMSON (2002): “Who Protected and Why? Tariffs the World Around 1870-1938,” Harvard University.
- BLINDER, A. S. (1998): *Central Banking in Theory and Practice*. MIT Press, Cambridge, Massachusetts.
- BOLTON, P., AND C. HARRIS (1999): “Strategic Experimentation,” *Econometrica*, 67, 349–374.
- BROCK, W. A., AND S. N. DURLAUF (2001a): “Discrete Choice Models with Social Interaction,” *Review of Economic Studies*, 68(2), 235–260.
- (2001b): *Interaction-Based Models* vol. 5. North Holland, Amsterdam.
- CHAMBERLAIN, G. (1980): “Analysis of Covariance with Qualitative Data,” *Review of Economic Studies*, 47(1), 225–238.
- CLEMENS, M. A., AND J. G. WILLIAMSON (2004): “Why Did the Tariff-Growth Correlation Change After 1950?,” *Journal of Economic Growth*, 9 (1), 5–46.
- COGLEY, T., AND T. J. SARGENT (2005): “The conquest of US inflation: Learning and robustness to model uncertainty,” *Review of Economic Dynamics*, 8(2), 528–563, available at <http://ideas.repec.org/a/red/issued/v8y2005i2p528-563.html>.
- CONLEY, T., AND C. UDRY (2005): “Learning About a New Technology: Pineapple in Ghana,” Yale University.
- CRAN, W. (2002): “Commanding Heights : The Battle for the World Economy,” Videorecording.
- DOAN, T., R. LITTERMAN, AND C. A. SIMS (1984): “Forecasting and Conditional Projection Using Realistic Prior Distributions,” *Econometric Review*, 3, 1–100.

- DOLLAR, D. (1992): “Outward-Oriented Developing Economies Really Do Grow More Rapidly: Evidence from 95 LDCs, 1976-85,” *Economic Development and Cultural Change*, 40 (3), 523–544.
- EDWARDS, S. (1998): “Openness, Productivity and Growth: What Do We Really Know?,” *Economic Journal*, 108 (2), 383–398.
- EL-GAMAL, M. A., AND R. K. SUNDARAM (1993): “Bayesian Economist ... Bayesian Agents I: An Alternative Approach to Optimal Learning,” *Journal of Economic Dynamics and Control*, 17, 355–383.
- FISCHER, S. (1995): “Comment on J. Sachs and A. Warner, ‘Economic Reform and the Process of Global Integration’,” *Brookings Papers on Economic Activity*, 1995 (1), 100–105.
- FOSTER, A. D., AND M. R. ROSENZWEIG (1995): “Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture,” *The Journal of Political Economy*, 103(6), 1176–1209.
- FOTHERINGHAM, S., C. BRUNSDON, AND M. CHARLTON (2002): *Geographically Weighted Regression: The Analysis of Spatially Varying Relationships*. Wiley, Chichester.
- GELMAN, A., J. B. CARLIN, H. S. STERN, AND D. B. RUBIN (2004): *Bayesian Data Analysis*. Chapman and Hall, Boca Raton, Florida.
- GROSSMAN, G. M., AND E. HELPMAN (1995): “The Politics of Free-Trade Agreements,” *American Economic Review*, 85(4), 667–90, available at <http://ideas.repec.org/a/aea/aecrev/v85y1995i4p667-90.html>.
- JONES, C. I. (1997): “On the Evolution of the World Income Distribution,” *Journal of Economic Perspectives*, 11, 19–36.
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 73(1), 39–68.
- KORNAI, J. (2002): *The Socialist System: The Political Economy of Communism*. Clarendon Press, Oxford.
- KREMER, M., A. ONATSKI, AND J. STOCK (2001): “Searching for prosperity,” *Carnegie-Rochester Conference Series on Public Policy*, 55(1), 275–303, available at <http://ideas.repec.org/a/eee/crcspp/v55y2001i1p275-303.html>.
- KRUEGER, A. O. (1997): “Trade Policy and Economic Development: How We Learn,” *The American Economic Review*, 87(1), 1–22.
- KRUSELL, P., AND J.-V. RIOS-RULL (1996): “Vested Interests in a Positive Theory of Stagnation and Growth,” *The Review of Economic Studies*, 63(2), 301–329.
- LITTERMAN, R. (1986): “Forecasting With Bayesian Vector Autoregressions – Five Years of Experience,” *Journal of Business and Economic Statistics*, 4, 25–38.
- LUCAS, R. E. (1981): “Methods and Problems in Business Cycle Theory,” in *Studies in Business-Cycle Theory*, ed. by R. E. Lucas. MIT Press, Cambridge, MA.
- MADDISON, A. (2006): *The World Economy*. Development Center of the Organisation for Economic Cooperation and Development, Paris, France.
- MANSKI, C. F. (1993): “Identification of Endogenous Social Interaction: The Reflection Problem,” *Review of Economic Studies*, 60(3), 531–542.
- MARSHALL, M. G., AND K. JAGGERS (2005): “Polity IV Project: Dataset Users’ Manual,” Mimeo, George Mason University.

- MORRIS, S. (1995): “The Common Prior Assumption in Economic Theory,” *Economics and Philosophy*, 11, 227–253.
- MUKAND, S. W., AND D. RODRIK (2005): “In Search of the Holy Grail: Policy Convergence, Experimentation, and Economic Performance,” *American Economic Review*, 95(1), 374–383.
- PIKETTY, T. (1995): “Social Mobility and Redistributive Politics,” *The Quarterly Journal of Economics*, 110 (3), 551–584.
- PRIMICERI, G. E. (2006): “Why Inflation Rose and Fell: Policymakers’ Beliefs and US Postwar Stabilization Policy,” *The Quarterly Journal of Economics*, 121(3), 867–901.
- RODRIGUEZ, F., AND D. RODRIK (2000): “Trade Policy and Economic Growth: A Skeptic’s Guide to the Cross-National Evidence,” *NBER Macroeconomics Annual*, 15, 261–325.
- SACHS, J. D., AND A. WARNER (1995): “Economic Reform and the Process of Global Integration,” *Brookings Papers on Economic Activity*, 1995(1), 1–118.
- SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton, New Jersey.
- SARGENT, T. J., N. WILLIAMS, AND T. ZHA (2006): “Shocks and Government Beliefs: The Rise and Fall of American Inflation,” *The American Economic Review*, 96(4), 1193–1124, mimeo, New York University.
- STRULOVICI, B. (2008): “Voting and Experimentation,” mimeo, Northwestern University.
- WACZIARG, R., AND K. H. WELCH (2008): “Trade Liberalization and Growth: New Evidence,” *World Bank Economic Review*, 22(2), 187–231, NBER working paper no. 10152.

UNIVERSITY OF CALIFORNIA, LOS ANGELES AND NBER

PENNSYLVANIA STATE UNIVERSITY

NORTHWESTERN UNIVERSITY, CEPR AND NBER