Learning about Growth and Democracy

April 30, 2019

Abstract

We develop and estimate a model of learning that explains the observed correlation between income and democracy as well as the clustering of democratization events. In our model, countries' own and neighbors' past experiences shape elites' beliefs about the effects of democracy on economic growth and their likelihood of retaining power. These beliefs influence the choice to transition into or out of democracy. We show that learning from past experiences is crucial to explaining observed transitions since the mid-twentieth century. Moreover, our model predicts reversals to authoritarianism if the world experienced a growth shock the size of the Great Depression.

Keywords: democracy, learning, growth, diffusion
Word Count: 11,881

1 Introduction

Empirical scholarship on the causes of democracy has sought to understand two patterns: the strong correlative relationship between levels of material well-being and democracy (Lipset, 1959; Przeworski et al., 2000; Boix and Stokes, 2003; Acemoglu et al., 2008, 2009), and the spatial and temporal clustering of democratization events (Huntington, 1993; Brinks and Coppedge, 2006; Gleditsch and Ward, 2006; Ahlquist and Wibbels, 2012; Houle, Kayser and Xiang, 2016). Existing studies have treated these as distinct objects of inquiry, separately assessing the impact of domestic and international factors on the propensity to democratize. In this paper, we develop and estimate a model of elite belief formation that combines both domestic and systemic features in order to jointly explain the correlation between economic development and democracy and the clustering of democratic transitions.

We explicitly model the choice by incumbent elites to promote or subvert democracy.[1] This choice impacts their likelihood of retaining power both directly and through its effect on economic growth. Incumbents are uncertain about the relationship between democracy and growth, and they rely upon worldwide economic history to update their beliefs. We allow beliefs to be spatially correlated so that incumbents may learn more from the experiences of more proximate (or similar) countries.[2] In accordance with their beliefs, incumbents pursue democracy or autocracy, seeking to maximize the probability they remain in power. For a panel of 151 countries, we use data from 1875-1950 to calibrate initial conditions and data from 1951-2000 to structurally estimate our model.[3]

To assess the ability of our learning model to explain observed patterns of economic growth and democracy adoption, we conduct a series of goodness-of-fit and out-of-sample

[1]We conceive elites broadly as the group or faction in charge of the executive. See Section 2 for a more detailed discussion. For an overview of the formal literature that emphasizes the endogenous nature of institutions, see Gehlbach, Sonin and Svolik (2016) and Svolik (2019).

[2]Geographic distance is highly correlated with various measures of similarity between countries. Our results are robust to accounting directly for other types of similarity.

[3]Methodologically, our approach is similar to that of recent studies in political science that estimate model parameters by maximizing a likelihood derived from the equilibrium conditions of a formal model (Ascencio and Rueda, 2018; Crisman-Cox and Gibilisco, 2018).

(2001-2010) prediction exercises that pit our model against a range of reduced-form panel regressions typical of the approach taken in the existing empirical literature on democratization (Acemoglu et al., 2008; Boix, 2011). We show that learning from past experiences is crucial to explaining observed transitions to and from democracy, delivering an improvement in predictive success of over 100% relative to the best-fitting specification that does not account for learning, even after allowing for other potential channels of diffusion.

The success of our learning model is rooted, first, in our estimates of the political implications of economic growth. In line with a sizable empirical literature in political economy, we find that democracies tend to reward incumbents for growth by keeping them in power (Hibbs, 1977; Alesina, Roubini and Cohen, 1997; Brender and Drazen, 2008). In contrast, we find that growth tends to be destabilizing in autocracies. That is, our estimates are consistent with the view that rapid economic expansion in autocracies produces actors—a middle class, for example—able to challenge the group in power (Olson, 1963; Huntington, 1968; Hirschman and Rothschild, 1973). Together, these results suggest that the cross-sectional correlation between levels of material well-being and democracy is largely driven by elites who seek to benefit politically from the economic consequences of institutional choice. In particular, democratic incumbents will subvert democracy when they come to believe that it does not produce sufficient economic growth to win a fair election, while, conversely, autocratic incumbents will only transition to democracy when expected rates of growth make them more likely to retain power via election than under continued authoritarian rule.

Importantly, we distinguish transitions of power, where only the identity of the incumbent changes, from transitions into or out of democracy, where the form of government changes. Sudden changes in economic conditions, for instance, may lead to generalized political instability regardless of the system of government in place. However, whenever an incumbent is replaced—be it through election, coup, or revolution—the new group in power again faces the choice to support or subvert democracy. Our focus is on this choice, not transitions of power between factions.

The second feature of our learning model that underpins its empirical success is its ability to capture the wave-like nature of democratic transitions. We find that beliefs about the relationship between democracy and economic growth are highly correlated both temporally and spatially, which provides a structural interpretation for the observed clustering of transitions. Our estimates indicate there is an approximately 5,000 kilometer radius within which learning occurs. Outside this distance, virtually no additional information is gleaned. This advances the literature on the diffusion of democracy, which has struggled to disentangle competing mechanisms.[4]

Our methodological approach allows us to conduct counterfactual experiments of three types, each of which highlights the importance of learning. First, it enables us to understand how systemic shocks to prosperity impact the worldwide prospect for democracy. Specifically, our model predicts considerable reversals to authoritarianism if the world were hit with a shock to growth the size of the Great Depression. Second, it allows us to ask retrospective questions about historical democratization events, such as whether Greece, Portugal, and Spain would have democratized when they did had western European democracies suffered a recession in the early 1970s. Third, our model allows us to prospectively explore conditions that would presently lead countries to transition to or from democracy. Overall, our results suggests that, ultimately, democracy is a fragile system of government, one which depends significantly upon its own economic success.

Taken together, our findings contribute to a substantial body of work on modernization and democracy. Dating to at least the mid-twentieth century, social scientists have debated whether increases in per capita income have a causal effect on the probability that a state democratizes.[5] We bring to bear two innovations. First, we do not rely on the instrumental-

[4]Proposed mechanisms include diffusion through international organizations (Pevehouse, 2005), direct emulation of neighbors (Gleditsch and Ward, 2006), diffusion through trade and economic exchange (Mansfield, Milner and Rosendorff, 2000), cultural linkages (Wejnert, 2005), and military coercion (Kadera, Crescenzi and Shannon, 2003). On the inability of this literature to empirically falsify any particular mechanism of diffusion, see Torfason and Ingram (2010).

[5]The modernization hypothesis—that higher incomes per capita cause countries to democratize—dates to, at least, Lipset (1959). For evidence in favor, see Londregan and Poole (1996); Barro (1996); Boix and Stokes (2003); Boix (2011). Against, see Przeworski and Limongi (1997); Przeworski et al. (2000); Acemoglu

variables or reduced-form selection-on-observables identification strategies deployed—with mixed results—in the extant scholarship. Rather, we propose an explicit model of the relationship between economic growth and democracy that we take directly to the data. Second, and perhaps more significantly, we approach modernization as a systemic phenomenon. Our results suggest that focusing on the within-country impact of economic development is too narrow an object of inquiry. Through its influence on neighboring agents' beliefs, a given country's economic performance affects not only its own likelihood of transitioning to or from democracy but also the prospect for democracy outside of its borders.

Of course, we are not the first to propose diffusion through learning as a cause of democratic transitions.[6] If direct and consistent measures of beliefs were available over a sufficiently long period and a wide enough set of countries, it would be conceivable to directly estimate the impact of changing beliefs on democratization.[7] Given the current lack of systematic or reliable data, our paper represents the first attempt at estimating this impact.

We are also not the first to examine the role of learning in policymaking more broadly.[8] Our paper is most closely related to Buera, Monge-Naranjo and Primiceri's (2011) structural analysis of the impact of learning on the adoption of market-oriented (versus state-interventionist) policies. We build on their framework to model the worldwide evolution and diffusion of beliefs about the economic consequences of competing policies. Yet, while they consider the problem of a welfare-maximizing social planner, we take a political economy perspective and study institutional design as the outcome of self-interested choices made by power-seeking elites—a natural next step in this line of inquiry.

et al. (2008, 2009).

[6]Dahl (1998); Diamond (2011); Miller (2016)

[7]For work that attempts to gauge beliefs and their consequences for democratization in a subset of countries, see Almond and Verba (1963); Norris (1999); Chen and Lu (2011).

[8]Foster and Rosenzweig (1995); Primiceri (2006); Conley and Udry (2010); García-Jimeno (2016).

2 A Learning Model of Democratization

2.1 Setup: Elites, Beliefs, and Learning

We consider the decision problem of the decisive group or political actor in power in country i at time t.[9] This agent faces a choice between autocracy, $D_{i,t} = 0$, and democracy, $D_{i,t} = 1$. The incumbent's objective is to retain power in period $t + 1$. Let $Y_{i,t}$ denote country i's per capita GDP in period t, and let $y_{i,t} \equiv \log(Y_{i,t}) - \log(Y_{i,t-1})$ denote its growth rate. At the beginning of period t, to maximize her expected probability of remaining in power, the incumbent solves

$$\max_{D_{i,t} \in \{0,1\}} E_{i,t-1} \left[\frac{\exp(\alpha_i + \theta^{D=D_{i,t}} y_{i,t} - K_{i,t} D_{i,t})}{1 + \exp(\alpha_i + \theta^{D=D_{i,t}} y_{i,t} - K_{i,t} D_{i,t})} \middle| D_{i,t} \right], \tag{1}$$

where the integrand represents the probability of retaining power in period $t + 1$ and the expectation is taken conditional on the information available in country i at the conclusion of period $t - 1$. The integrand is increasing in the index $\alpha_i + \theta^{D=D_{i,t}} y_{i,t} - K_{i,t} D_{i,t}$, where coefficient α_i establishes a baseline for country i, coefficients $\theta^{D=0}$ and $\theta^{D=1}$ respectively measure the (de)stabilizing effect of GDP growth on elite turnover under autocracy and democracy, and $K_{i,t}$ captures the political cost of democracy to the incumbent, i.e., its direct effect on her likelihood of retaining power.

The incumbent chooses $D_{i,t}$ at the start of period t, forming a subjective forecast of its effect on GDP growth, $y_{i,t}$, to solve (1).[10] Incumbents believe that the relationship between GDP growth and democracy takes the form

$$y_{i,t} = (1 - D_{i,t})\beta_i^{D=0} + D_{i,t}\beta_i^{D=1} + \epsilon_{i,t}, \tag{2}$$

where $\beta_i^{D=0}$ and $\beta_i^{D=1}$ denote country i's long-run GDP growth rates under autocracy and

[9] Both in democracies and autocracies, we treat this as the political party or faction (when parties do not exist) in control of the executive—see Section 2.2.

[10] In what follows, we build on the learning framework of Buera, Monge-Naranjo and Primiceri (2011).

democracy, respectively, and $\epsilon_{i,t}$ is an exogenous shock to growth that is uncorrelated over time but potentially correlated across countries.[11] Specifically, the vector $\epsilon_t \equiv [\epsilon_{1,t}, \ldots, \epsilon_{n,t}]'$ of GDP growth shocks across the n countries of the world is independently and identically distributed (i.i.d.) over time according to a mean zero Normal distribution with covariance matrix Σ, i.e.,

$$\epsilon_t \overset{\text{i.i.d.}}{\sim} N(0, \Sigma).$$

Incumbents do not know $\beta_i \equiv [\beta_i^{D=0}, \beta_i^{D=1}]'$, but they have perfect knowledge of all other features of the model, including $K_{i,t}$, at the time of their choice.

The timing of events is as follows. At the end of period $t-1$, incumbents collect data on worldwide GDP growth rates and systems of government, and they update their beliefs about long-run economic growth under autocracy and democracy accordingly. At the beginning of period t, incumbents observe $K_{i,t}$ and decide what system of government, $D_{i,t}$, to adopt that period. Growth shocks and corresponding growth rates, conditional on incumbents' choices, are then realized.

Learning. In period $t = 0$, incumbents start out with a Normal prior over the vector of unknown long-run GDP growth rates $\beta \equiv [\beta_1^{D=0}, \ldots, \beta_n^{D=0}, \beta_1^{D=1}, \ldots, \beta_n^{D=1}]'$,

$$\beta \sim N(\bar{\beta}_0, P_0^{-1}), \tag{3}$$

where $\bar{\beta}_0$ and P_0 denote, respectively, the prior mean and precision matrix. We assume that incumbents' initial beliefs assign no correlation and the same degree of uncertainty to growth under autocracy and democracy:

$$P_0^{-1} = I_2 \otimes (V \cdot R \cdot V),$$

[11] Existing studies on the causal impact of democracy on economic growth have highlighted various potential mechanisms—most notably, how democracy influences redistributive policy, which in turn affects investment (Alesina and Rodrik, 1994; Persson and Tabellini, 1994; Acemoglu et al., 2014). For tractability, we abstract from considering these explicitly but allow incumbents to have flexible, country-specific beliefs over their net long-run impact.

where $V = \text{diag}([v_1\sigma_1, \ldots, v_n\sigma_n])$ is a diagonal matrix whose ith diagonal entry measures prior uncertainty (standard deviation) about country i's long-run growth rate under autocracy/democracy, and R is the cross-country prior correlation matrix. Prior uncertainty is parameterized by $\{v_i\}_{i=1}^n$, normalized by the standard deviation of growth shocks in each country $\{\sigma_i\}_{i=1}^n$ (the square roots of the diagonal elements of Σ).

Our assumptions yield simple, recursive Bayesian updating formulas for beliefs in each period: letting $D_t \equiv [D_{1,t}, \ldots, D_{n,t}]'$ and $y_t \equiv [y_{1,t}, \ldots, y_{n,t}]'$,

$$
\begin{aligned}
P_t &= P_{t-1} + \mathbf{D}_t'\Sigma^{-1}\mathbf{D}_t, \\
\bar{\beta}_t &= \bar{\beta}_{t-1} + P_t^{-1}\mathbf{D}_t'\Sigma^{-1}(y_t - \mathbf{D}_t\bar{\beta}_{t-1}),
\end{aligned}
$$

where $\mathbf{D}_t \equiv [\text{diag}(1 - D_t), \text{diag}(D_t)]$ is an $n \times (2n)$ matrix such that the ith element of the vector $\mathbf{D}_t\bar{\beta}_{t-1}$ equals $(1 - D_{i,t})\bar{\beta}_{i,t-1}^{D=0} + D_{i,t}\bar{\beta}_{i,t-1}^{D=1}$. The impact of new data on the posterior mean of β is determined by $P_t^{-1}\mathbf{D}_t'\Sigma^{-1}$, which depends on three key factors. First, higher initial uncertainty in beliefs (higher $\{v_i\}_{i=1}^n$) raises the relative precision of new information, increasing its impact. Second, higher correlation in growth shocks across countries (off-diagonal elements of Σ) reduces the informational content of observed growth rates and slows down learning. Lastly, higher cross-country correlation in initial beliefs (off-diagonal elements of R) increases belief responsiveness to data from other countries.

We allow incumbents to potentially learn more from neighboring (or more similar) countries. Letting $Z_{i,j}$ denote a vector that may include various measures of distance (geographic or otherwise) between countries i and j, we write

$$R_{i,j} = \exp(-Z_{i,j}'\gamma),$$

where γ is constrained to be nonnegative to ensure correlations between 0 and 1.[12]

[12]This formulation also guarantees the positive definiteness of R (Matérn, 1960).

Incumbents' optimal choice. While incumbents observe the political cost of democracy, $K_{i,t}$, prior to choosing $D_{i,t}$, this cost is unobservable to the researcher. We assume that $K_{i,t}$ has the following structure:

$$K_{i,t} = f_i + X'_{i,t}\xi + \kappa_{i,t}. \tag{4}$$

Coefficient f_i establishes a country-specific baseline, and the control vector $X_{i,t}$ may include various observable economic and political characteristics of country i (e.g., lagged per capita GDP or incumbents' time in power). Every period, country i also experiences an exogenous idiosyncratic shock, $\kappa_{i,t}$, to the political cost of democracy, where

$$\kappa_{i,t} \sim N(0, \varsigma_i^2).$$

The volatility of shocks to the political cost of democracy, ς_i, is allowed to be country-specific, but $\kappa_{i,t}$ is assumed to be independently distributed over time and across countries.

As discussed, when choosing $D_{i,t}$ incumbents have perfect knowledge of $K_{i,t}$ and all features of the model except for the effect of their choice on GDP growth, $y_{i,t}$. Together, (1), (2), and (4) imply that the optimal choice for country i's incumbent at time t is

$$\begin{aligned} D_{i,t} = \mathbf{1}\Bigg\{ E_{i,t-1}\left[\frac{\exp(\alpha_i + \theta^{D=1}(\beta_i^{D=1} + \epsilon_{i,t}) - f_i - X'_{i,t}\xi - \kappa_{i,t})}{1 + \exp(\alpha_i + \theta^{D=1}(\beta_i^{D=1} + \epsilon_{i,t}) - f_i - X'_{i,t}\xi - \kappa_{i,t})}\right] \\ > E_{i,t-1}\left[\frac{\exp(\alpha_i + \theta^{D=0}(\beta_i^{D=0} + \epsilon_{i,t}))}{1 + \exp(\alpha_i + \theta^{D=0}(\beta_i^{D=0} + \epsilon_{i,t}))}\right]\Bigg\}, \end{aligned} \tag{5}$$

where the expectations are taken only with respect to β_i and $\epsilon_{i,t}$ in accordance with the incumbent's beliefs at the conclusion of period $t - 1$.

Opposition groups and strategic experimentation. To conclude the description of our model, we briefly discuss how we account for non-elite learning and strategic interactions between incumbents and potential challengers.

The common prior assumption for incumbents in our model extends to all potential stakeholders in each country. Opposition groups (elite or non-elite) observe the same worldwide history of economic growth and democracy, and they would be faced with solving (1)—in the event they came to power—using information identical to that available to the incumbent. As a result, in this shared learning environment, the identity of the incumbent only matters by shaping the political cost of democracy.

For tractability, we abstract from explicitly modeling the intricacies of within-country elite turnover. Nevertheless, objective (1) can be viewed as describing an equilibrium probability of staying in power resulting from a richer interaction between the incumbent and potential challengers, both elite and non-elite. Importantly, we distinguish transitions of power, where only the identity of the incumbent changes, from transitions into or out of democracy. Whenever an incumbent is overthrown by a rival elite faction or via a revolution from below, the new group in power again faces the choice to support or subvert democracy.[13] Understanding how this choice by self-interested elites—newly in power or entrenched—is shaped by the evolution of beliefs about the economic effects of democracy is the focus of this paper.[14]

Finally, objective (1) precludes incumbents from adopting a form of government with negative expected consequences solely for the purpose of learning from the experience. In our setting, the prospect of losing power and thus not reaping the benefits of such experimentation limits its appeal. Relatedly, incumbents in our model are myopic and focused only on their immediate survival. While a fully dynamic version of our model with shared learning by forward-looking agents would introduce strategic experimentation incentives that would render the model intractable,[15] we believe our model offers a good first-order approximation

[13] Examples abound of revolutions from below, inspired by ostensibly democratic goals, that failed to deliver on the promise of liberal democracy. For instance, Skocpol's (1979) three main cases—the French, Chinese, and Russian revolutions—all began as mass revolutionary movements with outwardly democratic motives and each, nonetheless, resulted in dictatorship.

[14] An implication of our model is that transitions to or from democracy constitute attempts by incumbent elites to hold on to power. Notably, incumbents in our data retain power 30% and 19% of the time following transitions to and from democracy, respectively, despite the instability typically associated with transition periods.

[15] Bolton and Harris (1999); Bramoulleé, Kranton and D'Amours (2014); Mossel, Sly and Tamuz (2015).

to optimal behavior by incumbents with longer time horizons.[16]

2.2 Empirical Strategy

Like incumbents in our model, we adopt a Bayesian inference approach to recover the unknown structural parameters of our model, listed in Table 1.[17] With a slight abuse of notation, let the vector $\boldsymbol{\theta}$ collect all the parameters in Table 1, and let $I_{i,t}$ be an indicator of whether the incumbent in country i retained power ($I_{i,t} = 1$) or not ($I_{i,t} = 0$) at the conclusion of period t. Denote by $W^T \equiv \{I_t, y_t, D_t, X_t\}_{t=1}^T$ the set of all data available up to period T, where $I_t \equiv [I_{1,t} \ldots, I_{n,t}]'$ and $X_t \equiv [X_{1,t} \ldots, X_{n,t}]'$. Our goal is to estimate the true value of $\boldsymbol{\theta}$ by computing the mode of the posterior distribution of the model parameters,

$$p(\boldsymbol{\theta}|W^T) \propto \mathcal{L}(W^T|\boldsymbol{\theta})\pi(\boldsymbol{\theta}),$$

given the likelihood of the data, $\mathcal{L}$, and our prior, π. We describe $\mathcal{L}$ and π in turn.

Table 1: Model Parameters

$\{\alpha_i\}_{i=1}^n$:	baseline incumbent stability
$\theta^{D=0}$:	effect of GDP growth on elite turnover under autocracy
$\theta^{D=1}$:	effect of GDP growth on elite turnover under democracy
$\{\bar{\beta}_{i,0}^{D=0}\}_{i=1}^n$:	prior mean of long-run GDP growth rate under autocracy
$\{\bar{\beta}_{i,0}^{D=1}\}_{i=1}^n$:	prior mean of long-run GDP growth rate under democracy
$\{\upsilon_i\}_{i=1}^n$:	prior uncertainty about economic effects of autocracy/democracy
γ:	coefficients of cross-country correlation of prior beliefs
$\{f_i\}_{i=1}^n$:	baseline political cost of democracy
ξ:	coefficients of economic/political controls for political cost of democracy
$\{\varsigma_i\}_{i=1}^n$:	volatility of political cost of democracy

Likelihood of the data. While the structure of our model described thus far specifies incumbents' beliefs about how the data are generated as well as their optimal choices given

[16] See Section 3 and Appendix A5.

[17] Following Buera, Monge-Naranjo and Primiceri (2011), to reduce the dimensionality of the model we set Σ equal to its estimated value from the "true" data generating process (see Section 4).

those beliefs, we have refrained from specifying the "true" data generating process (DGP). In Section 4, to perform counterfactual experiments, we discuss and specify the true DGP. Next, we only make one key assumption about the true DGP that simplifies inference about the model parameters.

We assume that observed outcomes are only affected by actual choices and not by the beliefs that led to those choices. That is, transitions of power (I_t), GDP growth (y_t), and other economic and political characteristics of countries (X_t) are shaped by realized institutions (D_t), but they are not directly affected by beliefs about the potential effects of transitioning into or out of democracy. Formally, this assumption implies that the parameters in Table 1 are only involved in the component of the likelihood that describes the conditional probabilities of countries' observed systems of government. While we relegate a full derivation of the likelihood to Appendix A1, with a slight abuse of notation—using $\mathcal{L}$ to denote arbitrary densities of the data—it can be written as

$$\mathcal{L}(W^T|\boldsymbol{\theta}) \propto \prod_{t=1}^{T}\prod_{i=1}^{n}\mathcal{L}(D_{i,t}|X_{i,t}, W^{t-1}, \boldsymbol{\theta}),$$

where

$$\mathcal{L}(D_{i,t}|X_{i,t}, W^{t-1}, \boldsymbol{\theta}) = \Phi\left(\frac{\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})}{\varsigma_i}\right)^{D_{i,t}}\left[1 - \Phi\left(\frac{\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})}{\varsigma_i}\right)\right]^{1-D_{i,t}}, \quad (6)$$

Φ denotes the standard Normal cumulative distribution function, and $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ is the threshold value of $\kappa_{i,t}$—the realized shock in period t to the political cost of democracy in country i—that leaves country i's incumbent indifferent between autocracy and democracy. Note that (6) resembles the likelihood of a standard binary-choice Probit model. However, $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ is a nonlinear function (with no closed-from expression) of the data up to period t and the model parameters that encodes how each country's propensity for democracy evolves with incumbents' beliefs.

Prior. Given the size of our model, we adopt an informative prior, π, to prevent overfitting. To do so in a principled manner, we calibrate our prior in the way agents in our model would, allowing the observed past to inform initial beliefs. We use data from 1875-1950 (excluding the two world wars), a period that immediately precedes our main sample, to set the prior mean and precision of the model parameters so as to match analogous empirical moments. For example, we ensure that our prior over incumbents' initial beliefs about the relationship between democracy and GDP growth is consistent with average annual growth rates among autocracies and democracies in the pre-sample period. Similarly, we use pre-sample history of elite turnover to inform our prior over the parameters describing the likelihood of retaining power. For a complete description of our prior and how it is calibrated, see Appendix A3.

Estimation and inference. Calculating $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ to evaluate the likelihood of the data is computationally expensive.[18] To avoid this burden, we follow the Mathematical Programming with Equilibrium Constraints (MPEC) approach of Su and Judd (2012) to compute our maximum-a-posteriori estimator of $\boldsymbol{\theta}$.[19] The idea behind this approach is simple: instead of calculating $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ at every trial value of $\boldsymbol{\theta}$, treat $\overline{\kappa}_{i,t}$ as auxiliary parameters and impose the optimality (or equilibrium) conditions of the model as feasibility constraints on the log-posterior maximization program. This considerably reduces the computational cost of estimating the model. We describe our estimation strategy in detail in Appendix A4.

Data. For our main analysis, we obtain data from three sources, each measured at the country-year. First, we obtain data on **GDP & Per Capita GDP** from Maddison (2010). The Maddison Project Database provides information on comparative economic growth and income levels over the very long run. The data give estimates of GDP and GDP per capita annually from 1875-2008 for all of the independent states in our sample of countries. We

[18]We provide an iterative algorithm for computing $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ in Appendix A2.

[19]Standard errors are parametrically bootstrapped.

obtain two additional years (2009-2010) of GDP per capita growth rates from the Penn World Table (Feenstra, Inklaar and Timmer, 2015) for out-of-sample predictions.

Second, we obtain our dichotomous measure of **Democracy** from Boix, Miller and Rosato (BMR, 2013). This dataset provides an annual coding of democracy for every country in the world from 1800 to 2010. If the following three criteria are met, then countries are coded as democratic:

1. The executive is directly or indirectly elected in popular elections and is responsible either directly to voters or to a legislature.
2. The legislature (or the executive if elected directly) is chosen in free and fair elections.
3. A majority of adult men has the right to vote.

If these three criteria are not met, a country is coded as not democratic. While various alternative measures of democracy have been used in previous studies, the BMR coding is the most comprehensive and consistent for the period we cover (1875-2010). Nonetheless, our results are robust to employing alternative codings (see Appendix A5).

Finally, for each country-year, we code the **Executive Faction** from Goemans, Gleditsch and Chiozza (2009). The Archigos database on leaders describes the date and manner of entry and exit for the executives of all countries in our sample from 1875-2015. With these data, we then code, using biographical information, the political party of each executive. If an executive is non-democratically elected and we cannot identify a political party (nearly all of these cases are military regimes), we identify the particular faction to which the executive belongs. In combination with the entry/exit dates, we construct our measure of change in the faction of the executive (elite turnover).

3 Estimation Results: The Importance of Learning

Before summarizing our structural parameter estimates, we subject our model to a series of goodness-of-fit and out-of-sample prediction tests that assess its ability to explain observed patterns of democracy adoption. We consider five alternative specifications of our model that differ in the number of covariates used to characterize the political cost of democracy. In our baseline specification with no covariates, the political cost of democracy, $K_{i,t}$, consists of simply a country-specific baseline, f_i, plus an idiosyncratic shock, $\kappa_{i,t}$. We then consider specifications where we successively control for (lagged) log-GDP per capita, the incumbent's time in power, (lagged) trade volume as percentage of GDP (Gleditsch, 2002), and years as democratic (negative when autocratic) to account for consolidation effects (Svolik, 2013). Across specifications, we use geographic distance between capitals, $Z_{i,j}$, to capture the cross-country correlation in initial beliefs.[20]

To quantify the importance of learning for our model's ability to fit the data, we also estimate a "no-learning" version of our model. For each specification, we constrain beliefs about long-run growth rates under autocracy and democracy to be constant over time, thus shutting down the learning mechanism. These no-learning specifications are otherwise identical to their learning counterparts.

We conduct our model performance tests as follows. With each estimated model, we compute one-year-ahead forecasts of the choice between autocracy and democracy for each country-year. That is, conditional on the state of the world at the end of year $t-1$ as recorded in our data, we use (5) for each model to predict $D_{i,t}$ worldwide. We produce forecasts for the in-sample period used to estimate each model (1951-2000) and for 10 additional out-of-sample years (2001-2010).

In Figure 1, we plot the actual (gray) and predicted percentage of world democracies. In the top panel, predictions are generated using our baseline specification with no covariates. We plot predictions with (blue) and without (red) learning for both the in-sample (solid) and out-of-sample (dashed) periods. In the lower panel, we present the same set of estimates using our model with two covariates (lagged log-GDP per capita and incumbents' time in power).

[20] Geographic distance is highly correlated with other measures of cultural, economic, or political similarity between countries (Buera, Monge-Naranjo and Primiceri, 2011). Our results are virtually unchanged (available upon request) if we model the correlation in initial beliefs as a function of geographic and genetic distance between countries, the latter as measured by Spolaore and Wacziarg (2009).

Note the vast improvement in predictive success, both in and out of sample, when we account for learning. Unsurprisingly, the no-learning specification of our model with no covariates performs worst as it produces a constant prediction for each country.[21] However, while the inclusion of covariates does markedly improve the accuracy of the no-learning model, these gains pale in comparison to the role of learning. Indeed, our baseline learning model with no covariates vastly outperforms any specification that does not account for learning. As many of the covariates we condition upon are themselves outcomes of the selection process we model (e.g., per capita GDP or elite turnover), their inclusion should not yield much improvement in predictive success. Our results confirm this intuition.[22]

Table 2 provides a numerical summary of our goodness-of-fit tests. Each set of columns corresponds to a different model specification, with (odd columns) and without (even columns) learning. The first row gives the percentage of country-year observations each model correctly predicts. Unsurprisingly, all models perform remarkably well on this dimension. The reason is that, as transitions into or out of democracy are quite rare (130 total in-sample events), country fixed effects go a long way in fitting the data. Indeed, our no-learning model with no covariates (second column), which produces a constant prediction for each country, has a success rate of almost 90%.

A much harder test—one that is considerably more revealing of the underlying causes of democracy—is whether a model can correctly predict transitions to and from democracy. In Table 2, we present two scenarios, assessing each model's accuracy in predicting transitions within ± 0 years (second row) and ± 2 years (third row) of the event. Here, the importance of learning is striking. Models that do not account for learning perform quite poorly, even within a 5-year window.[23] And, while including additional covariates does increase accuracy,

[21]The observed temporal variation is an artifact of the changing population of countries in our data.

[22]To show that the success of our model is not an artifact of aggregating predictions across countries, we provide in Appendix A5 nearly identical results disaggregated by four regions of the world: the Americas, Europe, Africa, and Asia-Oceania.

[23]The 5-year window predictions for models in the last two columns should be taken with care. When controlling for years as democratic (negative when autocratic) and aggregating over 5 years, transitions in the data are mechanically picked up by the models and turned into correct predictions. See below for a similar comment about models with a one-year lag of democracy. Notably, turning on the learning mechanism still

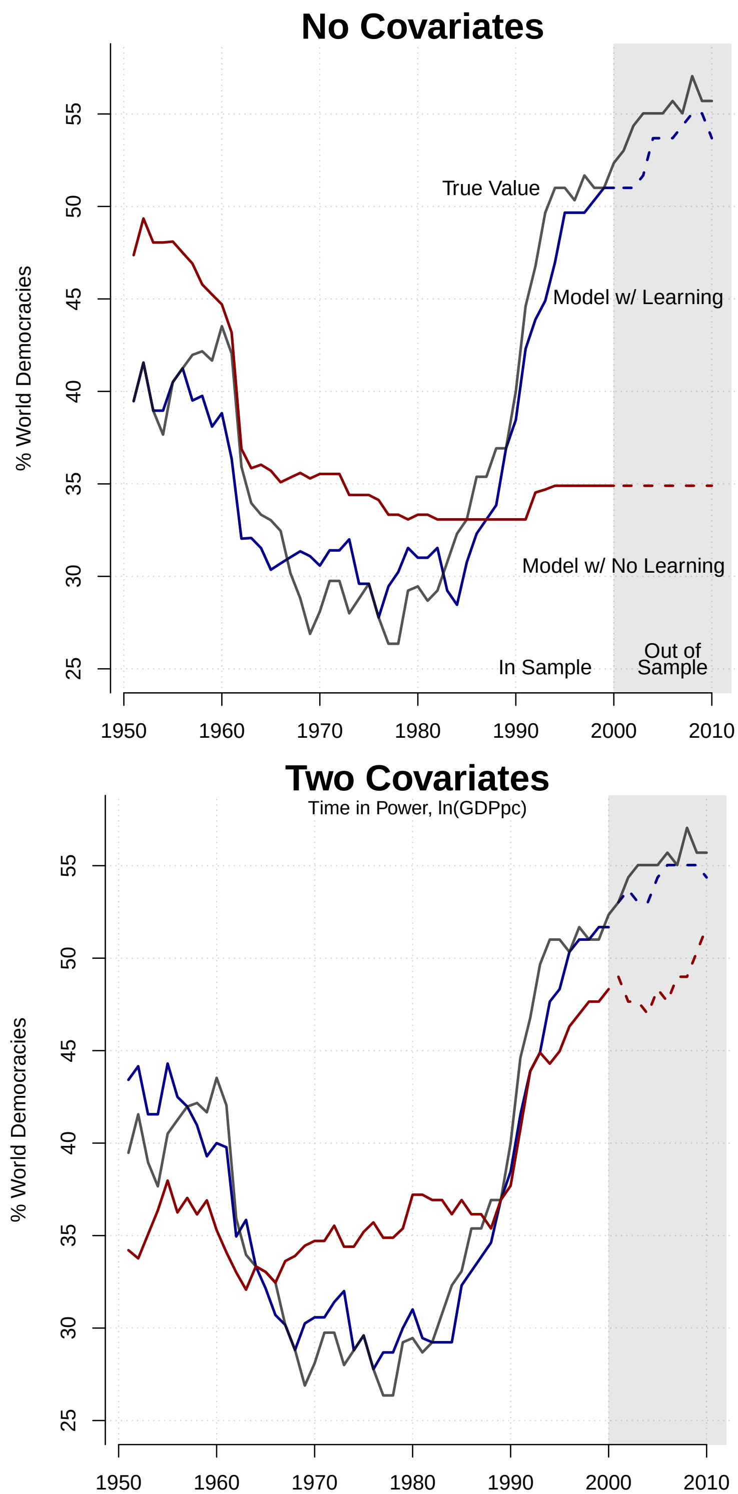

Figure 1: Observed versus Predicted Worldwide Prevalence of Democracy

Notes. This figure compares the true proportion of world democracies (gray) to in-sample (solid blue) and out-of-sample (dashed blue) estimates generated by our model. Additionally, we shut down learning in our model and produce both in-sample (solid red) and out-of sample (dashed red) predictions. In the top panel, estimates are generated using our baseline specification with no covariates. In the bottom panel, we control for lagged log-GDP per capita and incumbents' time in power.

Table 2: Model Fit

	No Covariates		One Covariate		Two Covariates		Three Covariates		Four Covariates	
	Learning	No Learning	Learning	No Learning	Learning	No Learning	Learning	No Learning	Learning	No Learning
Choices (% correct)	95.2	88.9	95.7	90.4	95.8	91.7	96.4	92.3	97.1	96.3
Transitions (% correct)										
± 0 years	9.3	0.0	6.2	1.6	7.0	3.9	11.0	7.1	6.3	0.8
± 2 years	41.1	0.0	43.4	7.8	48.1	20.2	52.0	23.6	70.1	53.5
Log-likelihood	-581.4	-1,390.8	-572.5	-1,176.8	-553.6	-1,061.8	-523.6	-995.1	-490.1	-704.4
Observations	5,925	5,925	5,925	5,925	5,925	5,925	5,845	5,845	5,845	5,845

Notes. This table reports various goodness-of-fit measures. Each set of columns corresponds to a different specification of our model, with (odd columns) and without (even columns) learning. Models in the first two columns use only country fixed effects to characterize the political cost of democracy. Models in the third and fourth columns control for lagged log-GDP per capita. Models in the fifth and sixth columns additionally control for incumbents' time in power. Models in the seventh and eighth columns also control for trade volume as a percentage of lagged GDP. Models in the last two columns add years as democratic (negative when autocratic) as a control. For each model, we report the percentage of correctly predicted in-sample system of government choices (first row). We similarly report the percentage of correctly predicted transitions to or from democracy within a 0-year window (second row) and a 5-year window (third row) of the event.

the marginal improvement is negligible. In contrast, turning on the learning mechanism in our model raises predictive success by over 100% in virtually all scenarios and all specifications. In fact, our baseline learning model with no covariates outperforms most no-learning specifications by a similar rate.[24]

To further benchmark our model, we present in Table 3 results from a series of panel regressions typical of the approach taken in the existing empirical literature on democratization. Using linear probability models, we regress our democracy measure against a full set of country fixed effects, a one-period lag of log-GDP per capita, and various controls. We exploit both annual data and, as in Acemoglu et al. (2008, 2009) and Boix (2011), five-year panels, which allow for the inclusion of covariates not available annually.[25] As before, with each specification we produce predictions for every country-period.

Again, as with our model, it is trivial to correctly predict close to 90% of country-period

delivers a sizable improvement in predictive power.

[24]The only exception is the model in the last column—see footnote 23.

[25]We also estimate five-year panel versions of our model (see Appendix A5), and results are essentially identical to those of Table 2, which may alleviate concerns about both the myopia of incumbents in our model and whether an annual timeframe is appropriate to study changes in system of government.

Table 3: Goodness of Fit of Reduced-Form Models

	I	II	III	IV	V	VI	VII	VIII	IX	X	XI	XII	XIII	XIV
Choices (% correct)	89.3	88.6	89.5	88.8	89.3	88.8	77.1	91.3	97.8	91.4	90.4	90.8	97.8	91.6
Transitions (% correct)														
±0 years	0.8		1.5		0.8				6.9		2.3		6.2	
±2 years	3.1	2.9	3.8	2.9	3.8	1.9	1.2	3.8	96.9	19.4	7.7	14.1	96.9	17.2
$\log(\text{GDPpc})_{t-1}$	0.115***	0.147***	0.116***	0.142***	0.112***	0.142***	0.033	0.032	0.024***	0.102***	0.060***	0.040	0.011	0.022
	(0.010)	(0.026)	(0.010)	(0.026)	(0.011)	(0.026)	(0.053)	(0.078)	(0.006)	(0.024)	(0.014)	(0.037)	(0.008)	(0.034)
Controls:														
Country Effects	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Time in Power			X	X	X	X	X	X	X	X	X	X	X	X
Trade/GDP					X	X	X	X	X	X	X	X	X	X
Education							X	X						
Labor Share								X						
Democracy_{t-1}									X	X			X	X
Time Effects											X	X	X	X
Observations	5,925	1,076	5,925	1,076	5,866	1,076	649	390	5,866	1,076	5,866	1,076	5,866	1,076
Panel Length	1	5	1	5	1	5	5	5	1	5	1	5	1	5

*p<0.1, **p<0.05, ***p<0.01

Notes. This table gives linear probability estimates of the impact of (lagged) log-GDP per capita on democracy. All models include country fixed effects. We successively add in controls for the incumbent's time in power, trade as a percentage of GDP, average years of schooling from Barro (1999), and labor share of value added from Rodrik (1999). In columns XI-XIV, we include year fixed effects. In columns IX, X, XIII, and XIV, we include a one-period lag of democracy. We present estimates with annual and five-year panels. Row 1 gives the percentage of country-period observations correctly predicted by each model. For models with annual data, we also report the success rate at predicting transitions within ±0 (row 2) and ±2 (row 3) years of the event. For the five-year panels, we produce success rates at predicting transitions just within the five-year window (row 3).

observations using country fixed effects. The important departure arises when we compare predictions of transitions derived from these reduced-form regressions to those generated by our learning model. Once more, we evaluate these predictions in exact (± 0) and five-year windows (± 2).

In terms of exact predictions, no reduced-form specification surpasses the predictive success of our baseline no-covariates learning model. In the five-year window, our model performs similarly well, save annual panel specifications with a lagged dependent variable, which correctly predict over 90% of transitions. However, aggregated over five years, the inclusion of the one-year lagged outcome leads any transition not picked up exactly by the model to be mechanically transformed into a correct prediction in subsequent years, thus yielding an extremely high predictive success rate in the five-year window and yet a low success rate within a single year. To further see this, note that, when we introduce a one-period lag in the five-year panel specifications, they successfully predict less than 20% of transitions, a rate which our model beats by more than 100%. More importantly, our model provides a structural interpretation for the observed persistence of systems of government that underpins the predictive success of these autoregressive specifications.

Of course, it may be the case that, rather than reflecting the learning process we describe, our model's success is simply an artifact of some alternative process of democratic diffusion that is indirectly picked up by our model's spatial and temporal flexibility. Examples of potential mechanisms proposed in the literature include direct emulation of neighbors (Gleditsch and Ward, 2006), cultural linkages (Wejnert, 2005), and diffusion through trade (Mansfield, Milner and Rosendorff, 2000). Due to data limitations, direct tests of each of these mechanisms would be impractical. Nonetheless, in the context of our model, the channel by which these alternative explanations could influence democratic transitions is through their impact on the political cost of democracy. To evaluate this possibility, we construct a distance-weighted measure of how democratic each country's neighborhood is over time, and we reestimate our baseline model using this measure as a control for direct diffusion effects. In Appendix A5, we find minimal increase in predictive success relative to our baseline model, which suggests that it is our proposed mechanism of learning about the economic effects of democracy—and not some alternative process of diffusion—driving our results.

In sum, our goodness-of-fit tests quantify and highlight the crucial role of learning in explaining worldwide democratization events and reversals to authoritarianism, overshadowing the usefulness of other explanatory variables typically employed in the literature. As discussed, this is not surprising given that many of these controls are themselves outcomes of the learning process we model. In light of these results, we hereafter focus our attention on the baseline no-covariates specification of our learning model, using it to conduct our counterfactual experiments in Section 4.

3.1 Structural Parameter Estimates

To understand how our model is fitting the data, we summarize our main parameter estimates and discuss their implications. We begin with our estimates of the (de)stabilizing effect of GDP growth on elite turnover under autocracy and democracy, $\theta^{D=0}$ and $\theta^{D=1}$, respectively. This question, itself, has been a separate subject of close academic inquiry for decades.[26] We find that the impact of growth on the likelihood that the incumbent group retains power indeed differs between autocracies and democracies.[27] The quantitative implications of our estimates are summarized in Figure 2, which plots the estimated probability (averaged across countries) that the incumbent remains in power at different rates of GDP growth. In blue, we plot our estimates under democracy and, in red, our estimates under autocracy.

In line with a substantial empirical literature in political economy, we find that, in democracies, economic growth is stabilizing for elites.[28] In other words, the party in government is more likely to win reelection when growth is high. In contrast, we find that growth is destabilizing in autocracies. This result comports with the view that rapid economic expansion in

[26] For a summary of the early literature on the topic, see Przeworski et al. (2000) ch. 1.

[27] Specifically, we estimate $\theta^{D=0} = -4.2213$ with a standard error of 2.2209, and $\theta^{D=1} = 8.8279$ with a standard error of 1.6368.

[28] Hibbs (1977); Alesina, Roubini and Cohen (1997); Brender and Drazen (2008).

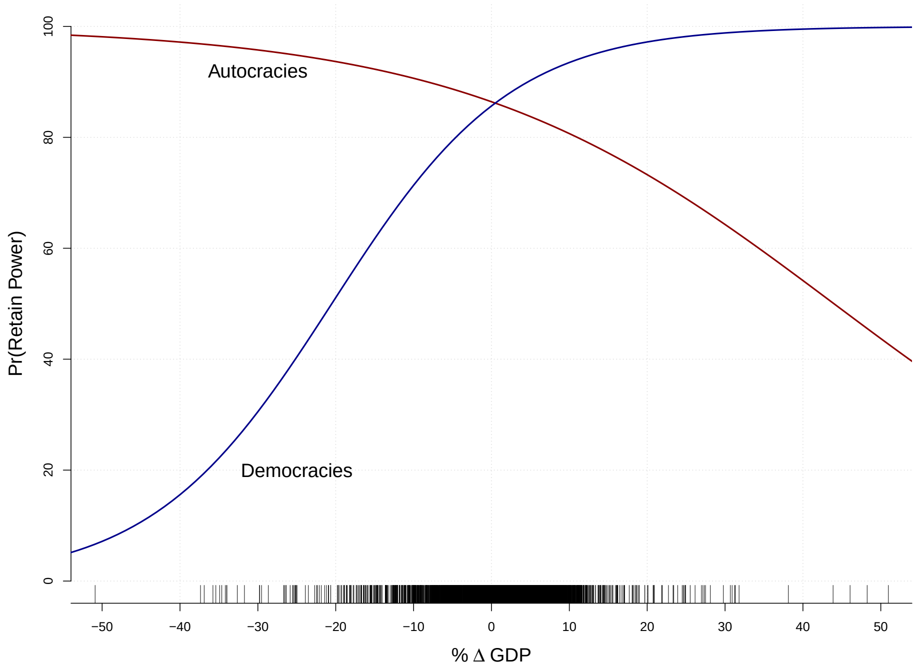

Figure 2: (De)Stabilizing Effect of GDP Growth on Elite Turnover

Notes. This figure plots the average (across countries) estimated probability of retaining power at different rates of GDP growth under autocracy (red) and democracy (blue). Observed growth rates between 1951-2000 are shown over the horizontal axis.

non-democracies creates inequalities of expectation or inequalities of outcomes that, in turn, engender attempts to subvert the political system.[29] That is, autocrats are less likely to remain in power when economic growth produces actors—a middle class, for example—able to place demands upon and challenge the authority of the group in power.

Coupling these results with learning helps explain the observed cross-sectional correlation between per capita GDP and democracy. Democracy becomes more appealing to incumbents as they come to believe that it is conducive to high rates of GDP growth. Conversely, autocracy entails an incentive to suppress economic growth.

Of course, there are notable exceptions. For example, over the past two decades China has experienced high rates of growth and nevertheless remained autocratic. Similarly, India for the first four decades of its independence experienced low rates of growth and yet remained democratic. Underlying these prominent cases is the country-specific political cost of

[29] Olson (1963); Huntington (1968); Hirschman and Rothschild (1973).

democracy. We recover estimates of the structural parameters describing the baseline political cost of democracy for each country, f_i, and plot them in Figure 3. Notably, the Chinese Communist party faces, all else equal, the lowest probability of remaining in power under democracy. In contrast, Congress at India's independence had the fourth highest ex-ante probability of retaining power under democracy.

Next, in Figure 4, we present estimates of the spatial decay of learning, i.e., the extent to which the cross-country correlation of beliefs depends on the geographic distance between capitals.[30] Consistent with the observed spatial clustering of democratization events noted in the literature, we find that learning is highly correlated geographically. As shown in Figure 4, at 5,000km (the approximate distance between the U.S. and Ecuador), the prior correlation of beliefs is only 0.12. At 10,000km (the distance between the U.S. and the Central African Republic), the potential for learning between countries is virtually absent. This estimated feature of our model reveals that elites learn from the experiences of relatively proximate countries, and it underpins the ability of our model to explain the spatial clustering of transitions to and from democracy.

To understand the dynamics of democracy adoption, Figure 5 presents our estimates of the evolution of beliefs about the economic effects of democracy. The top panel plots the evolution of the worldwide distribution of the mean of beliefs about $\beta_i^{D=1} - \beta_i^{D=0}$, the difference in long-run GDP growth rates under democracy versus autocracy. The bottom panel shows the evolution of worldwide uncertainty (standard deviation) about these beliefs. The initial state of beliefs in 1951 simply reflects our prior calibration exercise. In the pre-sample period (1875-1950), democracies grew, on average, about 0.4 percentage points faster than autocracies. The median of mean initial beliefs in the top panel of the figure is consistent with this statistic.

While estimated beliefs remain relatively flat for the first three decades of the in-sample period, in the 1980s and, even more dramatically, in the 1990s there is a sharp expansion of

[30] The estimated coefficient is $\gamma = 0.4234$ with a standard error of 0.2292.

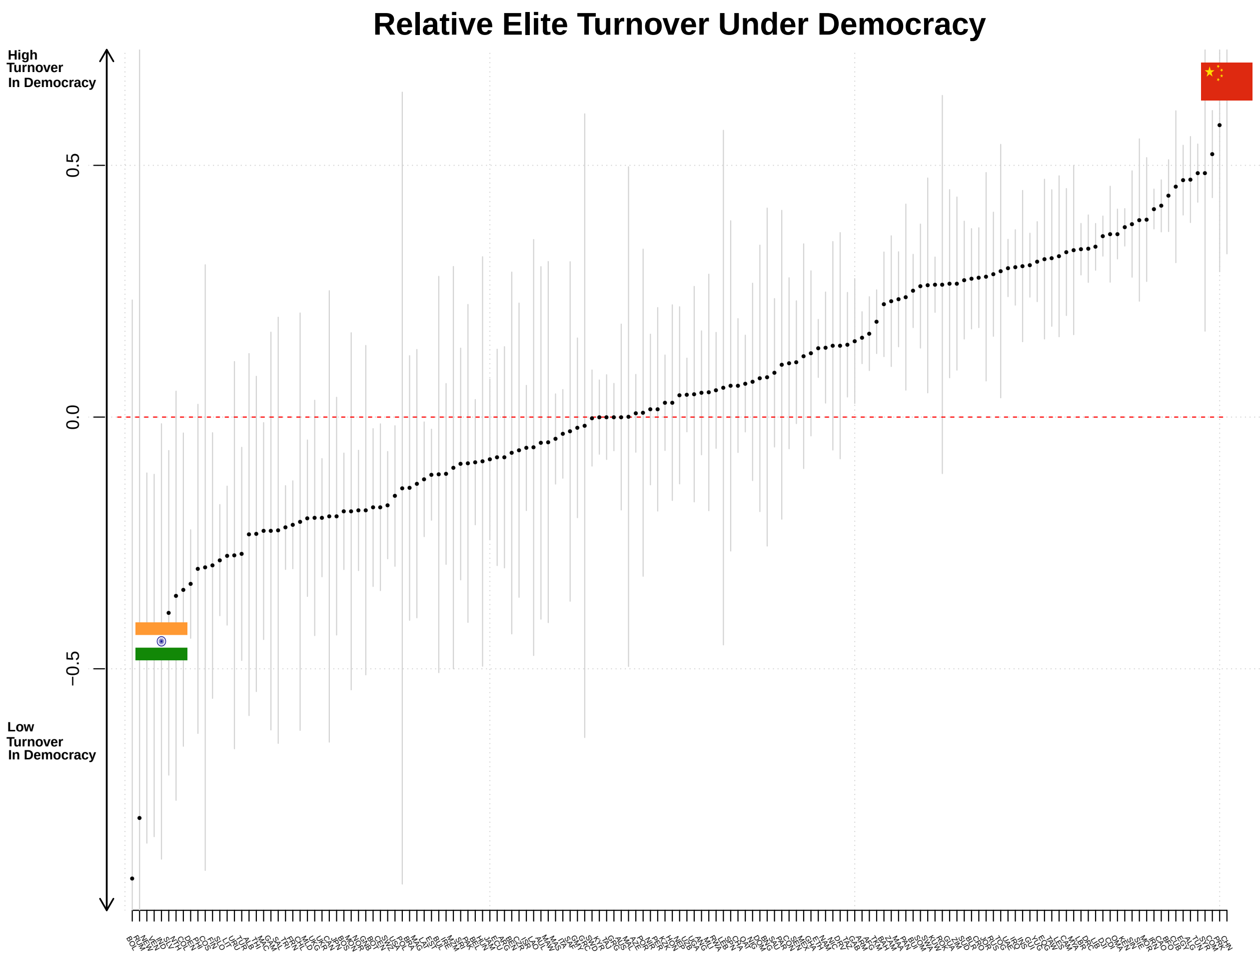

Figure 3: Political Cost of Democracy

Notes. This figure plots our estimates—along with 90% confidence intervals—of the baseline political cost of democracy, f_i, across countries. Positive (negative) values imply higher (lower) elite turnover, all else equal, under democracy than autocracy.

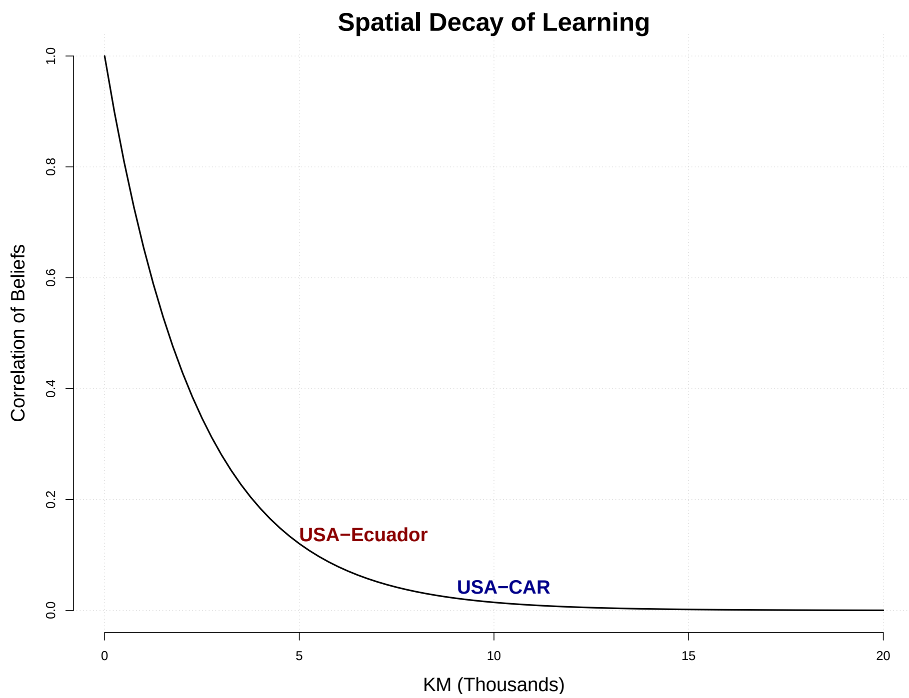

Figure 4: Cross-Country Correlation of Prior Beliefs

Notes. This figure shows the spatial decay of learning, plotting our estimate of the cross-country prior correlation of beliefs about the economic effects of democracy as a function of geographic distance between capitals. For reference, Ecuador and the Central African Republic are approximately five and ten thousand kilometers, respectively, from the U.S.

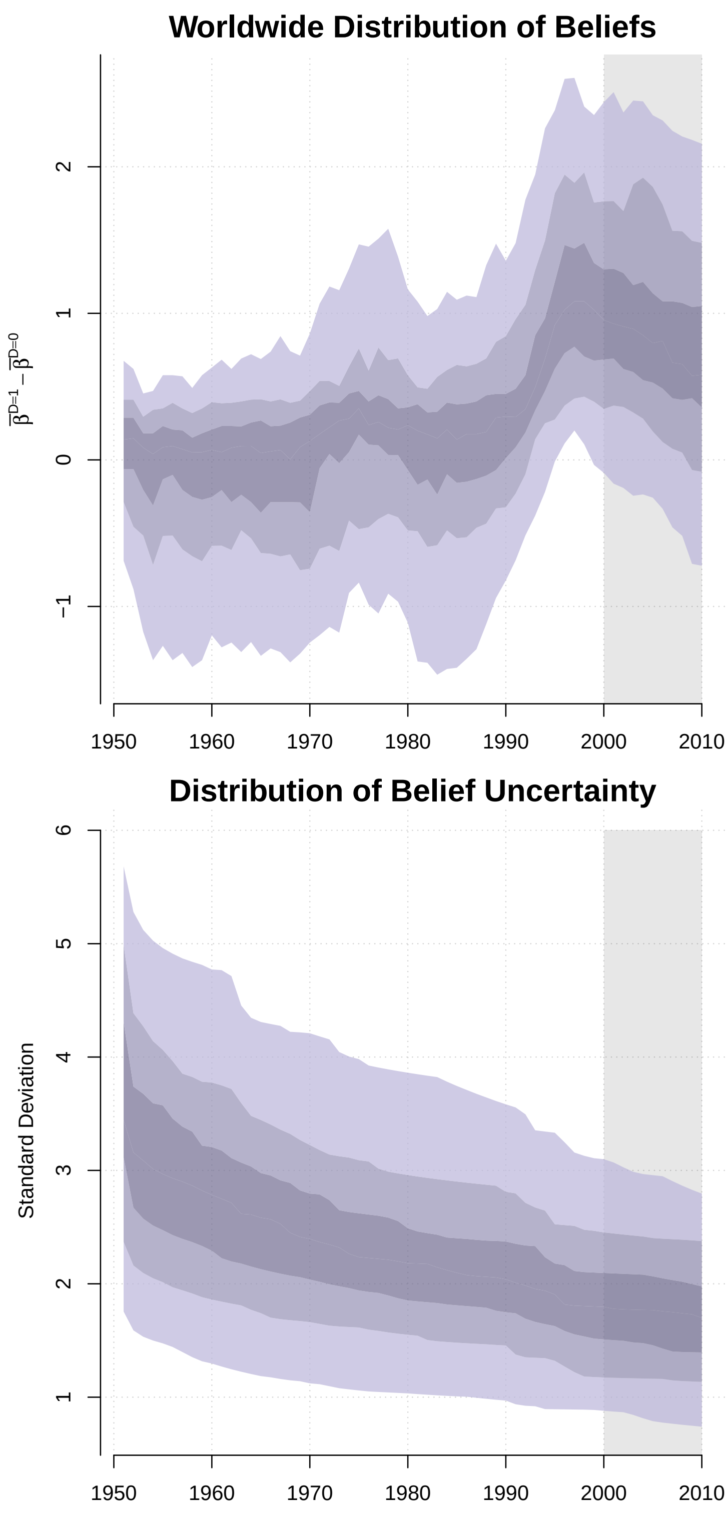

Figure 5: Worldwide Beliefs about Economic Effects of Democracy

Notes. The top panel of this figure plots estimates of the 20th-80th worldwide percentiles of the mean of beliefs about the percentage-point difference in long-run GDP growth rates under democracy versus autocracy. The lower panel shows estimates of the 20th-80th worldwide percentiles of uncertainty (standard deviation) in these beliefs.

beliefs in favor of democracy's superior potential to foster economic growth. This divergence in beliefs reflects an increasingly large gap between autocracies and democracies in observed economic performance. Between 1951-1979, little new information was revealed about the relative ability of democracy to generate growth: on average over this period, autocracies grew just 0.37 percentage points slower than democracies, roughly identical to the observed difference in our pre-sample period. By contrast, between 1980-2000, democracies produced, on average, 1.54 percentage points higher annual growth than autocracies.

This widening gap in economic performance—accelerating through the 1980s—reached a peak in 1987, when democracies outgrew autocracies by 3.3 percentage points. The resulting discrepancy in observed growth rates from countries' prior expectations led them to progressively update their beliefs. Together with our estimates of the (de)stabilizing effects of GDP growth on elite turnover, this change in worldwide beliefs helps explain our model's ability to correctly predict the striking rise in the percentage of world democracies observed in the same period (Figure 1).[31]

This process operated through changes in both the mean and precision of beliefs. After the oil crisis, as democracies came to outperform autocracies, this induced countries to revise up their estimates of democracy's superior economic potential, which drove transitions to democracy. As the first set of countries transitioned, democracy became less rare worldwide, reducing the uncertainty of beliefs about its relative economic merits. This helped solidify the growing consensus, leading to further democratization. The result was the cascade in beliefs between the 1980s and mid-1990s we highlight in Figure 5 and, ultimately, the wave of democratization that occurred over the same period.

[31] In our out-of sample period, we observe a decline in the relative performance of democracies, which grew just 0.67 percentage points faster than autocracies. As expected, we estimate a concomitant decline in average beliefs about the economic benefits of democracy.

4 Counterfactuals

Our model allows us to explore counterfactual experiments of three types. First, it enables us to understand how systemic shocks to prosperity impact the worldwide prospect for democracy. Second, it allows us to ask retrospective questions about historical democratization events. Third, our model allows us to prospectively explore conditions that would lead current elites to transition to or from democracy. We present results of each type in turn.

To conduct these experiments, it is necessary to specify and estimate the "true" data generating process. A considerable advantage of using our baseline no-covariates model to generate these counterfactuals is that only an estimate of the true relationship between GDP growth and democracy is required. Following Buera, Monge-Naranjo and Primiceri (2011), we assume that this relationship is described by a hierarchical linear model similar to (2), which we estimate using all available data between 1875-2000 (excluding the two world wars).[32] This specification is appealing for its flexibility and because it is consistent with our model of learning in that elites with beliefs (2) and prior (3) would eventually learn the truth over time.

4.1 A Second Great Depression

In the year before the market crash of 1929, 51% of the world's independent states were democracies. A year later this proportion dropped to just over 43%. By 1935, the fraction of countries that remained democratic decreased by another 5%, reaching a low of 36% by 1938. In Europe, democratic backsliding was even starker. At its theretofore high in 1920, twenty-six out of twenty-eight European states were democratic, but by 1938 thirteen of these countries had transitioned away from democracy.

[32] Specifically, we assume that

$$y_{i,t} = (1 - D_{i,t})b_i^{D=0} + D_{i,t}b_i^{D=1} + e_{i,t},$$
$$e_t \sim N(0, S \cdot Q \cdot S), \quad b \sim N(\bar{b}, I_2 \otimes S \cdot W \cdot S),$$

where S is a diagonal matrix, $Q_{i,j} = \exp(-Z_{i,j}\zeta_Q)$, and $W_{i,j} = \lambda \exp(-Z_{i,j}\zeta_W)$.

A substantial body of both historical and quantitative research has linked the global decline of democracy in the inter-war period directly to the economic downturn of the Great Depression.[33] Likewise, both in recent academic and popular discourse, the global economic recession of 2008 has been put forward as a contributing factor in the observed wave of recent democratic breakdowns.[34] Next, we provide evidence that these systemic economic crises indeed engender reversals to autocracy and, moreover, highlight how this is driven by changes in beliefs about the economic effects of democracy.

To that end, we simulate two sorts of crises. First, we generate a "short-deep" counterfactual crisis where, for our last in-sample year (2000), we simulate a 5.9% average worldwide contraction of per capita GDP, comparable to the worst year of the Great Depression (1931). In our second crisis, we construct a "long-shallow" counterfactual condition where we perturb growth by a smaller amount—1.7% annually (the average contraction between 1929-1933)—but extend this contraction over a five-year period (2000-2004). We present results where we concentrate these counterfactual conditions in autocratic and democratic countries, respectively. Under the "autocratic bias" condition, recessions are twice as deep in autocracies as in democracies, while keeping the worldwide average contraction consistent with our intervention. Conversely, under the "democratic bias" condition, recessions are twice as deep in democracies.

Model estimates of the worldwide percentage of democracies are shown in the top half of Figure 6. In the left-hand panel, we present estimates from the autocratic bias condition and, in the right-hand panel, estimates from the democratic bias condition. In both plots, the short-deep counterfactual is shown in green, and the long-shallow counterfactual is shown in purple. Note that both the short-deep and long-shallow counterfactual crises negatively impact the worldwide prevalence of democracy. While the large single-period decline in growth has a larger impact in the first year, thereafter the proportion of democracies begins to recover. In comparison, the smaller but longer crisis has a larger overall impact, with

[33] Frey and Weck (1983); De Bromhead, Eichengreen and O'Rourke (2013).
[34] Bartels (2013); Armingeon and Guthmann (2014).

the proportion of democracies continuing to decline through the duration of the economic contraction and recovering more slowly.

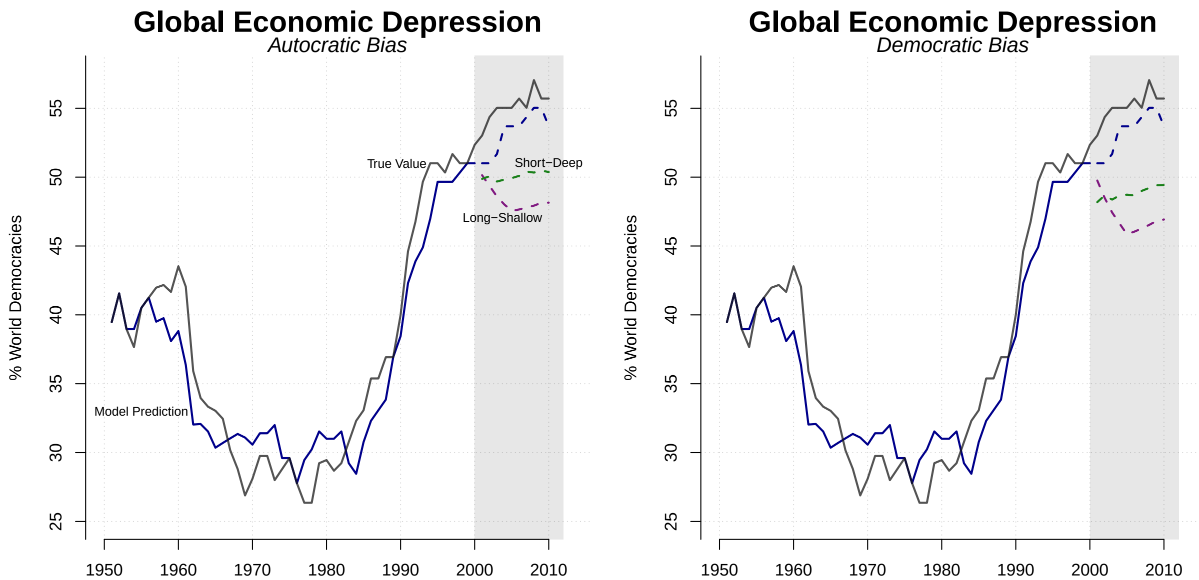

Figure 6: Second Great Depression Counterfactuals

Notes. This figure reports results from two counterfactual exercises. First, we present predictions following a short-deep crisis of a single year (2000) of 6% average worldwide economic contraction (green). Second, we present predictions based on a long-shallow crisis of 5 years (2000-2004) of 2% average contraction (purple). We plot the true percentage of world democracies (gray) as well as baseline model predictions (blue). We concentrate the contraction in autocracies in the left-hand panel, and in democracies in the right-hand panel.

Both patterns are consistent with our actors learning about the economic effects of democracy. In both counterfactuals, the initial shock forces agents to revise their beliefs. However, in the long-shallow counterfactual, as would be expected from a continued process of learning, our agents update their beliefs following each additional negative perturbation of worldwide growth and continue to select out of democracy accordingly. In contrast, in the short-deep scenario, we observe a single large drop in the percentage of democracies. Since after the initial shock to growth there is no "new" information revealed to our agents, there is little additional updating and the worldwide percentage of democracies starts to recover. Importantly, comparing effects across the autocratic and democratic bias conditions, it is clear that the reduction in world democracies is larger when the economic contraction is concentrated among democracies. This is consistent with the evolution of beliefs prior to our intervention, as discussed in Section 3. When democracies perform poorly, counter to the prevailing consensus, elites sharply revise their beliefs and select out of democracy.

4.2 The Third Wave

A number of studies highlight the influence of external actors on the prospects for democracy.[35] Especially for the early "third wave" democratization events in Greece, Portugal, and Spain, the potential for accession to the European Community (EC) has been put forward as a crucial determinant of their respective transitions.[36] Brussels' requirement that community members maintain a form of government consistent with liberal democracy, coupled with the economic benefits of access to the common market, generated an incentive to democratize. In this section, we show that, rather than serving as an institutional target, much of the impact the EC had upon third wave democratization operated through its constituent states' economic performance, which affected beliefs about the economic effects of democracy in potential member states.

To show this, we construct a counterfactual wherein we generate a recession in the EC's three largest economies, Britain, France, and Germany, of 2% average annual contraction for the two years preceding the first transition of Greece in 1974.[37] Then, to obtain an estimate of this counterfactual recession's impact on the transitions of Greece, Portugal, and Spain, we compare our predicted transitions in this counterfactual world to the truth. Our results are given in Figure 7.

For Portugal and Spain, the effect of lower growth in Britain, France, and Germany is considerable, delaying their transitions to democracy by almost 20 years. In contrast, we find no effect of this recession on Greece's transition. The reason for this becomes apparent once we compare the evolution of beliefs in our three example countries. We plot in Figure 8, for each of these cases, our estimates of beliefs under the observed economic conditions

[35] See, for example, Pevehouse (2005).
[36] Whitehead (1996); Powell (1996)
[37] Recessions in our sample last 2 years on average, with an average 2% drop in per capita GDP.

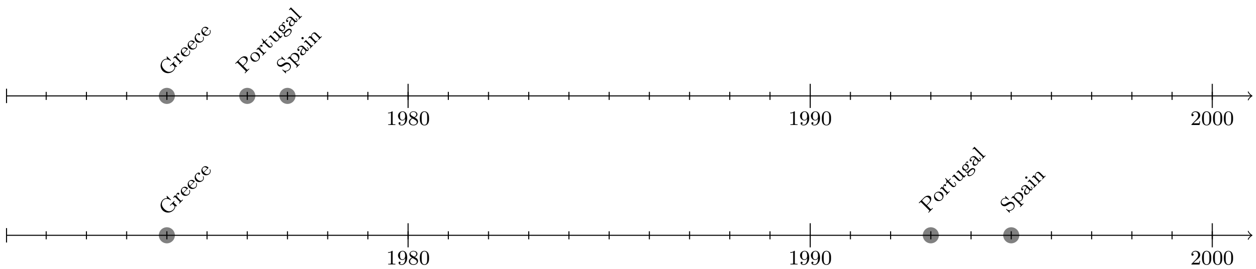

Figure 7: Counterfactual Third Wave

Notes. The top panel shows the true timeline of democratization of Greece, Portugal, and Spain. In the bottom panel, we present this timeline under our counterfactual scenario of a two-year recession (1972 and 1973) in Britain, France, and Germany of 2% average annual contraction.

(solid) and under our counterfactual timeline (dashed). Note that for Portugal and Spain there is a substantial divergence in beliefs between the observed and counterfactual timelines. Following our counterfactual recession, Portuguese and Spanish beliefs become markedly less favorable towards the potential for growth under democracy. In contrast, for Greece this is not the case: there is no substantial difference in their beliefs. The reason for this, according to our model, is that Greek elites pay little attention to the large, relatively distant Western economies we used to construct our counterfactual scenario. Britain, France, and Germany are simply too different from Greece to be used as a reference for learning.

Ultimately, our counterfactual experiment suggests that the distinguishing characteristic of Portugal and Spain in contrast to, for example, Brazil or Mexico is not their underlying propensity for democracy. In terms of their estimated baseline political cost of democracy, Portugal and Brazil and Spain and Mexico are fairly close.[38] Rather, Portugal and Spain democratized early because they learned from Western Europe, benefiting from proximity to successful liberal democracies.

[38]The baseline political cost of democracy for Portugal and Brazil is estimated at 0.01 and -0.14, and in terms of rank order they are 71 and 39, respectively. Estimates for Spain and Mexico are 0.06 and 0.12, yielding a rank order of 83 and 93, respectively.

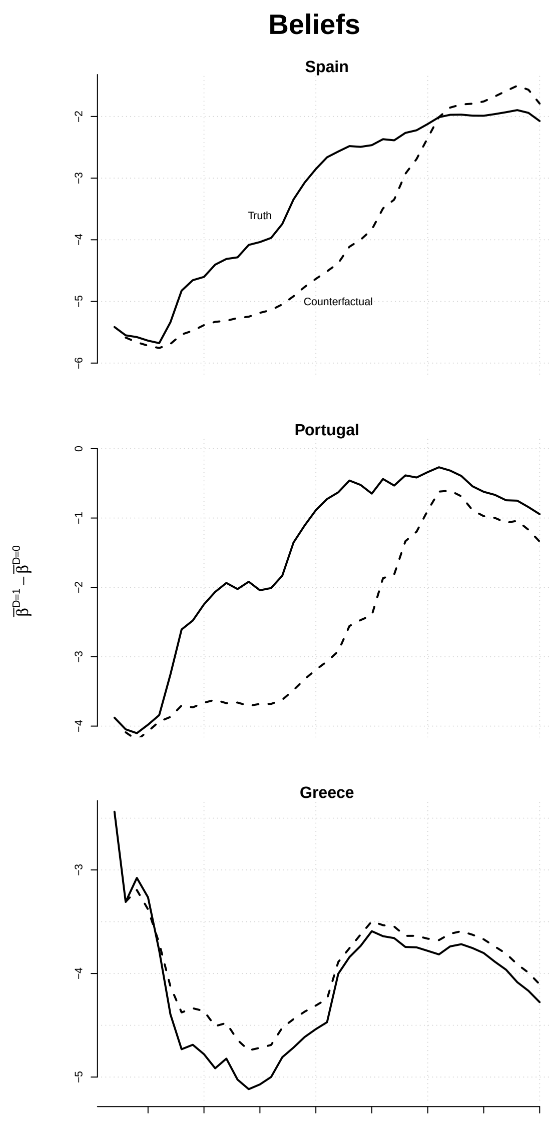

Figure 8: Counterfactual Third Wave Beliefs

Notes. We plot estimates of Spain, Portugal, and Greece's mean beliefs under the true growth rates (solid) and under a counterfactual two-year recession (1973 and 1974) in Britain, France, and Germany of 2% average contraction (dashed).

4.3 Chinese (and North Korean) Democracy

With an eye to contemporary politics, we evaluate the stability of a pair of geopolitically important autocracies. We explore conditions under which our model predicts that China and North Korea would democratize. We focus first on the Chinese case. Here, we look for the minimal average growth rate among China's democratic neighbors, in a five-year economic expansion, that would result in transitions of two types.[39] First, we find the rate of growth that would result in at least a single year of democracy. Second, we establish the growth rate that would deliver a permanent transition to democracy.

To obtain a predicted single-year Chinese transition to democracy, we estimate that it would take five years (2000-2004) of 6.5% average annual growth in China's democratic neighbors. After this single year of democracy (2005), our model predicts an immediate reversal to autocracy. To obtain a "permanent" transition—that is, a prediction of democracy until the end of our sample—we estimate that China's democratic neighbors would have to grow at an average annual rate of 11% between 2000-2004. In contrast, under the same set of counterfactual conditions, North Korea does not democratize for any period. North Korea would democratize for a single year following five years of 16.5% average growth in its democratic neighbors, and permanently following five years of 20.5% average growth.

To highlight the differences in learning, in Figure 9 we plot Chinese and North Korean beliefs under the "Chinese democracy" scenarios of 6.5% and 11% average growth in their democratic neighbors. While our intervention increases the perception that democracy outperforms autocracy in both countries, the change in beliefs in North Korea is too small to induce a transition. These results suggest that the prospects for Chinese and North Korean democracy are limited. It would take a large economic boom in Asian democracies for China to democratize, and an even larger implausible expansion to generate the same outcome in North Korea.

[39] Economic expansions in our sample last five years on average.

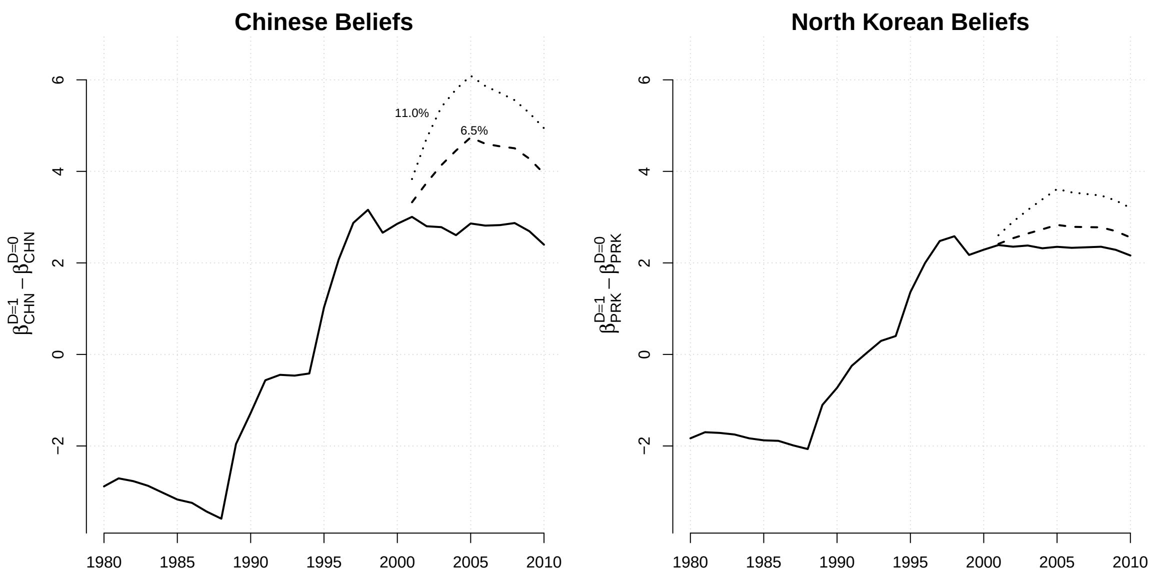

Figure 9: Chinese and North Korean Beliefs

Notes. We plot estimates of Chinese and North Korean beliefs about the economic effects of democracy under observed growth (solid) and under two counterfactual five-year expansions (2000-2004) in their democratic neighbors of 6.5% (dashed) and 11% (dots) average annual growth.

5 Conclusion

In this paper, we propose and estimate a learning model of democratization. We show that worldwide learning about the economic consequences of democracy is crucial to explaining observed patterns of democracy adoption. In particular, our model jointly rationalizes the cross-sectional correlation between income and democracy and the clustering of transitions to and from democracy.

Our parameter estimates indicate that, for office-holding elites, growth is stabilizing in democracies and destabilizing in autocracies. As such, autocracy entails an incentive to suppress economic growth, and democracy the converse. When democratic incumbents come to believe that democracy does not produce sufficient growth to win a fair election, they have an incentive to subvert popular rule. Conversely, autocratic incumbents may liberalize if they expect sufficiently fast subsequent economic development.

Furthermore, we show that learning is highly circumscribed geographically. Our model indicates that countries only update their beliefs using their most proximate neighbors as references. Learning dissipates rapidly outside of a 5,000 kilometer radius. As a consequence, we provide a structural explanation for the regional clustering of democratic transitions.

In combination, these features allow us to successfully predict, both in (1951-2000) and out of sample (2001-2010), much of the observed variation in democracy adoption. Moreover, our approach allows to evaluate the worldwide prospect for democracy. Rather than an "end of history," we show that democracy is only as resilient as the economic performance it engenders. Even in the 1990s, when the success of democratic systems made such proclamations seem reasonable, we find substantial variation in the worldwide distribution of beliefs about democracy's impact on growth and, notably, substantial uncertainty in these beliefs. As recent history suggests, and as our counterfactual experiments demonstrate, systemic economic crises, particularly those concentrated in democracies, have the potential to generate waves of autocratic reversals.

Finally, our paper contributes to a substantial literature on democracy and development. Our results indicate that focusing on the within-country impact of economic development on democratization is insufficient. Rather, modernization should be understood as a systemic phenomenon. Because countries learn from each other, a single country's economic performance affects not only its own likelihood of transitioning to or from democracy but also the prospect for democracy outside of its borders through its influence on the beliefs of neighboring agents.

References

Acemoglu, Daron, Simon Johnson, James A Robinson and Pierre Yared. 2008. "Income and democracy." *The American Economic Review* 98(3):808–842.

Acemoglu, Daron, Simon Johnson, James A Robinson and Pierre Yared. 2009. "Reevaluating the modernization hypothesis." *Journal of Monetary Economics* 56(8):1043–1058.

Acemoglu, Daron, Suresh Naidu, Pascual Restrepo and James A. Robinson. 2014. "Democracy Does Cause Growth." NBER Working Paper No. 20004.

Ahlquist, John S and Erik Wibbels. 2012. "Riding the Wave: World Trade and Factor-Based Models of Democratization." *American Journal of Political Science* 56(2):447–464.

Alesina, Alberto and Dani Rodrik. 1994. "Distributive Politics and Economic Growth." *Quarterly Journal of Economics* 109(2):465–490.

Alesina, Alberto, Nouriel Roubini and Gerald D Cohen. 1997. *Political cycles and the macroeconomy.* MIT press.

Almond, Gabriel Abraham and Sidney Verba. 1963. *The civic culture: Political attitudes and democracy in five nations.* Princeton University Press.

Armingeon, Klaus and Kai Guthmann. 2014. "Democracy in crisis? The declining support for national democracy in European countries, 2007–2011." *European Journal of Political Research* 53(3):423–442.

Ascencio, Sergio and Miguel R. Rueda. 2018. "Poll Watchers, Polling Stations, and Electoral Manipulation." Conditionally accepted at *American Political Science Review.*

Barro, Robert J. 1996. "Democracy and growth." *Journal of economic growth* 1(1):1–27.

Barro, Robert J. 1999. "Determinants of democracy." *Journal of Political economy* 107(S6):S158–S183.

Bartels, Larry M. 2013. "Political effects of the Great Recession." *The Annals of the American Academy of Political and Social Science* 650(1):47–76.

Boix, Carles. 2011. "Democracy, development, and the international system." *American Political Science Review* 105(04):809–828.

Boix, Carles, Michael Miller and Sebastian Rosato. 2013. "A complete data set of political regimes, 1800–2007." *Comparative Political Studies* 46(12):1523–1554.

Boix, Carles and Susan C Stokes. 2003. "Endogenous democratization." *World politics* 55(04):517–549.

Bolton, Patrick and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica* 67(2):349–374.

Bramoulleé, Yann, Rachel Kranton and Martin D'Amours. 2014. "Strategic Interaction and Networks." *American Economic Review* 104(3):898–930.

Brender, Adi and Allan Drazen. 2008. "How do budget deficits and economic growth affect re-election prospects? Evidence from a large panel of countries." *The American Economic Review* 98(5):2203–2220.

Brinks, Daniel and Michael Coppedge. 2006. "Diffusion is no illusion: Neighbor emulation in the third wave of democracy." *Comparative Political Studies* 39(4):463–489.

Buera, Francisco J., Alexander Monge-Naranjo and Giorgio E. Primiceri. 2011. "Learning the Wealth of Nations." *Econometrica* 79(1):1–45.

Chen, Jie and Chunlong Lu. 2011. "Democratization and the middle class in China: The middle class?s attitudes toward democracy." *Political Research Quarterly* 64(3):705–719.

Conley, Timothy G. and Christopher R. Udry. 2010. "Learning about a New Technology: Pineapple in Ghana." *American Economic Review* 100(1):35–69.

Crisman-Cox, Casey and Michael Gibilisco. 2018. "Audience Costs and the Dynamics of War and Peace." *American Journal of Political Science* 62(3):566–580.

Dahl, Robert Alan. 1998. *On Democracy.* Yale University Press.

Davison, Anthony C. and David V. Hinkley. 1997. Bootstrap Methods and their Application. In *Cambridge Series in Statistical and Probabilistic Mathematics.* New York: Cambridge University Press.

De Bromhead, Alan, Barry Eichengreen and Kevin H O'Rourke. 2013. "Political extremism in the 1920s and 1930s: Do German lessons generalize?" *The Journal of Economic History* 73(2):371–406.

Diamond, Larry. 2011. "Why democracies survive." *Journal of Democracy* 22(1):17–30.

Dubé, Jean-Pierre, Jeremy T. Fox and Che-Lin Su. 2012. "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation." *Econometrica* 80(5):2231–2267.

Feenstra, Robert C, Robert Inklaar and Marcel P Timmer. 2015. "The next generation of the Penn World Table." *The American Economic Review* 105(10):3150–3182.

Foster, Andrew D. and Mark R. Rosenzweig. 1995. "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture." *Journal of Political Economy* 103(6):1176–1209.

Frey, Bruno S and Hannelore Weck. 1983. "A statistical study of the effect of the Great Depression on elections: The Weimar Republic, 1930–1933." *Political Behavior* 5(4):403–420.

García-Jimeno, Camilo. 2016. "The Political Economy of Moral Conflict: An Empirical Study of Learning and Law Enforcement under Prohibition." *Econometrica* 84(2):511–570.

Gehlbach, Scott, Konstantin Sonin and Milan W. Svolik. 2016. "Formal Model of Nondemocratic Politics." *Annual Review of Political Science* 19:565–584.

Gleditsch, Kristian Skrede. 2002. "Expanded Trade and GDP Data (version 6.0)." *Journal of Conflict Resolution* 46(5):712–724.

Gleditsch, Kristian Skrede and Michael D Ward. 2006. "Diffusion and the international context of democratization." *International organization* 60(04):911–933.

Goemans, Henk E, Kristian Skrede Gleditsch and Giacomo Chiozza. 2009. "Introducing Archigos: A dataset of political leaders." *Journal of Peace research* 46(2):269–283.

Heiss, Florian and Viktor Winschel. 2008. "Likelihood approximation by numerical integration on sparse grids." *Journal of Econometrics* 144(1):62–80.

Hibbs, Douglas A. 1977. "Political parties and macroeconomic policy." *American political science review* 71(4):1467–1487.

Hirschman, Albert O and Michael Rothschild. 1973. "The Changing Tolerance for Income Inequality in the Course of Economic Development* With A Mathematical Appendix." *The Quarterly Journal of Economics* 87(4):544–566.

Houle, Christian, Mark A Kayser and Jun Xiang. 2016. "Diffusion or Confusion? Clustered Shocks and the Conditional Diffusion of Democracy." *International Organization* 70(4):687.

Huntington, Samuel P. 1968. *Political order in changing societies.* Yale University Press.

Huntington, Samuel P. 1993. *The third wave: Democratization in the late twentieth century.* Vol. 4 University of Oklahoma press.

Kadera, Kelly M, Mark JC Crescenzi and Megan L Shannon. 2003. "Democratic survival, peace, and war in the international system." *American Journal of Political Science* 47(2):234–247.

Lipset, Seymour Martin. 1959. "Some social requisites of democracy: Economic development and political legitimacy." *American political science review* 53(01):69–105.

Londregan, John B and Keith T Poole. 1996. "Does high income promote democracy?" *World politics* 49(1):1–30.

Maddison, Angus. 2010. "Statistics on world population, GDP and per capita GDP, 1-2008 AD." *Historical Statistics* pp. 1–36.

Mansfield, Edward D, Helen V Milner and B Peter Rosendorff. 2000. "Free to trade: Democracies, autocracies, and international trade." *American Political Science Review* 94(02):305–321.

Matérn, Bertil. 1960. "Spatial variation: Stochastic models and their application to some problems in forest surveys and other sampling investigations." *Meddelanden från Statens Skogsforskningsinstitut* 49(5):1–144.

Miller, Michael K. 2016. "Democracy by Example? Why Democracy Spreads When the World's Democracies Prosper." *Comparative Politics* 49(1):83–116.

Mossel, Elchanan, Allan Sly and Omer Tamuz. 2015. "Strategic Learning and the Topology of Social Networks." *Econometrica* 83(5):1755–1794.

Norris, Pippa. 1999. *Critical citizens: Global support for democratic government.* Oxford University Press.

Olson, Mancur. 1963. "Rapid growth as a destabilizing force." *The Journal of Economic History* 23(4):529–552.

Persson, Torsten and Guido Tabellini. 1994. "Is Inequality Harmful for Growth?" *American Economic Review* 84(3):600–621.

Pevehouse, Jon C. 2005. *Democracy from above: Regional organizations and democratization.* Cambridge University Press.

Powell, Charles. 1996. International aspects of democratization: The case of Spain. In *The international dimensions of democratization: Europe and the Americas*, ed. Laurence Whitehead. Oxford: Oxford University Press chapter 11, pp. 285–314.

Primiceri, Giorgio E. 2006. "Why Inflation Rose and Fell: Policymakers' Beliefs and US Postwar Stabilization Policy." *Quarterly Journal of Economics* 121(3):867–901.

Przeworski, Adam and Fernando Limongi. 1997. "Modernization: Theories and facts." *World politics* 49(02):155–183.

Przeworski, Adam, Michael Alvarez, Jose Cheibub and Fernando Limongi. 2000. "Democracy and Development: Political Institutions and Well-Being in the World, 1950-1990 (Cambridge Studies in the Theory of Democracy).".

Rodrik, Dani. 1999. "Democracies Pay Higher Wages." *Quarterly Journal of Economics* 114(3):707–738.

Skocpol, Theda. 1979. *States and social revolutions: A comparative analysis of France, Russia and China.* Cambridge University Press.

Spolaore, Enrico and Romain Wacziarg. 2009. "The Diffusion of Development." *The Quarterly Journal of Economics* 124(2):469–592.

Su, Che-Lin and Kenneth L Judd. 2012. "Constrained Optimization Approaches to Estimation of Structural Models." *Econometrica* 80(5):2213–2230.

Svolik, Milan W. 2013. "Learning to Love Democracy: Electoral Accountability and the Success of Democracy." *American Journal of Political Science* 57(3):685–702.

Svolik, Milan W. 2019. "Democracy as an equilibrium: rational choice and formal political theory in democratization research." *Democratization* 26(1):40–60.

Torfason, Magnus Thor and Paul Ingram. 2010. "The global rise of democracy: A network account." *American Sociological Review* 75(3):355–377.

Wejnert, Barbara. 2005. "Diffusion, development, and democracy, 1800-1999." *American Sociological Review* 70(1):53–81.

Whitehead, Laurence. 1996. *The international dimensions of democratization: Europe and the Americas.* Oxford University Press.

Online Appendix

A1 Likelihood of the Data

Recall that $W^T \equiv \{I_t, y_t, D_t, X_t\}_{t=1}^T$ denotes the set of all data available up to period T. With a slight abuse of notation—using $\mathcal{L}$ to denote arbitrary densities of the data—the likelihood function can be written as

$$\mathcal{L}(W^T|\boldsymbol{\theta}) = \prod_{t=1}^{T} \mathcal{L}(W_t|W^{t-1}, \boldsymbol{\theta}),$$

where $W_t \equiv \{I_t, y_t, D_t, X_t\}$ collects the data generated in period t. As discussed in Section 2.2, we assume that observed outcomes are only affected by actual choices and not by the beliefs that led to those choices. That is, transitions of power (I_t), GDP growth (y_t), and other economic and political characteristics of countries (X_t) are shaped by realized institutions (D_t), but they are not directly affected by beliefs about the potential effects of transitioning into or out of democracy. Formally, this assumption allows us to write

$$\begin{aligned}\mathcal{L}(W_t|W^{t-1}, \boldsymbol{\theta}) &= \mathcal{L}(I_t, y_t, D_t, X_t|W^{t-1}, \boldsymbol{\theta}) \\ &= \mathcal{L}(I_t|y_t, D_t, X_t, W^{t-1})\mathcal{L}(y_t|D_t, X_t, W^{t-1}) \cdots \\ &\quad \mathcal{L}(D_t|X_t, W^{t-1}, \boldsymbol{\theta})\mathcal{L}(X_t|W^{t-1}),\end{aligned}$$

which implies that $\mathcal{L}(W^T|\boldsymbol{\theta}) \propto \prod_{t=1}^{T} \mathcal{L}(D_t|X_t, W^{t-1}, \boldsymbol{\theta})$.

To compute $\mathcal{L}(D_t|X_t, W^{t-1}, \boldsymbol{\theta})$, notice from (5) that, given $(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$, there is a threshold value of $\kappa_{i,t}$—the realized shock in period t to the political cost of democracy in country i—such that $D_{i,t} = 1$ if and only if $\kappa_{i,t}$ falls below the threshold. This threshold value, denoted $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$, is defined implicitly by

$$E_{i,t-1}\left[\frac{\exp(\alpha_i + \theta^{D=1}(\beta_i^{D=1} + \epsilon_{i,t}) - f_i - X_{i,t}'\xi - \overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta}))}{1 + \exp(\alpha_i + \theta^{D=1}(\beta_i^{D=1} + \epsilon_{i,t}) - f_i - X_{i,t}'\xi - \overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta}))}\right] \\ = E_{i,t-1}\left[\frac{\exp(\alpha_i + \theta^{D=0}(\beta_i^{D=0} + \epsilon_{i,t}))}{1 + \exp(\alpha_i + \theta^{D=0}(\beta_i^{D=0} + \epsilon_{i,t}))}\right]. \tag{A1}$$

Appendix A2 describes an iterative algorithm for computing $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$. Since $\kappa_{i,t}$ is distributed independently across countries, the likelihood can be written as

$$\mathcal{L}(W^T|\boldsymbol{\theta}) \propto \prod_{t=1}^{T}\prod_{i=1}^{n}\mathcal{L}(D_{i,t}|X_{i,t}, W^{t-1}, \boldsymbol{\theta}),$$

where

$$\mathcal{L}(D_{i,t}|X_{i,t}, W^{t-1}, \boldsymbol{\theta}) = \Phi\left(\frac{\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})}{\varsigma_i}\right)^{D_{i,t}}\left[1 - \Phi\left(\frac{\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})}{\varsigma_i}\right)\right]^{1-D_{i,t}}$$

and Φ denotes the standard Normal cumulative distribution function.

A2 Computing $\overline{\kappa}_{i,t}(\cdot)$: A Fixed-Point Algorithm

We provide an algorithm for computing $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ from (A1). To simplify notation, let $V_0(\beta_i^{D=0}, \epsilon_{i,t}) \equiv \alpha_i + \theta^{D=0}(\beta_i^{D=0} + \epsilon_{i,t})$ and $V_1(\beta_i^{D=1}, \epsilon_{i,t}) \equiv \alpha_i + \theta^{D=1}(\beta_i^{D=1} + \epsilon_{i,t}) - f_i - X_{i,t}'\xi$. Define $g : \mathbb{R} \to \mathbb{R}$ by

$$g(\kappa) \equiv E_{i,t-1}\left[\frac{\exp(V_1(\beta_i^{D=1}, \epsilon_{i,t}) - \kappa)}{1 + \exp(V_1(\beta_i^{D=1}, \epsilon_{i,t}) - \kappa)}\right].$$

Differentiating under the integral, we have

$$g'(\kappa) = -E_{i,t-1}\left[\frac{\exp(V_1(\beta_i^{D=1}, \epsilon_{i,t}) - \kappa)}{1 + \exp(V_1(\beta_i^{D=1}, \epsilon_{i,t}) - \kappa)}\left(1 - \frac{\exp(V_1(\beta_i^{D=1}, \epsilon_{i,t}) - \kappa)}{1 + \exp(V_1(\beta_i^{D=1}, \epsilon_{i,t}) - \kappa)}\right)\right] < 0,$$

and $-g'(\kappa) < g(\kappa)$ as $\frac{\exp(V_1(\beta_i^{D=1}, \epsilon_{i,t})-\kappa)}{1+\exp(V_1(\beta_i^{D=1}, \epsilon_{i,t})-\kappa)} \in (0,1)$. Now, define $h : \mathbb{R} \to \mathbb{R}$ by

$$h(\kappa) \equiv \kappa + \log(g(\kappa)) - \log(U_0),$$

where $U_0 \equiv E_{i,t-1}\left[\frac{\exp(V_0(\beta_i^{D=0}, \epsilon_{i,t}))}{1+\exp(V_0(\beta_i^{D=0}, \epsilon_{i,t}))}\right]$. Notice that $h'(\kappa) = 1 + \frac{g'(\kappa)}{g(\kappa)} \in (0,1)$, which implies that, for any $\kappa', \kappa \in \mathbb{R}$,

$$|h(\kappa') - h(\kappa)| = \left|\int_{\kappa}^{\kappa'} h'(k)dk\right| < |\kappa' - \kappa|.$$

Thus, h is a contraction. By the Contraction Mapping Theorem, h has a unique fixed point, which by construction is equal to $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ and can be computed iteratively as follows:

Step 0. Initialize $\kappa_0 = -f_i - X'_{i,t}\xi$ and set $l = 1$.

Step l. Set $\kappa_l = h(\kappa_{l-1})$. If $\kappa_l = \kappa_{l-1}$, stop; otherwise, continue to step $l+1$.

As logarithms are computationally more expensive than exponentials, $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ can be computed more efficiently using an exponential version of the algorithm.[40]

[40] Iteratively evaluate $h_e(w) \equiv w\left(\frac{g_e(w)}{U_0}\right)$, where $g_e(w) \equiv E_{i,t-1}\left[\frac{\exp(V_1(\beta_i^{D=1}, \epsilon_{i,t}))/w)}{1+\exp(V_1(\beta_i^{D=1}, \epsilon_{i,t}))/w}\right]$, and compute $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ as the natural logarithm of the unique fixed point of h_e. The starting point for the algorithm in this case should be constrained to satisfy $g_e(w_0)^2 < U_0$.

A3 Prior

We set our prior over the model parameters in Table 1 as follows. We assume that

$$\begin{aligned}
\alpha_i &\overset{\text{i.i.d.}}{\sim} N(\underline{\alpha}, \omega_\alpha^2), \\
\theta^{D=0,1} &\overset{\text{i.i.d.}}{\sim} N(\underline{\theta}, \omega_\theta^2), \\
\bar{\beta}_{i,0}^{D=0} &\overset{\text{i.i.d.}}{\sim} N(\underline{\beta}_0^{D=0}, \omega_\beta^2), \\
\bar{\beta}_{i,0}^{D=1} &\overset{\text{i.i.d.}}{\sim} N(\underline{\beta}_0^{D=1}, \omega_\beta^2), \\
v_i &\overset{\text{i.i.d.}}{\sim} \text{IG}(s_v, d_v), \\
f_i &\overset{\text{i.i.d.}}{\sim} N(\underline{f}, \omega_f^2), \\
\varsigma_i &\overset{\text{i.i.d.}}{\sim} \text{IG}(s_\varsigma, d_\varsigma), \\
\gamma &\sim \text{Uniform}, \\
\xi &\sim \text{Uniform},
\end{aligned}$$

where $\text{IG}(s, d)$ denotes the Inverse-Gamma distribution with shape parameter s and scale parameter d. We calibrate our prior using pre-sample data from 1875-1950 (excluding the two world wars):

- We set $\underline{\beta}_0^{D=0} = 0.0180$ and $\underline{\beta}_0^{D=1} = 0.0218$, the average annual growth rates among autocracies and democracies, respectively, in the pre-sample period. We then set $\omega_\beta = 0.02$, allowing for considerable uncertainty about the mean of initial beliefs.

- We select $s_v = 3$ and $d_v = 0.7423$ so that the prior mean and standard deviation of $v_i\sigma_i$ equal the standard deviation of average growth rates, $\bar{y}_i$, in the pre-sample period. A pre-sample estimate of the mean of σ_i (equal to 0.0531) is obtained from the residuals of a regression of GDP growth on country and time fixed effects. We then set the prior mean of v_i equal to $\sqrt{\text{Var}(\bar{y}_i)}/0.0531 = \sqrt{0.0004}/0.0531 = 0.3711$.

- We set $\underline{\theta} = 0$ to adopt an agnostic starting point about whether GDP growth has a stabilizing or destabilizing effect on elite turnover across forms of government, and we normalize $\omega_\theta = 1$.

- To adopt an agnostic starting point regarding the political cost of democracy, we set $\underline{f} = 0$. We describe our choice of ω_f below.

- To ensure prior correlations of beliefs between 0 and 1, we adopt a flat (improper) prior over $\gamma \geq 0$. For the political cost of democracy, we center the variables in $X_{i,t}$ around their sample means so that $K_{i,t}$ has an expected value of zero (in line with our agnostic view of f_i), and we adopt a flat (improper) prior over ξ.

- Letting $\overline{\theta y}_i$ and $\overline{KD}_i$ denote the within-country means of $\theta^{D=D_{i,t}} y_{i,t}$ and $K_{i,t} D_{i,t}$, respectively, and noting that $\alpha_i + \overline{\theta y}_i - \overline{DK}_i$ approximately equals the log-odds of staying in power in country i, we select $\underline{\alpha}$, ω_α, and ω_f to match the first two moments of these log-odds across countries in the pre-sample period. Since $E(\alpha_i + \overline{\theta y}_i - \overline{DK}_i) = \underline{\alpha}$, we set $\underline{\alpha} = 1.8432$, the average log-odds in the pre-sample period.[41] Noting that the variance of the log-odds among autocracies is approximately equal to $\omega_\alpha^2 + \mathrm{Var}(\bar{y}_i)$, while the variance among democracies is approximately equal to $\omega_\alpha^2 + \mathrm{Var}(\bar{y}_i) + \omega_f^2$, we set $\omega_\alpha^2 + 0.0005 = 0.722$ and $\omega_\alpha^2 + 0.0007 + \omega_f^2 = 1.0802$, so $\omega_\alpha = 0.8494$ and $\omega_f = 0.5984$.

- Finally, to discourage the model from fitting the data with large (absolute) realizations of the unobserved political cost shock $\kappa_{i,t}$, we set $s_\varsigma = 3$ and $d_\varsigma = 0.2992$ so that ς_i has a prior mean and standard deviation of $\omega_f/4 = 0.1496$.

A4 Maximum-A-Posteriori Estimator: MPEC Approach

As discussed in Section 2.2—and as can be appreciated from Appendix A2—calculating $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ from (A1) to evaluate the likelihood of the data is computationally expensive. To avoid this burden, we follow the Mathematical Programming with Equilibrium

[41] For countries that experienced no elite turnover in the pre-sample period, we limit the probability of staying in power to equal the maximum among countries with turnover (0.95).

Constraints (MPEC) approach of Su and Judd (2012) to compute our maximum-a-posteriori (MAP) estimator of $\boldsymbol{\theta}$. The idea behind this approach is simple: instead of calculating $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ at every trial value of $\boldsymbol{\theta}$, treat $\overline{\kappa}_{i,t}$ as an auxiliary parameter and impose (A1)—the optimality (or equilibrium) condition of the model—as a feasibility constraint on the log-posterior maximization program. Accordingly, we estimate $\boldsymbol{\theta}$ by solving $\max_{\boldsymbol{\theta},\overline{\boldsymbol{\kappa}}} \log\left(p(\boldsymbol{\theta}, \overline{\boldsymbol{\kappa}}|W^T)\right)$ subject to $\overline{\kappa}_{i,t}$ satisfies (A1) for all i and t.

As shown by Su and Judd (2012), MPEC and the standard approach of directly maximizing $\log\left(p(\boldsymbol{\theta}|W^T)\right)$ yield theoretically-identical estimates of $\boldsymbol{\theta}$. Computationally, MPEC's advantage arises from the fact that modern optimization algorithms do not enforce constraints until the final iteration of the search process. Thus, the computationally expensive condition (A1) is satisfied exactly once rather than at every trial value of $\boldsymbol{\theta}$. Moreover, for this reason, MPEC is robust to sensitivity issues that may arise from not setting a sufficiently stringent convergence criterion when computing $\overline{\kappa}_{i,t}(X_{i,t}, W^{t-1}, \boldsymbol{\theta})$ (Dubé, Fox and Su, 2012). A potential disadvantage is that, by introducing $\overline{\boldsymbol{\kappa}}$ as additional parameters, MPEC increases the size of the optimization problem. However, this concern is mitigated by the sparsity that results from each auxiliary parameter $\overline{\kappa}_{i,t}$ entering a single constraint.

To reap the computational benefits of the MPEC approach, it is essential to employ optimization software tailor-made to handle large-scale problems—with thousands of variables and nonlinear constraints. Accordingly, we implement our MPEC-MAP estimator using the industry-leading software Knitro.[42] Due to memory and computational constraints—our baseline model with no covariates features 8,459 variables—our implementation relies on Knitro's Interior/Direct algorithm with their limited-memory quasi-Newton BFGS approximation of the Hessian of the Lagrangian. Nevertheless, we provide exact first derivatives of the log-posterior and constraints.[43] With a 3.0 GHz machine, it takes about 5-6 days to estimate our model once.

[42]https://www.artelys.com/en/optimization-tools/knitro

[43]Knitro offers a derivative-check option—which our implementation passes—to test the code for exact derivatives against finite-difference approximations.

To mitigate concerns about potential local maxima, for each model specification we randomly draw 5 sets of starting values for the optimization algorithm from the prior distribution of the model parameters described in Appendix A3. We then select, for each specification, the solution that achieves the highest log-posterior value. Reassuringly, there is very little divergence in solutions across starting values.

Standard errors for our parameter estimates are parametrically bootstrapped (Davison and Hinkley, 1997). Due to the considerable computational cost of estimating our model, we only compute standard errors for our baseline specification with no covariates. This has the added advantage that only estimates from the true DGP described in Footnote 32 are necessary to generate bootstrap samples.

A final notable computational challenge is that the integrals in (A1) have no closed-form solution. Rather than employing a Monte Carlo approximation, which would require independent draws across all i and t to prevent simulation error from propagating, we rely on sparse-grid integration as implemented by Heiss and Winschel (2008). This approach is much more efficient and delivers virtually exact integral computations for integrands that are well approximated by polynomials—as is the case for the integrals in (A1).[44]

A5 Additional Results

In this appendix, we describe in detail various additional results mentioned in the paper.

Model fit by world region. In Figure A1, we present goodness-of-fit and out-of-sample prediction results disaggregated by four regions of the world: the Americas, Europe, Africa, and Asia-Oceania. As in Figure 1, we again compare the true proportion of democracies (gray) in each region to our learning model's predictions, both in (solid blue) and out of sample (dashed blue). The top panel presents results from our baseline model with no covariates; the bottom panel, from our model with two covariates. Notably, both specifica-

[44]Our implementation computes exact integrals of fifteenth-degree polynomials.

tions perform well at this, or indeed any, level of geographic aggregation. And, while not included in Figure A1 to avoid clutter, our learning model still significantly outperforms any alternative that ignores the role of learning.

Direct diffusion of democracy. As discussed in Section 3, to evaluate the possibility that our model's success is simply an artifact of some alternative process of democratic diffusion that is indirectly picked up by our model's spatial and temporal flexibility, we construct a distance-weighted measure of how democratic each country's neighborhood is over time, and we reestimate our baseline model using this measure as a control for direct diffusion effects on the political cost of democracy. Specifically, we include in $X_{i,t}$ the weighted average

$$\bar{D}_{i,t-1} = \frac{\sum_{j \neq i} \exp(-\delta d_{i,j}) D_{j,t-1}}{\sum_{j \neq i} \exp(-\delta di, j)},$$

where $d_{i,j}$ denotes the distance between i and j's capitals.

To reduce the computational burden, instead of estimating δ—the parameter determining the size of each country's effective neighborhood—we consider five scenarios. For our "medium neighborhood" scenario, we set δ equal to the estimated value of the parameter governing the spatial decay of learning in our baseline specification, described in Section 3.1. We also consider a "smaller" and "smallest" neighborhood scenario, where we increase the value of δ two-fold and five-fold, respectively, and a "larger" and "largest" neighborhood scenario, where we divide δ by two and five, respectively. Table A1 presents the results of this exercise, following the format of Table 2. We find that controlling for direct diffusion provides little increase in predictive power, which suggests that it is our proposed mechanism of learning about the economic effects of democracy—and not some alternative process of diffusion—driving our results.

Robustness to alternative measure of democracy and timeframe. Finally, we explore the sensitivity of our results to (i) our preferred measure of democracy and (ii) our

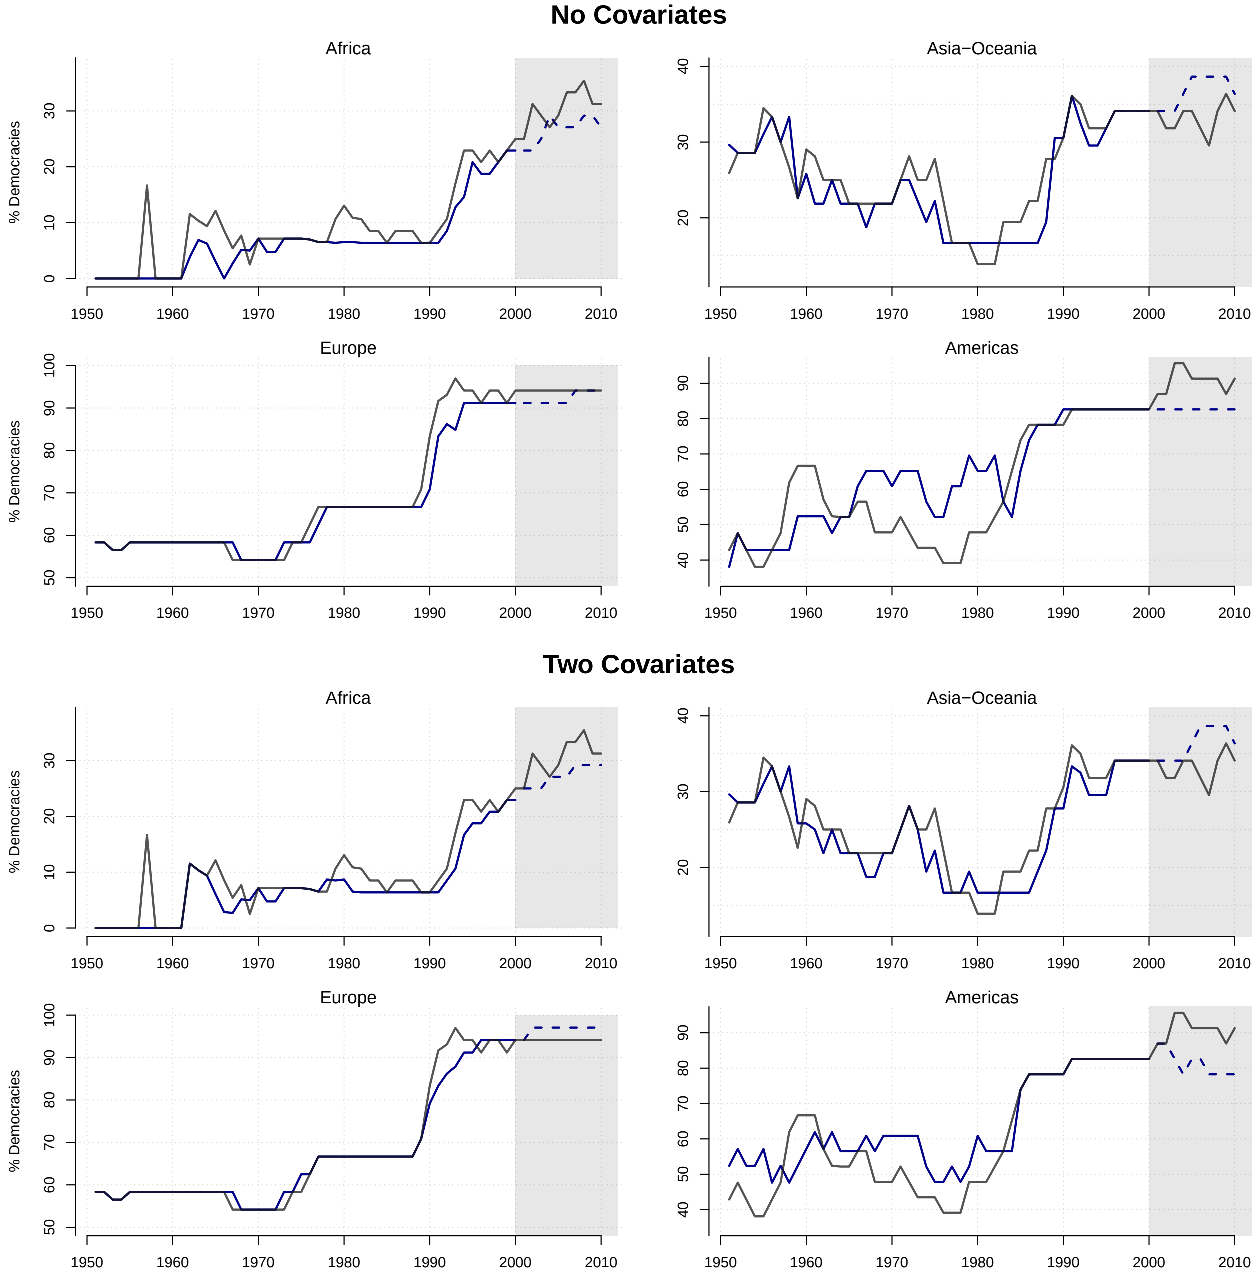

Figure A1: Observed versus Predicted Prevalence of Democracy by World Region

Notes. By world region, this figure compares the true proportion of democracies (gray) to estimates generated by our learning model for both the in-sample (solid blue) and out-of-sample (dashed blue) periods. The top panel presents results using our baseline specification with no covariates. In the lower panel, we control for lagged log-GDP per capita and incumbents' time in power.

Table A1: Direct Diffusion of Democracy

	Smallest Neighborhood		Smaller Neighborhood		Medium Neighborhood		Larger Neighborhood		Largest Neighborhood	
	Learning	No Learning	Learning	No Learning	Learning	No Learning	Learning	No Learning	Learning	No Learning
Choices (% correct)	96.3	92.0	96.0	92.5	96.3	92.3	95.9	92.1	96.7	91.9
Transitions (% correct)										
± 0 years	11.6	5.4	12.4	7.0	13.2	8.5	20.2	7.0	18.6	4.7
± 2 years	41.9	18.6	54.3	25.6	54.3	27.1	50.4	26.4	56.6	22.5
Log-likelihood	-536.6	-1,019.1	-539.7	-964.3	-511.4	-931.7	-553.6	-939.3	-481.5	-946.2
Observations	5,925	5,925	5,925	5,925	5,925	5,925	5,925	5,925	5,925	5,925

Notes. From left to right, respectively, models in each "neighborhood" scenario control for direct diffusion effects on the political cost of democracy using distance weights $\delta = 2.5$, $\delta = 1$, $\delta = 0.5$, $\delta = 0.25$, and $\delta = 0.1$. For each model, we report the percentage of correctly predicted in-sample system of government choices (first row). We similarly report the percentage of correctly predicted transitions to or from democracy within a 0-year window (second row) and a 5-year window (third row) of the event.

chosen timeframe.

First, as noted in Section 2.2, various alternative measures of democracy have been employed in the literature. To test the robustness of our results to this feature of the data, we take Acemoglu et al.'s (2014) preferred measure and reestimate our baseline model.[45] Since the BMR coding is more comprehensive than this alternative measure, to keep the results as comparable as possible—in particular, to avoid having to modify the time span of the sample—we use the BMR coding to fill any gaps in Acemoglu et al.'s (2014) data.

Second, to address potential concerns about the myopia of incumbents in our model and whether an annual timeframe is appropriate to study changes in systems of government, we estimate a five-year panel version of our baseline model (and true DGP). Following Acemoglu et al. (2008), we take the observation of democracy every fifth year as our measure for the five-year panel.

Table A2 summarizes the results of these robustness exercises, following the format of Table 2. Our findings are virtually unchanged.

[45] To that end, we also reestimate the true DGP (see Footnote 32) to obtain a new estimate of Σ.

Table A2: Robustness to Democracy Measure and Timeframe

	Alternative Democracy Measure		Five-Year Panel	
	Learning	No Learning	Learning	No Learning
Choices (% correct)	95.4	87.5	95.5	87.3
Transitions (% correct)				
±0 years	12.5	0.0		
±2 years	46.1	0.0	56.8	0.0
Log-likelihood	-647.9	-1,555.4	-160.37	-388.1
Observations	5,925	5,925	1,139	1,139

Notes. Models in the second and third columns are estimated using Acemoglu et al.'s (2014) preferred measure of democracy with no covariates. Models in the last two columns are estimated using a five-year panel version of our data with no covariates. For each model, we report the percentage of correctly predicted in-sample system of government choices (first row) and the percentage of correctly predicted transitions to or from democracy within a 0-year window (second row) and a 5-year window (third row) of the event.