

DISCUSSION PAPER SERIES

DP13763

**THE ROLE OF ELECTORAL INCENTIVES
FOR POLICY INNOVATION: EVIDENCE
FROM THE U.S. WELFARE REFORM**

Andreas Bernecker, Pierre Boyer and Christina
Gathmann

PUBLIC ECONOMICS

THE ROLE OF ELECTORAL INCENTIVES FOR POLICY INNOVATION: EVIDENCE FROM THE U.S. WELFARE REFORM

Andreas Bernecker, Pierre Boyer and Christina Gathmann

Discussion Paper DP13763

Published 28 May 2019

Submitted 27 May 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **PUBLIC ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Andreas Bernecker, Pierre Boyer and Christina Gathmann

THE ROLE OF ELECTORAL INCENTIVES FOR POLICY INNOVATION: EVIDENCE FROM THE U.S. WELFARE REFORM

Abstract

This paper shows that electoral incentives matter for the decision to implement novel policies. Our empirical setting is the period prior to and following the U.S. welfare reform in 1996, which marked the most dramatic shift in social policy since the New Deal. Our findings indicate that governors with strong electoral support are less likely to experiment than governors with little support. Yet, governors who cannot be reelected actually experiment more than governors striving for reelection. These findings are robust to controlling for ideology, preferences for redistribution, the state legislature, and cross-state learning.

JEL Classification: D72, D78, O35, H75, I38

Keywords: Policy innovation, Electoral incentives, Welfare reform, Spillovers

Andreas Bernecker - bernecker@fh-aachen.de
FH Aachen University of Applied Sciences

Pierre Boyer - pierre.boyer@polytechnique.edu
CREST, Ecole Polytechnique and CEPR

Christina Gathmann - christina.gathmann@awi.uni-heidelberg.de
University of Heidelberg

The Role of Electoral Incentives for Policy Innovation: Evidence from the U.S. Welfare Reform*

Andreas Bernecker[†] Pierre C. Boyer[‡] Christina Gathmann[§]

This version: May 27, 2019

Abstract

This paper shows that electoral incentives matter for the decision to implement novel policies. Our empirical setting is the period prior to and following the U.S. welfare reform in 1996, which marked the most dramatic shift in social policy since the New Deal. Our findings indicate that governors with strong electoral support are less likely to experiment than governors with little support. Yet, governors who cannot be reelected actually experiment more than governors striving for reelection. These findings are robust to controlling for ideology, preferences for redistribution, the state legislature, and cross-state learning.

*We are grateful to Toke Aidt, Manuel Arellano, Henning Bohn, Antonio Ciccone, Matz Dahlberg, Peter Fredriksson, Vincenzo Galasso, Björn Tyrefors Hinnerich, Mario Jametti, Helene Lundqvist, Andreas Madestam, Dylan Minor, Espen Moen, Jörg Oechssler, Paanu Poutvara, Torsten Persson, Per Pettersson-Lidbom, Federico Revelli, Christian Riis, Elisabeth Schulte-Runne, Albert Solé-Ollé, Clémence Tricaud, Christoph Vanberg, Galina Zudenkova, and seminar participants at Freiburg, Heidelberg, IEB in Barcelona, Innsbruck, Konstanz, Mannheim, Norwegian Business School, Stockholm, the Silvaplane Workshop on Political Economy, the Conference on Evaluation of Political Reforms in Mannheim, the European Public Choice Society Meeting, the CESifo Conference on Public Sector Economics, the IIPF, and the EEA Annual Meeting for many helpful suggestions. We thank Greg Shaw, Dave Andersen, and Paul Ehmann for sharing their data. Jan Scherer, Michael Hellwig, and Robert Aue provided outstanding research assistance. We are responsible for all remaining errors.

[†]FH Aachen University of Applied Sciences, bernecker@fh-aachen.de

[‡]CREST, École Polytechnique, pierre.boyer@polytechnique.edu

[§]Department of Economics, University of Heidelberg, IZA and CESifo, christina.gathmann@awi.uni-heidelberg.de (*corresponding author*)

1 Introduction

“The country needs and, unless I mistake its temper, the country demands bold, persistent experimentation. It is common sense to take a method and try it: If it fails, admit it frankly and try another. But above all, try something”

Franklin D. Roosevelt

Just as technological innovations are crucial for an economy’s long-run performance, policy innovations are important for a country’s adaptability to changing environments. In both cases, decisions whether to adopt innovative policies or technologies involve sizable uncertainty. Policy-makers, for instance, often do not know the benefits of the reform for their local jurisdiction or when these potential benefits will manifest. If little is known about the effectiveness of novel policies or technologies, there are potentially large gains from experimentation and possibly large costs of failure.

Indeed, many policy innovations started out as experiments – like the deregulation of the airline industry or the reform of the telecommunications sector in the United States. And yet, there are many situations in which policy-makers do actually *not* ‘try something’ as suggested by F.D. Roosevelt’s quote above. On the contrary, examples abound of reforms *never* implemented; or policies that persist despite perceived failure. Hence, a central question in political economy is: under what conditions do politicians actually engage in experimentation and when do they shy away from it?

Theoretical models of political agency highlight that electoral incentives may be one reason why decision-makers are hesitant to engage in policy experiments when both voters and politicians are uncertain about the mapping from policies to realized outcomes (see, e.g., Rogoff, 1990; Harrington, 1993; Majumdar and Mukand, 2004). If experimentation affects a politician’s reelection chances, a conflict of interest between voters and politicians may arise when *ex ante* profitable policies are tried. A popular incumbent with strong electoral support, for instance, has a lot to lose in case a policy experiment turns out to be a failure. As such, she might be hesitant to engage in policy experimentation even if it is beneficial for the electorate.

To explore the role of electoral incentives for the decision to experiment with novel policies, we develop a simple political agency model in which voters are uncertain about the best policy, but politicians may have some information depending on their prior experience. Politicians care about voters’ welfare as well as their reelection chances.

During her term in office an incumbent can decide whether to stick to the known status quo policy or experiment with a new and uncertain policy. A politician, uncertain whether the experiment is beneficial, weighs the (risky) reputation gains from an experiment against her safe initial reputation. The higher her initial reputation (hence, her reelection chances *ex ante*), the more an incumbent has to lose if the policy experiment fails. As a result, reputation concerns make politicians hesitant to engage in risky experimentation. Politicians who cannot be reelected, in turn, might actually be more willing to experiment compared to the average politician.

The model also has something to say when a policy experiment might be reversed. If an experiment fails in the short term, it is either because the policy is wrong for the current state of the world or because of bad luck. A politician with little knowledge whether the experiment is beneficial in the current state of the world is more likely to make a mistake when experimenting than a politician who knows the state of the world. Yet, larger potential welfare gains of an experiment, which increases the incumbent's willingness to gamble with voters' welfare, reduce the likelihood of a policy reversal.

Little is known empirically whether and when policy-makers engage in experimentation – and what role electoral incentives play. The main contribution of this paper is to fill this gap. Our empirical setting is the period before and after the 1996 welfare reform, the most important shift in U.S. social policy since the New Deal. The 1996 reform marked a radical turn in the provision of income support for the poor, which emphasized workfare rather than welfare. The reform decentralized most of the authority to design welfare policies from the federal government to the U.S. states. Even prior to 1996, states could apply for welfare waivers to experiment with new welfare rules under the old 'Aid to Families with Dependent Children' (AFDC) program. After the adoption of the 'Temporary Aid to Needy Families' (TANF) program in 1996, states were largely free to choose their own welfare policies within some broad federal rules (see Harvey et al., 2000). Yet, many of the new policy options, like eligibility criteria, work requirements while on welfare or the type of sanctions in case of non-compliance, had never been tried before. Further, state policy-makers had little prior knowledge whether and which novel policies would benefit their local jurisdiction. Studying U.S. welfare policy over the past decades thus provides an ideal setting to explore why (and when) state policy-makers

engaged in policy experimentation.¹

For the empirical analysis, we collect detailed information on experimentation with waivers and TANF welfare rules and combine it with a rich state-level panel dataset of political and socio-demographic characteristics between 1978 and 2007. The empirical approach relies on panel fixed effects regressions controlling for aggregate influences, state-specific unobservables, and state-specific linear trends. We perform a large number of robustness checks to rule out concerns about omitted variables and reverse causality. Further, we assess several alternative explanations for the extent and timing of policy experimentation: governor ideology, voter preferences for redistribution, the role of the state legislature, and spillover across states.

Our empirical results provide strong support for the view that reelection concerns influence the decision to experiment with novel welfare policies. Strong electoral support for a governor reduces experimentation for governors who may be reelected, but actually increases experimentation among lame duck governors. The influence of reelection concerns can explain a sizable fraction of the variation in the number of waivers or changes in TANF rules implemented in a state. Experimentation with redistributive policies might also be influenced by ideological differences between policy-makers. Republicans, who have long pushed for the idea of workfare rather than welfare, certainly played an important role in the federal Welfare Reform of 1996. While the ideology of the governor matters, electoral incentives continue to play an important role even conditional on governor ideology.

In contrast to governor ideology, we find little influence of voter ideology, the preferences for redistribution or the demand for welfare among the electorate. Our results are further robust to controlling for characteristics of the state legislature as measured by party composition, legislative polarization, or electoral competition. Politicians who need to decide on unknown policies might look to their neighbors for inspiration or engage in welfare competition instead. Defining “neighbors” in terms of geography, population size or ideology, we find that cross-state spillovers matter but do not alter our conclusions that electoral incentives are important determinants of policy experimentation.

¹The decision to decentralize reflects the idea of laboratory federalism that in a world of imperfect information, decentralization is beneficial to foster experimentation in different jurisdictions and to learn about the best policy implemented (see Oates, 1999, for a survey).

2 Related literature

A vast literature has investigated the consequences of the U.S. welfare reform for caseloads, employment, earnings, poverty, marriage patterns and other socio-economic outcomes (see Blank, 2002; Grogger and Karoly, 2005; Moffitt, 2002, for surveys). The question why certain states adopted new welfare policies and experimented with novel policies – like eligibility rules, requirements during benefit receipt and sanctions in the case of non-compliance – has received little attention so far. Our article provides the first empirical test that reelection concerns are important to understand the dynamics of welfare policy-making in the states.²

The analysis in this paper also contributes to the literature on electoral incentives (e.g., Besley and Case, 1995a, 2003; List and Sturm, 2006; Besley, 2007; Alt et al., 2011; Ferraz and Finan, 2011; Gagliarducci and Nannicini, 2013; de Janvry et al., 2012; Martinez-Bravo et al., 2014).³ Existing studies mostly focused on how accountability is strengthened by elections.⁴ Our paper shows for the first time that electoral incentives also matter for policy innovations. The analysis in this paper suggests a potential downside of electoral incentives, which has so far not been investigated. Term limits may benefit voters because politicians who cannot be reelected worry less about their reputation and hence are more likely to experiment with risky, but potentially beneficial policies.⁵

Our theoretical setup builds on a large literature on reputation concerns following the seminal work of Holmström (1982, 1999).⁶ Closer to us are models where voters and politicians are both uncertain about the mapping from chosen policies to realized outcomes (Harrington, 1993; Majumdar and Mukand, 2004; Fu and Li, 2014; Willems, 2013;

²In political science, Lieberman and Shaw (2000) relate the choice of welfare rules to local demand conditions, while Soss et al. (2001) analyze the relationship between a state’s racial composition or ideology and selected policy rules under TANF.

³Besley and Case (1995a), for instance, show that reelection concerns indeed keep politicians in check: public spending and taxes are lower in a governor’s first term of office but increase when the governor faces a binding term limit. More recently, Ferraz and Finan (2011) demonstrate that electoral accountability reduces corruption and rent-seeking, while de Janvry et al. (2012) find evidence for a more successful implementation of a national welfare program among local politicians in Brazil.

⁴Electoral accountability may benefit voters through two channels. Elections allow voters to select those incumbents whose prior track record suggests they are of high ability. In addition, electoral accountability constrains the opportunistic behavior of incumbents.

⁵In related theoretical work, Smart and Sturm (2013) demonstrate that term limits may benefit the voter if some politicians are “public-spirited”, i.e. have payoffs that coincide with those of the electorate. As in our model, electoral accountability and the value of future office may discourage public-spirited politicians from taking actions that are in the interests of voters because it reduces their reelection prospects.

⁶An early application to the political arena is the political agency model by Rogoff (1990).

Dewan and Hortala-Vallve, 2017).⁷ As in our setting, politicians trade off the potential welfare and electoral gains when undertaking a policy reform against the reputation loss that a failure would deliver. The main focus in our paper is to bring a simple version of these models to the data.

Further, our study speaks to the literature on policy persistence, which analyzes why some policies persist even if they are known to be a failure.⁸ We focus on the reputation costs of reversals as one explanation why policies persist; most importantly, we provide empirical evidence under what conditions policies persist during the TANF reform.

Finally, the empirical analysis is also related to the literature on policy innovations more broadly, be it the adoption of state lotteries, new taxes or labor market regulation (e.g. Walker, 1969; Berry and Berry, 1990, 1992).⁹ Most of this literature does not test a particular channel encouraging or inhibiting policy experimentation. Our paper focuses on the role of reelection concerns as one mechanism why policy experiments might or might not be implemented.

3 U.S. Welfare Policy and Experimentation

3.1 The Old AFDC Program

Since the New Deal, the program ‘Aid to Families with Dependent Children’ (AFDC) provided financial assistance to needy children lacking parental care or support (see Grogger and Karoly, 2005; Moffitt, 2008). The program was jointly administered by the federal and state governments but the rules of the program were set at the federal level. States could only choose the level of monthly benefits but had otherwise little room to shape

⁷Recent papers by Callander (2011a,b) and Callander and Hummel (2014) model policy experimentation with a richer policy space: the “world” there is not restricted to a two-state world and policies can take more values than the binary policy choice (reform or not) adopted in most of the earlier literature. Callander and Harstad (2015) present a model where heterogeneous districts choose both whether to experiment and the type of policy experiments. Reelection concerns of the policy makers are absent in their analysis. Instead, we study a setup with two states of the world and a binary decision in order to analyze the consequences of reelection concerns for policy experimentation and reversals.

⁸Potential mechanisms stressed in the previous literature are individual-specific uncertainty about the winners and losers of a reform (Fernandez and Rodrik, 1991; Jain and Mukand, 2003; Ciccone, 2004), conflict between different groups about who will bear the costs of reform (Alesina and Drazen, 1991; Drazen and Grilli, 1993), political constraints on economic reform plans (Dewatripont and Roland, 1992), electorate’s demand for reform (Prato and Wolton, 2018), or vested interests benefiting from the status quo (Coate and Morris, 1999).

⁹Karch (2007); Berry and Berry (2007) provide surveys of this political science literature.

welfare policy.¹⁰

Starting in the early 1960s, state governments obtained more autonomy to experiment with welfare rules. Section 1115 of the Social Security Act allowed the Secretary of Health and Human Services to waive federal AFDC rules and regulations. States could petition the U.S. Department of Health and Human Services (DHHS) to change welfare rules in temporary pilot projects. Waivers were primarily used to implement policies that would reduce welfare dependency and encourage employment. Most waivers implemented multiple changes to welfare rules simultaneously, which frequently included stricter work requirements and tougher sanctions in case of non-compliance. To apply for a waiver, state governments submitted a detailed application, which rules and program elements they planned to modify, and what regulations they wanted to implement instead. The proposed waiver provisions were then reviewed, and sometimes altered, by the federal offices with jurisdiction over the state's proposal. The Secretary of Health and Human Services made the final decision whether to approve a waiver, request changes or deny it. In most cases, the waiver was approved as is or after some adjustments.

To measure policy experimentation during the AFDC period, we code whether a state applied for a waiver. Information on waiver applications in each state is available from Koerper (1996), Crouse (1999), and Lieberman and Shaw (2000). We count all waiver applications irrespective of whether the state planned to implement it only in specific counties or the state as a whole. To reduce any bias through the federal review process, we include all waiver applications whether they have been approved and implemented, denied by the Secretary of Health and Human Services or withdrawn by the state. If a state did not apply for a waiver in a given year, the variable is set to zero.

Welfare waivers were rare until the late 1980s when AFDC caseloads approached almost 4 million families or around 11 million recipients nationwide. Waivers became more popular under President Reagan when federal aid for AFDC declined and government funds tightened during the 1990-1991 recession. Fifteen waivers were approved in 14 states under the Reagan administration and another 15 applications from 12 states under the Bush administration (see Harvey et al., 2000). As reform efforts accelerated during the

¹⁰Single-parent families were the primary beneficiaries of the AFDC program. For a family of three with no other income, for instance, monthly benefits in 1995 varied from \$119 in Mississippi to \$720 in Alaska. States could also implement a few eligibility rules like "fit parent" or "suitable home" provisions in order to limit payments to families with unsatisfactory behavior.

first term of the Clinton administration, the federal government approved 83 waivers (see Figure A1 for the evolution of waiver activity between the late 1970s and 1996).

3.2 The 1996 Welfare Reform and State TANF Programs

The passage of the Personal Responsibility and Work Opportunity Reconciliation Act (PRWORA) in August 1996 abolished the old AFDC program and replaced it with the ‘Temporary Aid to Needy Families’ (TANF) program. The 1996 reform reshaped the landscape of U.S. welfare policy along at least three dimensions: first, the reform decentralized the authority to design welfare programs to the state level. State governments were now free to choose their own welfare rules, however generous or restrictive, bound only by some federal requirements.¹¹ As states had to decide on many new policy rules – like eligibility, work requirements, sanctions or time limits – policy rules have changed even after the adoption of TANF at the state level and do so until today.

Second, states now carried the full fiscal responsibility for their welfare policies. After 1996, federal funding changed from a matching grant into a single block grant. States who spent more than the federal block grant had to shoulder all the additional costs of more generous provisions or higher caseloads.¹² The federal government further introduced financial penalties, in the form of percentage reductions in the block grant, on states with excessive caseloads or states that failed to comply with the federally mandated work participation rates of recipients. States therefore had strong incentives to keep their spending levels and welfare caseloads in check. These incentives are reflected in the evolution of caseloads, which fell nationwide by more than 56% between 1994 and 2000.¹³

Finally, TANF had the explicit goal of reducing welfare dependency and boosting work participation. By setting time limits to benefit receipt, for example, TANF pushed the idea that government support was a temporary measure rather than a permanent source

¹¹In practice, waiver provisions remained in place until they expired or were replaced by new statewide TANF programs. State TANF programs replacing existing AFDC programs were implemented between September 1996 (Massachusetts, Michigan, and Vermont) and as late as January 1998 (California).

¹²The block grant for each state was fixed at the average federal spending level in the 1992-1995 fiscal years. States were further required to contribute to other program components like subsidized child care. Under the old AFDC program, in contrast, a matching grant by the federal government co-financed state and local contributions such that the costs of rising caseloads were shared between federal and state governments.

¹³The literature suggests that the booming U.S. economy of the 1990s is responsible for about a third of the caseload reduction; the remainder is likely due to states’ incentives to tighten access to welfare benefits (see, e.g., Council of Economic Advisors, 1999; Fang and Keane, 2004).

of income for poor families.¹⁴ Work requirements became much stricter under TANF than under AFDC, and focused on active participation in the labor market rather than training and education. Recipients who failed to comply with the new work requirements also faced much harsher sanctions compared to the old AFDC program.

To identify policy experiments during the TANF period is challenging as states had many policy dimensions to experiment with. We focus on four key areas that were at the center of the public and academic debate during the reform process (see, e.g., Crouse, 1999; Fang and Keane, 2004; Grogger and Karoly, 2005): (1) Whether the state adopted a family cap, which limits welfare benefits if a child was conceived and born while the parent is on welfare. (2) The rules imposed for time limits on benefit receipt. (3) The type of work requirements adults had to fulfill in order to remain eligible for benefit receipt. (4) The sanctions imposed if a benefit unit did not fulfill the work requirements. We collect information on these welfare rules in each state and year from the Welfare Rules Database of the Urban Institute, which provides the most comprehensive and up-to-date information on state policy rules for the post-1996 period (Urban Institute, 2015). The data appendix and Table A1 provide more details about each policy rule analyzed.

To illustrate how we measure policy experiments or reversals, take the example of family caps. Between 1996 and 2010, twenty-two states (among them, California, Florida, and Illinois) adopted a family cap, while five states (among them, Illinois) later abolished it. We thus code a policy experiment as one if a state first adopts a family cap in any year after 1996 and zero if no family cap was adopted; we code a policy reversal if a state abolishes the family cap in any year between 1997 and 2010. The coding for other welfare rules follows a very similar logic: a tightening of a policy rule relative to the status quo in a specific year is coded as an experiment, abolishing that restriction in a later year is coded as a policy reversal.¹⁵ Figure A2 shows that many policy experiments occur shortly after the PRWORA reform. The Figure also demonstrates that policy reversals (the dashed line) are rare throughout the whole period.

¹⁴The federal government set a lifetime time limit for benefit receipt of 60 months. States that offered longer time limits have to use state funds for recipients beyond the federal time limit.

¹⁵Note that we can track changes in specific policy rules only since 1996. Rule changes during the waiver period are not captured systematically in the Welfare Rules Database of the Urban Institute. That will generate some measurement error in our post-1996 policy experimentation measure. Suppose a state imposes a family cap under a waiver, but did not have a family cap in 1996 recorded in the Welfare Rules Database. If the state again adopted a family cap after 1996, we would code this as an experiment even though the state has tried this policy rule during the waiver period.

3.3 Role of Governors

State governments became key decision-makers during the era of U.S. welfare reform. Governors in particular played an important role in initiating welfare waivers. Wisconsin's governor Tommy Thompson, for instance, became a leading figure as he made welfare reform a top priority and immediately created a task force to reform the existing AFDC system (Mead, 2004; Kaplan, 2000). The Thompson administration applied for three waivers in 1988, 1992, and 1993. In 1993, Governor Thompson used his line item veto to implement the new Wisconsin Works, better known as W-2. With its emphasis on labor market participation, stricter work requirements and harsh sanctions for noncompliance, the W-2 program became a blueprint for the federal welfare reform in 1996. Governors in Delaware, Michigan, Minnesota, Nebraska, and Ohio played similar defining roles for the speed and direction of welfare reform (see Weissert, 2000; Winston, 2002).¹⁶

Like all policy-makers, state governors faced a lot of uncertainty, which welfare policies to choose. Even if a governor had a general sense of the direction of reform she still had to decide which and how to implement many new rules. A governor who wanted to reduce welfare caseloads, for instance, could restrict eligibility, exclude certain groups, impose stricter employment requirements or impose harsh sanctions upon non-compliance (or a combination thereof). Even if governors could build on some evidence from waiver states after 1996, it was unclear whether a similar policy bundle would achieve the same results if implemented in a state with a different population structure, demand for welfare or economic situation (Gelbach, 2017).

4 Theoretical Framework

To guide our empirical analysis on the decision to experiment with (and possibly reverse) uncertain welfare rules during the AFDC and TANF period, we develop a political agency model inspired by Majumdar and Mukand (2004). Our main goal is to highlight under which conditions a politician engages in experimentation and the role played by electoral incentives. To simplify the exposition, all proofs are relegated to the Online Appendix.

¹⁶State governors also played a prominent role in pushing and keeping federal welfare reform on the political agenda, mostly through lobbying of the National Governors' Association (Weaver, 2000).

4.1 Setup

Policies. Consider an elected politician who has an opportunity to enact a reform. The incumbent needs to decide whether to continue with the safe status quo policy a_S or experiment with the new policy a_N . Both policies affect overall welfare, which is publicly observable. Different policies are appropriate for different environments; hence, the success of the new policy a_N is contingent on the underlying state of the world. The state of the world may be one of two types denoted by S and N , with a_S and a_N being the appropriate policies for the two environments, respectively. If the underlying state of the world is S , then enacting the new policy a_N causes a net loss in welfare. The state N (resp. S) occurs with probability p (resp. $1 - p$).

The status quo policy a_S is assumed to deliver a fixed welfare, normalized to 0, independently of the state of the world. No benefits from the new policy occur in state S . In state N , the new policy initiative delivers a benefit Δ . However, this benefit occurs either in the short term or in the long term. The benefit realizes in the short term (resp. in the long run) with probability q (resp. $1 - q$).

The total cost of enacting and continuing the new policy initiative is c , which is incurred independently of the success of the reform. This cost may represent the additional training of the bureaucracy to effectively administer the new policy or other costs of implementation. The total cost c consists of short- and long-run costs; hence, $\frac{c}{2}$ has to be paid when the reform is implemented and again if the reform is maintained in the long term. If the reform is reverted back to a_S instead, then the long-term cost is not incurred. The benefit of a reform that is reversed is zero, however.¹⁷

We make the following assumptions:

$$p\Delta - c > 0, \tag{A1}$$

$$pq\Delta - \frac{c}{2}(1 + pq) > 0, \tag{A2}$$

and

$$\beta\Delta - \frac{c}{2} < 0, \tag{A3}$$

¹⁷If there is a positive benefit occurring with some probability instead, politicians would face a higher cost of “gambling” with the welfare of the citizens. This assumption would increase the likelihood that a policy experiment is reverted and hence, would make the incentives of politicians more congruent with the preferences of the electorate.

where $\beta = \frac{p(1-q)}{1-pq} < p$ is the probability that the state of the world is N conditional on a short-term failure of the reform.

Assumption (A1) implies that a reform that is continued has positive expected welfare gains ex ante. In our setting, this assumption can be justified on three grounds: first, the perception that the old AFDC program was too permissive and needed reform was widespread among policy-makers in the late 1980s and early 1990s (see e.g. Haskins, 2006; Reintsma, 2007). Second, decentralized policies would improve average welfare if local policies provided a better fit to local needs (compared to a common policy for all). Even if the experiment tightened rules and thus reduced recipients' welfare, average state welfare might go up if the additional money is used to benefit the average inhabitant. Assumption (A2) says that undertaking the reform and reverting back to the status quo in the face of short-term failure has positive expected benefits ex ante. Finally, assumption (A3) implies that, if the benefit does not realize in the short term, the expected future benefit from continuing with the reform is negative.¹⁸

Timing. The game has three periods, $T = \{0, 1, 2\}$. At the beginning of the game ($T = 0$), the incumbent politician has the opportunity to change the existing policy a_s . She receives a private signal about the appropriateness of the policy, i.e. the state of the world, and then faces the choice of either maintaining the status quo policy a_s or enacting the new policy a_N . If she decides to maintain the status quo, then the ex-post welfare is 0. On the other hand, if the new policy is enacted, a cost $\frac{c}{2}$ is initially incurred. Then, the politician learns about the realization of short-term benefit from the reform at the beginning of $T = 1$. If the new policy turns out to be successful, the politician knows that the state is N , and therefore a_N is the appropriate policy. This implies that the benefit Δ is secured by continuing the reform and incurring the long-term cost $\frac{c}{2}$ at the beginning of period $T = 2$. However, in the case where the benefit does not realize in the short term, it is unclear whether the failure is because of bad luck or because the policy is not appropriate. The politician now has to decide whether to continue with the policy initiative a_N or to revert back to the initial status quo a_s . If she reverts back to a_s , ex-post welfare is $-\frac{c}{2}$. If she continues with a_N , the society again incurs a cost $\frac{c}{2}$ and gets the long-run benefit only if the state of the world is N . The ex-post welfare is observed

¹⁸We implicitly assume that the parameters of the model take values such that assumptions (A1), (A2), and (A3) hold simultaneously. This is the case for a large range of parameter values.

before the next election that takes place at the end of period $T = 2$, corresponding to the end of the incumbent's term in office.

Politicians. There are many politicians that differ in their capability to assess the appropriateness of policies for the economy, and can be either of high quality or of low quality. The incumbent politician knows her type. For simplicity, we assume that a high-quality incumbent knows the state of the world. The high-quality politician serves as a benchmark to compare the decisions taken by different politicians depending on their electoral incentives and welfare considerations. On the other hand, a low-quality politician does not know anything *ex ante*, i.e. she only knows that the state of the world N is realized with probability p , and that in this state the benefit realizes in period 1 (resp. period 2) with probability q (resp. $1 - q$).¹⁹

The incumbent politician cares about both welfare of her constituency as well as her own future electoral prospects. The incumbent's objective function is

$$\gamma(\text{Welfare}) + (1 - \gamma)(\text{Probability of Reelection}), \quad (1)$$

where $\gamma \in [0, 1]$ is the relative weight that she puts on welfare.

Finally, at the beginning of the game ($T = 0$) the incumbent has earned a reputation among the electorate represented by $\lambda \in [0, 1]$. We assume that this reputation does not depend on future decisions taken by the incumbent and remains constant until the next election.²⁰

Voters and Reelection Rule. The electorate consists of a representative voter who chooses to reelect the incumbent or vote for a challenger at the end of the term $T = 2$. The voter cares only about *ex-post* welfare. We assume that the voter does not know that there are high- or low-quality politicians so that he does not think that the incumbent and the many potential challengers differ in competence. However, the voter may take into account the reputation of the politician when deciding to reelect the incumbent. Assuming the representative voter does not know that politicians differ in abilities simplifies the analysis considerably. Under this assumption, there is no signaling game in the policy

¹⁹We discuss below a version of the model where politicians have an ideological bias in favor or against the reform.

²⁰See Majumdar and Mukand (2004) for an analysis where the reputation of the incumbent changes with the decision to undertake reforms.

choice of the incumbent: both challenger and incumbent are perceived to be identical by the voter except in the case where the incumbent's reputation is positive.²¹

We assume that the representative voter follows a simple reelection rule: reelect the incumbent with probability 1 if ex-post welfare is positive; never reelect the incumbent if ex-post welfare is negative; and reelect the incumbent with probability λ if ex-post welfare is 0. This reelection rule is in the spirit of agency models where voters punish the incumbent for bad performance (see Persson and Tabellini, 2000).

Benchmark: Policy Choices of a High-quality Politician. To provide a benchmark, we first analyze the policy choices of a high-quality incumbent.

Proposition 1. *For any relative welfare weight γ and any reputation λ , a high-quality incumbent always undertakes a reform if and only if the underlying state of the world calls for a reform (state N). In addition, a reform undertaken is never reverted back to the status quo.*

A high-quality politician knows the state of the world. If the state of the world is S , she will never experiment because there are neither welfare nor electoral gains from a reform. Proposition 1 shows that a high-quality incumbent has her objective fully aligned with the interest of the representative voter: there is no conflict of interest between maximizing voter welfare and her reelection chances.

4.2 Policy Choices of a Low-quality Politician

We now show that the decision-making of a low-quality incumbent may exhibit a conflict between maximizing welfare and reelection chances. We first focus on the decision to revert an experiment.

Proposition 2. *A low-quality incumbent reverts back to the status quo in face of a short-term failure of the reform if and only if $\gamma \geq \gamma^*$, where*

$$\gamma^* := \frac{1}{1 + \frac{c}{2\beta} - \Delta}. \quad (2)$$

²¹The signaling game where the representative voter makes inferences on quality of the incumbent based on the sequence of policy choices and the realized outcomes is analyzed in Majumdar and Mukand (2004). The predictions they derive are similar to the ones we obtain in our simplified model.

A low-quality politician who observes the short-term failure of a reform does not know whether the reform will be successful in the long term. This effect pushes her to revert back to the status quo to increase expected welfare. However, reelection concerns pull her towards sticking with the reform. Proposition 2 disentangles these two conflicting forces: if the relative welfare weight is above the threshold γ^* defined in equation (2), the low-quality incumbent reverses the reform as welfare considerations outweigh reelection concerns.

Corollary 1. *A low-quality incumbent is more likely to revert in the face of short-term failure of the reform when γ increases and Δ decreases.*

While high-quality politicians never revert an experiment (Proposition 1), Corollary 1 shows that low-quality politicians make their reversal decision dependent on the welfare weight and potential gains from the experiment.

We now analyze whether an uninformed politician who cares about both welfare and reelection ($\gamma \in [0, 1]$) undertakes the reform.

Proposition 3. *(I.) When $\gamma \geq \gamma^*$, then a low-quality incumbent undertakes a reform if and only if*

$$\gamma \left(pq\Delta - \frac{c}{2}(1 + pq) \right) + (1 - \gamma)pq \geq (1 - \gamma)\lambda. \quad (3)$$

(II.) When $\gamma < \gamma^$, then a low-quality incumbent undertakes a reform if and only if*

$$\gamma(p\Delta - c) + (1 - \gamma)p \geq (1 - \gamma)\lambda. \quad (4)$$

Proposition 3 follows from a comparison of the politician's payoff if she undertakes the reform (for the two cases where she will or will not revert back to the status quo in the face of short-term failure) with the payoff from sticking with the status quo and running the next election on her initial reputation instead.

If the policy choices of the incumbent are driven only by welfare considerations ($\gamma = 1 > \gamma^*$), the decision-making policy described in Proposition 3 is the optimal one from a welfare perspective under the veil of ignorance. On the opposite, a low-quality incumbent who is purely office-motivated ($\gamma = 0 < \gamma^*$) implements the reform if and only if $p \geq \lambda$.

A low-quality politician might not maximize voter welfare for two reasons. First, the politician's initial reputation can conflict with the decision to experiment with beneficial

reforms. An uninformed politician with a high enough initial reputation will not experiment even though it would maximize ex ante welfare. As a consequence, reputation concerns result in too little policy experimentation from a welfare perspective. Second, in the face of short-term failure a purely office-motivated politician does not take into account the expected welfare loss of continuing with the reform and “gambles” for reelection with the welfare of the electorate (Proposition 2). This happens because a policy reversal back to the status quo implies that the politician would lose the election for sure, whereas if the politician persists with the reform, she can win the election if the long-term benefit realizes (which occurs with probability $\beta > 0$).

Corollary 2. *A low-quality incumbent is more likely to experiment when Δ increases, γ increases, and λ decreases.*

Corollary 2 states that low-quality politicians experiment more when the potential welfare gains are high, when they care more about welfare (and less about reelection) and when the initial reputation is low (because then the incumbent has little to lose in terms of her reelection chances).

Ideology and Experimentation. So far there is no role for ideology in the model as politicians agree on the decision to undertake the reform given the state of the world. We now extend the model to allow for some ideological differences between politicians: we assume that Republicans assign a higher probability that the reform is the right thing to do, i.e. that the state of the world calls for a reform. This assumption captures the common perception of the welfare reform, and certainly the federal PRWORA reform, that it was fueled by a conservative agenda to curb welfare spending. Formally, we assume that Republicans who do not observe the state of the world believe that the state of the world N (where an experiment is welfare-enhancing) occurs with probability $p_R > p$.²² In contrast, Democrats who do not observe the state of the world continue to have the belief that the state of the world is N with probability p .²³ Since the representative voter cares only about ex-post welfare, the reelection rule remains unaffected by the bias of the politicians.

²²Democrats and Republicans could also differ in the assessment of Δ even if they agree on p . In this case, the relevant comparative statics are those in Corollaries 1 and 2.

²³The analysis yields the same prediction if Democrats underestimate the probability of the state being N relative to Republicans instead.

Proposition 4. *If Republicans believe that the state of the world where an experiment is welfare-enhancing occurs with higher probability than the Democrats, Republicans will be more likely to experiment with the reform and to persist in the face of short-term failure relative to Democrats.*

Proposition 4 states that Republicans will become more likely to experiment with the reform (which increases expected welfare) and to persist in the face of short-term failure (which decreases expected welfare) relative to a Democrat when they perceive that the reform is more likely to be the right thing to do than Democrats.

Empirical Predictions. The model highlights how electoral incentives and ideology may influence policy experimentation. The first channel is that politicians with higher initial electoral support among the electorate are less likely to experiment, especially if politicians are purely office-motivated. Incumbents who seek reelection are hesitant to risk their electoral support in order to pursue experiments that might fail.²⁴

The second channel is that the influence of electoral incentives should be weaker for politicians with a lower weight on reelection. Politicians who are more concerned about welfare are more likely to experiment than politicians who care mostly about their own reelection. Furthermore, the model predicts that higher potential gains from a reform encourages experimentation. Finally, if Republicans are more likely to believe that an experiment is the right thing to do in the current state of the world, they should experiment more than Democrats.

The model also speaks to the likelihood of observing a policy reversal. The initial electoral support of the politician should have no effect on the likelihood of a reversal. The reason is that a politician who reverts back to the status quo is voted out of office for sure. A politician who does not revert is either successful in the long run (in which case she is reelected for sure) or is not successful (in which case she is voted out of office). In all of these cases, the initial electoral support has no impact on reelection chances and hence, does not matter for the decision to revert an experiment. A politician is less likely to revert if she cares more about reelection and the potential welfare gains from

²⁴In the Online Appendix, we extend the model to allow that voters derive some welfare from the reputation of the politician even if an experiment is implemented. In this extension, we predict a non-monotone effect of reputation on the likelihood to experiment: politicians with very low, intermediate, and very high reputation potentially exhibit different experimentation behavior. However, we find no evidence for a non-monotone effect in our empirical analysis, see columns (6) and (7) of Table 3.

an experiment are high. Finally, Republicans are more likely to persist in the face of short-term failure than Democrats.

5 Data Sources and Estimation Approach

5.1 Data Sources

For our empirical analysis, we require measures for the electoral support of a governor (an empirical counterpart to λ), the importance of reelection concerns (a proxy for $1 - \gamma$) and the potential gains from experimentation (an empirical counterpart to Δ). All of these concepts are difficult to measure as they are ultimately unobservable to the researcher.

We use the vote margin in the past gubernatorial race to proxy the governor's electoral support. Given the incumbency effect, governors who were elected with a larger majority are more likely to be reelected conditional on the state of the world and the type of challenger. To rule out that the past vote margin just captures electoral competition, we control for the competitiveness of state elections using the Holbrook-Van Dunk index (Holbrook and Van Dunk, 1993), among others. To measure reelection concerns among governors, we use information on whether a state governor faces a binding term limit in her current term in office, i.e. whether she is a lame duck. Lame ducks might still have some electoral incentives if they plan to run for another office at a higher level. All we require is that lame ducks care somewhat less about their electoral prospects than governors who can and want to be remain in office. A lame duck might differ from other governors along other, possibly unobservable dimensions like competence or experience (Alt et al., 2011). Below, we show that controlling for governor age, experience, and educational credentials has few effects on our estimates.

Further, we need a proxy for the potential gains of a policy experiment, which is challenging to measure. An important motivation to apply for a waiver prior to 1996 was to reduce caseloads and hence, welfare spending. The incentive to curb spending became even stronger during the TANF period because states now had to cover all welfare expenditures exceeding the federal block grant (see Section 3.2). We thus use a state's past AFDC/TANF spending to proxy the potential gains from policy experimentation. The idea here is that the block grant encourages politicians with growing welfare expenditures to experiment with rules that effectively reduce the state's welfare burden and free up

resources for other public goods and services.

To control for governor ideology, we use the governor’s party and Berry’s ideology measure, which varies from 0 to 100 with larger values indicating more liberal views (see Berry et al., 1998). We digitized additional information for each governor, like years of experience in prior political offices and educational background from online biographies. The decision to experiment might also be influenced by voter preferences or their demand for welfare. We use Berry’s measure of citizen ideology as well as Cook’s Partisan Voting Index (PVI) to proxy for voter ideology. In addition, we employ individual-level information on preferences for redistribution and preferences over welfare spending from the General Social Survey to control for state-level preferences for redistribution.

To capture the role of the legislature, we include measures for party composition and polarization. Polarization is calculated as the difference between the DW-NOMINATE scores of each party where the scores are calculated from the state’s representatives in the House or Senate (Poole and Rosenthal, 1984; 1997). We further code whether the state government is politically divided between Democrats and Republicans. States might also imitate or learn from each other. To analyze such cross-state spillovers, we define geographic neighbors (states sharing a boundary); states similar in population size (based on their population in 1978) and ideological neighbors (neighboring states with a governor belonging to the same party). Our measure of cross-state spillovers is then the average number of experiments undertaken in neighboring states in the previous year.

Table 1 contains summary statistics of all variables. The Online Appendix contains additional details on the source and coding of our dependent and independent variables.

5.2 Estimation Approach

Our empirical analysis follows much of the literature on electoral incentives in using fixed effects panel estimators. Specifically, we estimate variants of the following model:

$$Exp_{st} = \alpha_1 EleSup_{st'} + \phi_1 LD_{st'} + \pi_1 LD_{st} * EleSup_{st'} + \mu_1 Spend_{st-2} + \gamma_1' X_{st} + t_t + \theta_s * \tau_t + \varepsilon_{1st}, \quad (5)$$

where Exp_{st} denotes our measure of policy experimentation, which we pool over the AFDC (1978-1996) and the TANF period (1996-2007). $EleSup_{st'}$ is the governor’s electoral support in the last election, LD_{st} is an indicator whether the governor faces a binding term limit and $Spend_{st-2}$ represents the potential gains from experimentation (lagged by 2 years). We control for state population, its age structure and governor age X_{st} . State

fixed effects, θ_s , absorb any state differences in the propensity to innovate. Year fixed effects, t_t , control for aggregate time trends in experimentation, while state-specific trends, $\theta_s * \tau_t$, adjust for state-specific dynamics in experimentation over time.

From the model in Section 4, we expect experimentation to decrease with electoral support for the governor, i.e. $\alpha_1 < 0$ as popular governors are hesitant to engage in risky policy experiments. In addition, electoral incentives should be weaker for governors who cannot be reelected and hence worry less about the risk of policy experimentation, so $\pi_1 > 0$. Also, we expect experimentation to increase with the potential gains from experimentation, i.e. $\mu_1 > 0$.

In the absence of quasi-experimental variation, there are several potential concerns with the model in equation (5). First, electoral support and spending might suffer from reverse causality. Empirically, we deal with this concern in two ways. Our first approach is to use the vote margin in the past election and lagged spending on the welfare program as regressors. Lags help to cope with contemporaneous shocks that might influence both the decision to experiment and program spending, for instance. Lags would, however, not help with persistent shocks or systematic omitted variables that affect both sides of the equation. In the absence of plausible instruments, we augment (5) with a lagged dependent variable to account for persistent factors driving policy experimentation and use an Arellano-Bond panel estimator (Arellano and Bond, 1991).²⁵

Endogeneity in past welfare spending is likely to bias our estimates toward zero rather than inflate the estimate. The model suggests that higher AFDC or TANF spending should raise the likelihood of experimentation. Reverse causality would suggest that past experimentation, by tightening sanctions, work requirements and eligibility criteria, reduces future welfare spending because of declining caseloads. For electoral support, the potential bias from reverse causality is a-priori less clear. If a governor campaigns on policy experimentation, the decision to experiment with welfare rules would boost her future vote margin. Our model, in contrast, predicts that past vote margin reduces experimentation because of the potential electoral costs of failure. Yet, campaigning on

²⁵An alternative method to address endogeneity would be to restrict the sample to close elections, which generates quasi-random variation in the party of the governor elected into office (Lee et al., 2004). The basic idea is then to compare the behavior of a governor who was just elected into office in say 1992 and reelected in 1996 to the behavior of a candidate who won the seat in 1996 from an incumbent of the opposite party. We would expect that the first governor has more electoral support than the second governor because of the incumbency advantage. Unfortunately, a sample of close elections with a vote margin below 2 percentage points has less than 100 observations, which renders a RDD infeasible.

cutting welfare transfers could also reduce electoral support among potential beneficiaries of income support. We show below that controlling for the demand for welfare and preferences for redistribution does not affect our estimates. If anything, these alternative explanations will then make the coefficients on welfare spending and electoral support conservative estimates of the true effect.²⁶ Another concern with our empirical model in equation (5) could be omitted variables, which might bias our effects upward or downward. We address this concern by controlling for a large number of potential confounders, like the demand for welfare in the population, people’s preferences for redistribution, governor ideology, the political structure of the legislature and spillover effects from neighbors. In addition to our state and year fixed effects and state-specific linear, we further allow for state-specific quadratic and cubic trends or decade-specific state dummies to absorb any unobserved non-linear trends in experimentation. Finally, we also explore the sensitivity of our results to non-linearities in the independent variables. None of our qualitative results are affected by these alternative specifications.

Finally, we analyze policy reversals using the following model:

$$Reversal_{st} = \mu_2 Spend_{st} + \phi_2 LD_{st} + \gamma'_2 X_{st} + t_t + \theta_s + \theta_s * \tau_t + \varepsilon_{2st}, \quad (6)$$

where all variables are defined as before. The higher the gains from experimentation and the more concerned the governor is about reelection, the less likely a politician will revert a policy; so, $\mu_2 < 0$ and $\phi_2 > 0$.

6 Empirical Results

6.1 Electoral Concerns and Policy Experimentation

We first show the basic relationship graphically. Figure 1 plots policy experiments against the past vote margin. The linear predictions show that there is a clear negative correlation for governors who can be reelected (the blue line), while the relationship is much weaker for lame duck governors (the red line) – just as predicted by our model. We next investigate the role of electoral incentives for policy experimentation according to

²⁶Alternatively, suppose governors learn from past experiments and voters support politicians if they experimented. We would then expect negative serial correlation in experimentation over time and higher electoral support for the governor. Yet, we find little empirical support for this alternative explanation: past experimentation does not discourage future experimentation and there is no statistically significant relationship between past experimentation and electoral support.

(5). The dependent variable in Table 2 is the number of policy experiments in a state and year during the 1978-2008 period.²⁷ All specifications control for state and year fixed effects as well as state characteristics and governor age. Even columns also add state-specific trends in order to control for idiosyncratic upward or downward trends in experimentation. Columns (1) and (2) of Table 2 document that governors with high electoral support indeed experiment less than governors with less electoral support. Our model provides one explanation for this relationship: politicians, especially if they are purely office-motivated, are hesitant to reform because they weigh the risky electoral gains from an experiment against their safe initial electoral support.

We next investigate whether the influence of electoral support on experimentation is muted for governors who cannot be reelected. Columns (3) and (4) add whether a governor is a lame duck and its interaction with past electoral support to the specification. The estimates show that electoral support still has a negative effect on experimentation among governors who can be reelected. Governors who face a term limit, in turn, are actually more likely to experiment as the sum of the coefficient on the vote margin and the interaction effect of vote margin and lame duck is with 0.02 positive.

The last two columns of Table 2 show that higher potential gains (as measured by past AFDC/TANF spending) are associated with more experimentation. An increase in welfare spending by one standard deviation increases experimentation by 73% of a standard deviation.²⁸ Furthermore, the effect of electoral support on experimentation remains economically sizable even after controlling for the potential gains from experimentation. An increase in the past vote margin by one standard deviation reduces policy experiments among governors who can stand for reelection by about 10% of a standard deviation.²⁹ Lame duck governors, in turn, are somewhat more likely to engage in risky experimentation compared to governors who can be reelected. The effect based on the estimates in column (6) is an increase by 5.3% of a standard deviation $((-0.007+0.011)*9.4/0.705)$.

To address the potential endogeneity of the past vote margin or past welfare spending, we also estimate an augmented model with lagged dependent variable to absorb persis-

²⁷We find very similar results if we use a count model (see column (1) of Table A2) and statistically weaker effects if we use a dummy variable for experimentation (see column (2) of Table A2).

²⁸The standard deviation of real welfare expenditures is 1027.3 (US\$ million). Hence, we get $0.001*1027.3/0.75=0.73$.

²⁹The standard deviation of the past vote margin in our time period is 9.4 percentage points, while the standard deviation of policy experimentation is 0.705. Hence, $(-0.007*9.4)/0.705=-0.093$.

tence in policy experimentation. The results for the Arellano-Bond estimator shown in columns (3)-(5) of Appendix Table A2 are very similar to the baseline estimates in Table 2.³⁰ Table 3 further probes the robustness of our estimates to omitted variables and alternative specifications. To control for state-specific, nonlinear unobservable shocks, we add state-specific quadratic and cubic trends (in column (1)) and state-specific decade dummies (in column (2)) to control for differential developments in policy innovation over time. The effects on electoral support and its interaction with being a lame duck do not change. A potential concern with our empirical model in equation (5) is that electoral support just proxies for the competitiveness of state elections. We therefore include whether the governor has changed in the past two electoral cycles (in column (3)); the Holbrook and Van Dunk index of competitiveness (in column (4)); and both measures simultaneously (in column (5)). The estimates in columns (3)-(5) of Table 3 show no statistically significant influence of competition on policy experimentation. Even more importantly, the coefficients on the past vote margin remain unchanged compared to the baseline. Hence, the past vote margin does not capture that swing states or states with tougher electoral competition experiment less. Instead, there are fewer policy experiments because the current governor has more to lose if her electoral support is high and the experiment fails.

The past vote margin might affect experimentation in a nonlinear fashion; governors with very high electoral support might be willing to experiment more, for instance, because they get reelected almost surely. To test for such nonlinearities, we add a quadratic term (in column (6)) or split electoral support into low, medium, and high support (in column (7)). The results do not suggest any nonlinearities; governors with very high electoral support experiment less than governors who won the election at a more narrow margin conditional on our control variables. We next assess whether our results are robust to the definition of lame duck governors. We drop states with a one term limit (in column (8)) to rule out that the positive effect of lame ducks on experimentation is explained by governors from these states. In column (9) of Table 3, we drop states with no term limit to check whether states with term limit differ from those without.

³⁰The test statistic for the error structure of the residuals further suggests that we do not have persistent autocorrelation in the error term. Finally, the Sargan test statistic cannot reject the null hypothesis that the overidentifying restrictions from the instruments (lags from $t - 2$ to $t - 4$ of the dependent and predetermined independent variables) are valid.

Identification of the lame duck effect now comes from a comparison of governors in their first and second term conditional on controlling for their general experience. For both specifications, the estimates remain unchanged compared to the baseline.

We next explore whether our results are sensitive to how we measure the potential gains from experimentation. The decision to experiment might depend on past growth rates rather than the level of welfare spending. We thus replace past levels with the past growth rate in spending (in column (10)). High growth rates in welfare spending in the past, just like the level of past spending, encourage experimentation suggesting that governors consider experimentation more attractive the more scope there is for reducing expenditures.³¹ Finally, the gains from experimentation could depend on the demand for welfare, which in turn depends on the population structure in a state. To check for these demand-side influences, we include in column (11) of Table 3 the size of the immigrant population, share of Blacks, the share of unmarried births and income inequality to the baseline specification. None of these variables have an effect on policy experimentation and our baseline estimates are again unaffected. Hence, states with higher welfare spending are more likely to experiment even conditional on the size of the immigrant population or the share of single moms, for instance.

Thus far, we assumed that all governors had little information about the state of the world. Yet, some governors might be better in recognizing and discerning valuable information from noisy cues. To explore such heterogeneity, we use two proxies for a governor's ability to predict the state of the world: her prior political experience and her educational background. Experience in prior political offices should be a good indicator of the governor's ability if elections select the most able candidate, or if governors improve their political competence while in office (see, e.g., Alt et al., 2011, for evidence supporting the latter view). Alternatively, we use whether a governor has a Master's or higher degree or not as a proxy for ability (see Galasso and Nannicini, 2011; Gagliarducci and Nannicini, 2013; Kotakorpi and Poutvaara, 2011, for similar approaches).³² The first specification

³¹To further reduce concerns about endogeneity in past AFDC/TANF spending, we also use total welfare spending, which includes spending on Food Stamps as well, as a measure of potential gains. We also find a positive effect here suggesting that high spending on welfare programs makes governors more likely to experiment with policy rules.

³²As these characteristics are not randomly assigned, we should interpret the coefficients as correlations. While voters may observe tenure as a governor, they are unlikely to know exactly how much political experience a governor has accumulated prior to being elected. Voters might be more aware of the educational credentials of a governor, however.

in Table A3 includes an indicator whether governors have less than a Master’s degree (in column (1)) or little prior political experience (in column (3)). According to the model in Section 4, politician with little information about the state of world may be more or less likely to experiment than politicians with better information.³³ The results reflect this ambiguity: Less educated governors experiment less, while governors with little prior experience experiment more. Neither of the coefficients reach statistical significance, however. The second specification (in columns (2) and (4) of Table A3) interacts the indicator for low education or experience with being a lame duck. There is some evidence that governors with little prior experience are more likely to experiment when they are lame ducks, while the net effect for low-educated lame ducks is zero. Most importantly, Table A3 reveals that electoral incentives matter for policy experimentation even after we account for a governor’s ability to assess the state of the world.

6.2 Effect of Governor Ideology

A common perception of the welfare reform, and certainly the federal PRWORA reform, is that it was fueled by a conservative agenda to curb welfare spending. Yet, can ideology among governors also explain experimentation with welfare rules at the state level? Most importantly, are ideological differences among governors responsible for the role electoral incentives play for policy experimentation as shown in Tables 2 and 3? To answer these questions, we first study the relationship between governor ideology and experimentation by including an indicator whether the governor is a Republican (in column (1) of Table 4) or Berry’s measure of party ideology to account for ideological differences within parties across states (in column (2) of Table 4). Berry’s measure varies from 0 to 100 where larger values indicate more liberal views. As expected, Republican governors are more likely to experiment than Democratic governors and more liberal governors experiment less. These findings are consistent with the idea that governor ideology matters for the decision to experiment.

We next add our measures of electoral incentives to the specification: the past vote margin, whether the governor is a lame duck and the interaction between the two as well as past welfare spending (in columns (3) and (4) of Table 4). The results show that the

³³The high-quality politician experiments when the state of the world is N , which occurs with probability p . The low-quality politician is more likely to experiment if $\gamma = 1$, but less likely to experiment if $\gamma = 0$ and $\lambda \geq p$.

past vote margin influences policy experimentation even conditional on governor ideology: a higher vote margin reduces experimentation for governors who can be reelected and slightly increases it for governors who cannot. As before, past welfare spending encourages policy experimentation. To illustrate the relationship between ideology, electoral support and policy experimentation, Figure 2 plots experimentation against the past vote margin separately for Democratic and Republican governors. The pattern is striking: there is a negative relationship between experimentation and the past vote margin for both types, but Republican governors experiment more.

Overall, ideology shifts the level of experimentation as suggested by our theoretical model; but electoral incentives remain an important factor even after accounting flexibly for the ideological stance of the governor. To assess the relative importance of the two channels, we use the estimates in columns (3) of Table 4. Accordingly, one standard deviation in governor ideology (0.5) can account for about 6% of a standard deviation ($0.0866 \cdot 0.5 / 0.705$) in policy experimentation. Hence, governor ideology explains less of the variation in policy experimentation than electoral incentives, which account for 14% of a standard deviation in experimentation.

6.3 Alternative Explanations

Overall, the findings provide strong evidence that electoral incentives of the governor influence the decision to experiment with welfare policy. Yet, our analysis has so far abstracted from other, potentially important drivers of welfare policy: the preferences of voters, the political influence of the legislature or spillover effects across states. We now investigate each of these factors in turn.

Heterogeneity in Voter Preferences We have so far abstracted from voter heterogeneity. The model assumes for simplicity that there is no disagreement about the potential gains from experimentation as there is a single representative voter. Citizens might differ, however, in their assessment of the optimal amount of welfare spending: some might find it too high, others too low.

We first include standard measures like Berry’s measure of voter ideology (in column (1) of Table 5) and Cooks’ Partisan Voting Index (in column (2)) to our baseline. Both specifications show no systematic relationship between state-level shifts in voter ideology

and policy experimentation.³⁴ These broad measures of voter ideology might be too crude to pick up preferences for redistribution. As an alternative strategy, we use data on the preferences for redistribution from the General Social Survey. In column (3) of Table 5, we add an indicator equal to one if respondents in a state are more likely to support redistribution than respondents in the median state as well as its interaction with past welfare spending. In column (4), we use a binary indicator equal to one if respondents in a state are more likely to say that their government spends too little on welfare relative to the median state as well as its interaction with past welfare spending. Experimentation is higher in states whose electorate prefer more welfare spending. Most importantly, our baseline estimates are unaffected by controlling for voter preferences or heterogeneity in beliefs about the benefits of welfare transfers.

Influence of the Legislature So far, we have restricted attention to the governor as the relevant decision-maker. Yet, the legislature can be agenda setter or veto player by delaying policy experimentation, for instance. Previous evidence suggests that party composition and tightness of legislative elections seem to have an effect on welfare spending under AFDC (e.g. Barrilleaux et al., 2002).

In Table 6, we add controls for the political structure of the legislature to our baseline: the party composition in the lower and upper house (in column (1)); the degree of polarization in the parliament, which is often viewed as an indicator for legislative competition (in column(2)); whether the government is divided (in column (3)); and all controls simultaneously (in column (4)). Policy experimentation is less likely if the state parliament is dominated by Democrats; yet, experimentation is actually more likely if the state has a divided government (see Bernecker, 2016, for an empirical investigation). Most importantly, the coefficients on the past vote margin, its interaction with being a lame duck and past welfare spending are unaffected by controls for the state legislature.

Spillover Effects across States Our analysis has also abstracted from spillover effects across states. If potential welfare recipients are geographically mobile, states may engage in welfare competition (see Figlio et al., 1999; Saavedra, 2000; Wheaton, 2000). Spillovers can also arise if policy choices in neighboring states provide information to voters as in

³⁴The zero effect does not necessarily mean that voter ideology has no effect on experimentation as voter preferences are likely to influence governor ideology as well.

models of yardstick competition (Besley and Case, 1995b). In fact, a central insight of the literature on laboratory federalism is that decentralization promotes learning and spillover of knowledge in a federal system (see Oates, 1999; Strumpf, 2002).

We create three measures to capture learning effects. First, we define neighbors as states with similar population size and calculate their mean number of experiments in the previous year. Our second measure uses the average number of experiments by geographic neighbors in the previous year. Finally, we use the average number of experiments in neighboring states whose governor comes from the same party to capture spillovers. Table 7 shows that states learn from their geographic neighbors irrespective of ideological proximity. Most importantly, allowing for cross-state learning does not affect our qualitative conclusions that the governor’s electoral incentives matter for experimentation.

6.4 Evidence on Policy Reversals

We now turn to the decision to revert a policy experiment. Our theoretical framework suggests that policy reversals should be a rare event because a governor who reverses an experiment is not reelected for sure. Table A4 in the appendix confirms this pattern using descriptive evidence on policy reversals during the TANF period.³⁵ For each of the eleven policy rules, columns (1) and (2) of Table A4 list how many experiments and reversals were implemented. Policy reversals are indeed a rare event: while 22 states experimented with a family cap, for instance, only 5 states reversed the family cap later on. Across all rules we analyze, only about 20% of experiments get reversed.

Column (3) of Table A4 indicates that few policy reversals happen within the same electoral cycle as the experiment; in fact, the mean number of years between an experiment and a reversal is 5.5 years (see column (4)). One explanation for this pattern is that governors shy away from reversals because it demonstrates to voters that the governor made a mistake (as stressed in our theoretical framework). Reversals might then only be observed if a governor from the opposite side of the ideological spectrum is voted into power in a later cycle. A second explanation why reversals within the same electoral cycle are rare is that it takes time to observe the success or failure of an experiment. The last two columns of Table A4 indicate that, in many cases, reversals are implemented by governors from the same party as the initial experimenter. This finding suggests that it

³⁵As discussed in Section 2, we cannot measure policy reversals systematically during the waiver period.

likely takes more time than an electoral cycle to learn about and evaluate the outcome of an experiment. We find little support for the argument that most policy experiments are reversed after the political opponent is voted into office (see also Gelbach, 2017).

To investigate reversals more systematically, Table 8 relates the number of policy reversals to our baseline variables using equation (6). An important caveat is that policy experiments should only be reversed when the experiment fails. Unfortunately, we do not observe whether the experiment failed to deliver short term benefits or not. Our reversal measure therefore contains substantial measurement error, which reduces the precision of our estimates. We thus view the results in Table 8 as suggestive in nature. In line with our model, the potential gains from experimentation are negatively correlated with reversals (see column (1)). Governors prefer to gamble and stick to a policy experiment with large potential gains hoping that the experiment turns out to be successful in the long run. Column (2) shows that lame ducks are not more or less likely to revert an experiment, while in the model, lame ducks are more likely to revert because a reversal is less costly for voter welfare if the experiment failed in the short term. Column (3) indicates that Republican governors are less likely to revert an experiment, which is in line with our model. Finally, column (4) includes the past vote margin of the governor, which, according to our model, should not play a role for the decision to reverse an experiment. And indeed, the correlation between electoral support and a policy reversal is very weak (relative to the effect on policy experimentation in Table 2, for example) and never close to statistically significant. The zero correlation suggests that voters do not use observed policy choices to update their beliefs about a politician, possibly because they are not able to fully observe or evaluate these actions within a short time period.³⁶

7 Conclusion

We analyze the role of electoral incentives for policy innovations. Our empirical setting is the U.S. welfare reform in 1996, which marked the most important shift in social policy since the New Deal. Both during the waiver and the post-1996 period, state governments faced many new policy options which had never been implemented before. To guide our

³⁶The result that initial electoral support does not matter for the decision to revert a policy differs from the prediction in the full signaling game of Majumdar and Mukand (2004). There, voters update their belief about the politician's quality after observing policy choices and their outcomes. In that case, the initial reputation would also have an effect on reversal decisions.

empirical analysis, we develop a political agency model, which predicts that reelection concerns make politicians hesitant to engage in experimentation.

Our results confirm that reelection concerns reduce the probability of experimentation during the AFDC and TANF period. We also show that governors who face binding time limits and thus worry less about electoral support behave systematically different from governors who face reelections. We further find that ideological differences among policy-makers influence experimentation; yet, our main result that electoral incentives matter for policy experimentation holds even after accounting for ideological differences. Finally, we show that none of the alternative factors we consider – voter demand and preferences for redistribution, the influence of the legislature, and learning from other states – invalidates our conclusion that electoral incentives are important to understand the dynamics of welfare policy-making in the states.

References

- Alesina, A. and A. Drazen (1991). Why are stabilizations delayed? *American Economic Review* 81(5), 1170–1188.
- Alt, J., E. B. de Mesquita, and S. Rose (2011). Disentangling accountability and competence in elections: Evidence from U.S. term limits. *Journal of Politics* 73(1), 171–186.
- Arellano, M. and S. Bond (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *Review of Economic Studies* 58, 277–297.
- Barrilleaux, C., T. Holbrook, and L. Langer (2002). Electoral competition, legislative balance, and American state welfare policy. *American Political Science Review* 46(2), 415–427.
- Bernecker, A. (2016). Divided we reform? Evidence from US welfare policies. *Journal of Public Economics* 142, 24–38.
- Berry, F. S. and W. D. Berry (1990). State lottery adoptions as policy innovations: An event history analysis. *American Political Science Review* 84(2), 395–415.
- Berry, F. S. and W. D. Berry (1992). Tax innovation in the states: Capitalizing on political opportunity. *American Journal of Political Science* 36(3), 715–742.
- Berry, F. S. and W. D. Berry (2007). Innovation and diffusion models in policy research. In P. A. Sabatier (Ed.), *Theories of the Policy Process*, second edition, pp. 223–260. Boulder, CO: Westview.
- Berry, W. D., E. J. Ringquist, R. C. Fording, and R. L. Hanson (1998). Measuring citizen and government ideology in the American States, 1960–93. *American Journal of Political Science* 42(1), 327–348.
- Besley, T. (2007). Principled agents?: The political economy of good government. *Oxford University Press*.
- Besley, T. and A. Case (1995a). Does electoral accountability affect economic policy choices? Evidence from gubernatorial term limits. *Quarterly Journal of Economics* 110(3), 769–798.
- Besley, T. and A. Case (1995b). Incumbent behavior: Vote-seeking, tax-setting, and yardstick competition. *American Economic Review* 85(1), 25–45.
- Besley, T. and A. Case (2003). Political institutions and policy choices: Evidence from the United States. *Journal of Economic Literature* 41(1), 7–73.
- Blank, R. M. (2002). Evaluating welfare reform in the United States. *Journal of Economic Literature* 40, 1105–66.
- Callander, S. (2011a). Searching and learning by trial and error. *American Economic Review* 101(6), 2277–2308.

- Callander, S. (2011b). Searching for good policies. *American Political Science Review* 105(04), 643–662.
- Callander, S. and B. Harstad (2015). Experimentation in federal systems. *Quarterly Journal of Economics* 130(2), 951–1002.
- Callander, S. and P. Hummel (2014). Preemptive policy experimentation. *Econometrica* 82(4), 1509–1528.
- Center for Economic and Policy Research (2012). *Current Population Survey ORG Uniform Extracts*, Volume Version 1.7. Washington, D.C.
- Ciccone, A. (2004). Resistance to reform: Status quo bias in the presence of individual-specific uncertainty: Comment. *American Economic Review* 94(3), 785–795.
- Coate, S. and S. Morris (1999). Policy persistence. *American Economic Review* 89(5), 1327–1336.
- Council of Economic Advisors (1999). The effects of welfare policy and economic expansion on welfare caseloads: an update. Technical report, Executive Office of the President of the United States.
- Crouse, G. (1999). Implementation of major changes to welfare policies, 1992–1998. http://www.aspe.hhs.gov/hsp/Waiver-Policies99/policy_CEA.htm (accessed on February 7, 2014).
- de Janvry, A., F. Finan, and E. Sadoulet (2012). Local electoral incentives and decentralized program performance. *Review of Economist and Statistics* 94(3), 672–685.
- Dewan, T. and R. Hortala-Vallve (2017). Electoral competition, control and learning. *British Journal of Political Science*.
- Dewatripont, M. and G. Roland (1992). Economic reform and dynamic political constraints. *Review of Economic Studies* 4, 703–730.
- Drazen, A. and V. Grilli (1993). The benefit of crises for economic reforms. *American Economic Review* 83(3), 598–607.
- Fang, H. and M. P. Keane (2004). Assessing the impact of Welfare Reform on single mothers. *Brookings Papers on Economic Activity* 1, 1–95.
- Fernandez, R. and D. Rodrik (1991). Resistance to reform: Status quo bias in the presence of individual-specific uncertainty. *American Economic Review* 81(5), 1146–1155.
- Ferraz, C. and F. Finan (2011). Electoral accountability and corruption in local governments: evidence from audit reports. *American Economic Review* 101(4), 1274–1311.
- Figlio, D., V. Kolpin, and W. Reid (1999). Do states play welfare games? *Journal of Urban Economics* 46, 237–254.
- Fu, Q. and M. Li (2014). Reputation-concerned policy makers and institutional status quo bias. *Journal of Public Economics* 110, 15 – 25.
- Gagliarducci, S. and T. Nannicini (2013). Do better paid politicians perform better? Disentangling incentives from selection. *Journal of the European Economic Association* 11(2), 369–398.
- Galasso, V. and T. Nannicini (2011). Competing on good politicians. *American Political Science Review* 105(01), 79–99.
- Gelbach, J. (2017). *Uncontrolled Experiments from the Laboratories of Democracy: Traditional Cash Welfare, Federalism, and Welfare Reform*. Cheltenham, UK: Edward Elgar.
- Grogger, J. and L. A. Karoly (2005). *Welfare Reform: Effects of a Decade of Change*. Harvard University Press.
- Harrington, J. (1993). Economic policy, economic performance, and elections. *American Economic Review* 83(1), 27–42.
- Harvey, C., M. J. Camasso, and R. Jagannathan (2000). Evaluating welfare reform waivers under Section 1115. *Journal of Economic Perspectives* 14, 165–188.
- Haskins, R. (2006). *Work over Welfare: The Inside Story of the 1996 Welfare Reform Law*. Brookings Institution Press.
- Holbrook, T. and E. Van Dunk (1993). Electoral competition in the American States. *American Political Science Review* 87(4), 955–962.
- Holmström, B. (1982). Managerial incentive problems - A dynamic perspective. In *Essays in Economics and Management in Honor of Lars Wahlbeck*. Swedish School of Economics.
- Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *Review of Economic Studies* 66(1), 169–182.

- Jain, S. and S. W. Mukand (2003). Redistributive promises and the adoption of economic reform. *American Economic Review* 93(1), 256–264.
- Kaplan, T. (2000). Wisconsin works. In S. F. Liebschutz (Ed.), *Managing Welfare Reform in Five States: The Challenge of Devolution*. Rockefeller Institute.
- Karch, A. (2007). *Democratic Laboratories: Policy Diffusion among the American States*. University of Michigan Press.
- Klarner, C. (2003). Measurement of partisan balance of state government. *State Politics and Policy Quarterly* 3, 310–319.
- Koerper, K. (1996). Weekly tracking update - welfare reform: Section 1115 Waiver Activity. Office of Planning, Research and Evaluation.
- Kotakorpi, K. and P. Poutvaara (2011). Pay for politicians and candidate selection: An empirical analysis. *Journal of Public Economics* 95(7–8), 877 – 885.
- Lee, D. S., E. Moretti, and M. J. Butler (2004). Do voters affect or elect policies? Evidence from the U. S. House. *Quarterly Journal of Economics* 119(3), 807–859.
- Leip, D. (2012). Dave Leip’s Atlas of U.S. Presidential Elections. <http://www.uselectionatlas.org>. (accessed on February 7, 2014).
- Lieberman, R. C. and G. M. Shaw (2000). Looking inward, looking outward: The politics of state welfare innovation under devolution. *Political Research Quarterly* 53(2), 215–240.
- List, J. A. and D. M. Sturm (2006). How elections matter: Theory and evidence from environmental policy. *Quarterly Journal of Economics* 121(4), 1249–1281.
- Majumdar, S. and S. W. Mukand (2004). Policy gambles. *American Economic Review* 94(4), 1207–1222.
- Martinez-Bravo, M., G. Padro I Miquel, N. Quian, and Y. Yao (2014). Elections in China. Technical report, NBER working paper No. 18101.
- Mead, L. M. (2004). *Government Matters. Welfare Reform in Wisconsin*. Princeton University Press.
- Moffitt, R. A. (2002). Welfare programs and labor supply. In A. J. Auerbach and M. Feldstein (Eds.), *Handbook of Public Economics*, Chapter 34, pp. 2393–2430. Elsevier.
- Moffitt, R. A. (2008). Welfare reform: The US experience. *Swedish Economic Policy Review* 14 (2), 11–54.
- Oates (1999). An essay on fiscal federalism. *Journal of Economic Literature* 37(3), 1120–1149.
- Persson, T. and G. E. Tabellini (2000). *Political Economics: Explaining Economic Policy*. MIT Press.
- Poole, K. T. and H. Rosenthal (1984). The polarization of american politics. *Journal of Politics* 46, 1061–79.
- Poole, K. T. and H. Rosenthal (1997). *Congress: A Political-Economic History of Roll Call Voting*. New York: Oxford University Press.
- Prato, C. and S. Wolton (2018). Rational ignorance, populism, and reform. *European Journal of Political Economy* 55, 119–135.
- Reintsma, M. (2007). *The Political Economy of Welfare Reform in the United States*. Edward Elgar.
- Rogoff, K. (1990). Equilibrium political budget cycles. *American Economic Review* 80(1), 21–36.
- Saavedra, L. (2000). A model of welfare competition with empirical evidence from AFDC. *Journal of Urban Economics* 47, 248–279.
- Smart, M. and D. M. Sturm (2013). Term limits and electoral accountability. *Journal of Public Economics* 107, 93–102.
- Soss, J., S. F. Schram, T. P. Vartanian, and E. O’Brien (2001). Setting the terms of relief: Explaining state policy choices in the devolution revolution. *American Journal of Political Science* 45(2), 378–395.
- Strumpf, K. S. (2002). Does government decentralization increase policy innovation? *Journal of Public Economic Theory* 4(2), 207 – 241.
- United States Census Bureau (2011). Statistical Abstract of the United States: 2012 (131st edition).
- Urban Institute (2015). Welfare Rules Database. Washington, D.C. (accessed on May 7, 2015).
- U.S. Department of Health and Human Services, Administration for Children and Families (1998-2009). TANF Annual Reports to Congress, I-VIII. Technical report, GAO.
- U.S. Department of Homeland Security (2011). *Yearbook of Immigration Statistics*. Washington, D.C.: Office of Immigration Statistics.

- Walker, J. L. (1969). The diffusion of innovations among the American States. *American Political Science Review* 63(3), 880–899.
- Weaver, K. (2000). *Ending Welfare as we Know it*. Brookings Institution.
- Weissert, C. S. (2000). *Learning from Leaders: Welfare Reform Politics and Policy in five Midwestern States*. Rockefeller Institute Press.
- Wheaton, W. C. (2000). Decentralized welfare: Will there be underprovision? *Journal of Urban Economics* 48, 536–555.
- Willems, T. (2013). Political accountability and policy experimentation: Why to elect left-handed politicians? *Mimeo. Princeton University*. mimeo.
- Winston, P. (2002). *Welfare Policymaking in the States: The Devil in Devolution*. Georgetown University Press.

Online Appendix

A Theory

Proofs

Proof of Proposition 1. The proof of Proposition 1 follows from the optimization problem of the incumbent politician. A high-quality incumbent politician knows the state of the world. At $T = 0$, the incumbent chooses to reform or not (with the option to revert back to a_S if the reform is undertaken) in order to maximize equation (1).

In state S , no gain can be obtained from the reform ex post: the ex-post welfare of the reform in state S is either $-c$ or $-\frac{c}{2}$ if reverted back to a_S . Therefore, the reform should not be undertaken from a welfare perspective. Since the ex-post welfare of a reform in state S is negative, the probability of reelection (from the reelection rule) in state S is (weakly) positive if and only if no reform is undertaken. The optimal decision of the high-quality incumbent is then to stick with the status quo policy a_S independently of γ and λ .

In state N , undertaking the reform maximizes ex-post welfare since there are positive welfare gains ex post with probability 1, i.e. the ex-post welfare with a reform in state N is equal to $\Delta - c$, which is strictly positive from (A1), and is never reversed independently of the realization of the benefit in the short term. In addition, the probability of reelection is also maximized by implementing the reform since a positive ex-post welfare ensures reelection with probability 1, which is (weakly) greater than λ for any $\lambda \in [0, 1]$.

Proofs of Proposition 2 and Corollary 1. The proof of Proposition 2 directly follows from comparing the politician's payoffs in case of reverting the reform and of continuing the reform after a short-term failure. Reverting the reform is at least as good as continuing it if and only if

$$\gamma \left(\frac{-c}{2} \right) \geq \gamma (\beta \Delta - c) + (1 - \gamma) \beta,$$

or

$$-\beta \geq \gamma \left(\beta \Delta - \frac{c}{2} - \beta \right).$$

Since the term in parenthesis is negative from (A3), the last inequality can be rewritten to obtain the threshold γ^* in equation (2). The proof of Corollary 1 directly follows from differentiating equation (2) with respect to γ and Δ , respectively.

Proofs of Proposition 3 and Corollary 2. The proof of Proposition 3 directly follows from the comparison of the politician's payoff if she undertakes the reform (for the two cases where she will or will not revert back to the status quo in the face of short-term failure) with the payoff from sticking with the status quo and running the next election on her initial reputation instead.

The proof of Corollary 2 follows directly from differentiating equations (3) and (4) with respect to Δ , γ and λ respectively.

Proof of Proposition 4. The proof of Proposition 4 directly follows from the comparison of the behavior of the politicians which believe that the reform is the right thing

to do with different probabilities.

Extension of the theoretical framework: Reputation, experimentation, and ex-post welfare

We extend our model to allow for the possibility that voters derive some welfare from the reputation of the politician even when an experiment is implemented.

Following the spirit of retrospective voting models, we now assume that the total cost of enacting and continuing the new policy initiative is $c(\lambda) \geq 1$, which is incurred independently of the success of the reform. The parameter λ still represents the reputation of the politician among the electorate with $\lambda \in [0, 1]$. The function $c(\cdot)$ is assumed to be strictly decreasing in λ ($c'(\cdot) < 0$). Hence, voters prefer to have a high-reputation politician in office as it reduces the cost of experimentation. This assumption captures that ex-post welfare now depends on the benefit from the experimentation plus the governor's reputation in a simple way. The total cost $c(\lambda)$ still consists of short- and long-run costs; hence, $\frac{c(\lambda)}{2}$ has to be paid when the reform is implemented and again if the reform is maintained in the long term. If the reform is reverted back to a_S instead, then the long-term cost is not incurred. The benefit of a reform that is reversed is zero, however.

The incumbent politician still cares about both welfare of her constituency as well as her own future electoral prospects. The objective function of the incumbent is

$$\gamma(\text{Welfare}) + (1 - \gamma)(\text{Probability of Reelection}), \quad (7)$$

where $\gamma \in [0, 1]$ is the relative weight on welfare as before. In contrast to our model in Section 4, the reputation of the governor now enters the "Welfare" term through the reduction of the cost of experimentation.

We still assume that the representative voter follows a simple reelection rule: reelect the incumbent with probability 1 if ex-post welfare is positive; never reelect the incumbent if ex-post welfare is negative; and reelect the incumbent with probability λ if ex-post welfare is 0.

We also make the following assumptions:

$$p\Delta - c(0) > 0, \quad (A1')$$

$$pq\Delta - \frac{c(0)}{2}(1 + pq) > 0, \quad (A2')$$

and

$$\beta\Delta - \frac{c(1)}{2} < 0, \quad (A3')$$

where $\beta = \frac{p(1-q)}{1-pq} < p$ is the probability that the state of the world is N conditional on a short-term failure of the reform.

Assumption (A1') implies that a reform that is continued has positive expected welfare gains ex ante even for the politician with the highest cost ($\lambda = 0$). Assumption (A2') says that undertaking the reform and reverting back to the status quo in the face of short-term failure has positive expected benefits ex ante even for the politician with the highest cost ($\lambda = 0$). Finally, assumption (A3') implies that, if the benefit does not realize in the short term, the expected future benefit from continuing with the reform is negative even for the politician with the lowest cost ($\lambda = 1$).

Benchmark: Policy Choices of a High-quality Politician. The benchmark (Proposition 1) remains unaffected by the change in the specification of the cost function.

Policy Choices of a Low-quality Politician. We first focus on the decision to revert an experiment.

Proposition A.1. *A low-quality incumbent reverts back to the status quo in face of a short-term failure of the reform if and only if $\gamma \geq \gamma^*(\lambda)$, where*

$$\gamma^*(\lambda) := \frac{1}{1 + \frac{c(\lambda)}{2\beta} - \Delta}. \quad (8)$$

A low-quality politician who observes the short-term failure of a reform does not know whether the reform will be successful in the long term. This effect pushes her to revert back to the status quo to increase expected welfare. However, reelection concerns pull her towards sticking with the reform. Proposition A.1 disentangles these two conflicting forces as in Proposition 2 in the paper.

Interestingly, given a welfare weight γ , politicians are more likely to reverse an experiment in the case of short-term failure when reputation decreases ($\frac{\partial \gamma^*(\lambda)}{\partial \lambda} > 0$). When an experiment is implemented, the politician's reputation conflicts with the welfare of the voters because reputation now decreases the cost of the reform ex post and thus mitigates the cost of failure. We obtain a first testable hypothesis from our extension: conditional on experimenting, politicians are more likely to revert as reputation decreases.

We now analyze whether an uninformed politician who cares about both welfare and reelection ($\gamma \in [0, 1]$) undertakes the reform.

Proposition A.2. *(I.) When $\gamma \geq \gamma^*(\lambda)$, then a low-quality incumbent undertakes a reform if and only if*

$$\gamma \left(pq\Delta - \frac{c(\lambda)}{2}(1 + pq) \right) + (1 - \gamma)pq \geq (1 - \gamma)\lambda. \quad (9)$$

(II.) When $\gamma < \gamma^(\lambda)$, then a low-quality incumbent undertakes a reform if and only if*

$$\gamma(p\Delta - c(\lambda)) + (1 - \gamma)p \geq (1 - \gamma)\lambda. \quad (10)$$

Proposition A.2 follows from a comparison of the politician's payoff if she undertakes the reform (for the two cases where she will or will not revert back to the status quo in the face of short-term failure) with the payoff from sticking with the status quo and running the next election on her initial reputation instead.

Interestingly, the extension delivers new insights on the role of reputation for experimentation: first, if the welfare weight is zero (purely office-motivated, i.e. $\gamma = 0 < \gamma^*(\lambda)$) the results in the paper remain valid: a low-quality incumbent who is purely office-motivated implements the reform if and only if $p \geq \lambda$. However, as soon as there is a positive welfare weight, the reputation is now both on the left and right hand side of equations (9) and (10). As a result, a governor's reputation has now potentially a non-monotone effect on experimentation. Whereas in Section 4 an incumbent is more likely to experiment when λ decreases, the effect in the extended model depends on the curvature of the cost function.

In an electorate where the cost function is very concave in λ , an increase of reputation will make the politicians more likely to experiment for very high levels of reputation. At the same time it will make politician with intermediate reputation less likely to experiment. If the cost function is a very convex function of λ , a politician with little reputation will become more likely to experiment as their reputation increases, whereas this will not be the case for a politician with very high reputation.

Hence, a second testable hypothesis from this extension is a non-monotone effect of reputation on the likelihood to experiment: politicians with very low, intermediate, and very high reputation have potentially different experimentation behavior.

B Data

B.1 Measuring Policy Experimentation during the TANF period

To measure policy experimentation and reversals for the post-1996 period, we rely on the Welfare Rules Database by the Urban Institute. The database provides detailed information about states' TANF policies obtained from caseworker manuals and regulations, which are typically more detailed and up-to-date than the official plans state submit periodically to the federal government. While the database contains hundreds of rules on eligibility, benefit calculation, and many other aspects of welfare reform, we focus on a set of rules in policy areas that were at the center of the public and political debate surrounding welfare reform (see Table A1 for a list of the rules and how we code policy experiments and policy reversals).

Family caps. Under AFDC, benefit levels increased with family size. Under a family cap, additional benefits that an assistance unit would receive for a child born into the benefit unit while on welfare would be capped. Some states provide a percentage of the increase to the unit, while others provide no additional funds to the unit for the additional child. In several states, a family is never able to regain benefits for a capped child, even after the case has been closed for a period of time. In others, a family cap can be removed (and hence, the child can be included in the benefit computation should the family apply for assistance again) if the assistance unit remains off welfare for some time. We code a binary indicator equal to one if a state has a family cap and zero if not.

Work requirements. Under AFDC, states could require recipients to participate in the Job Opportunities and Basic Skills Training (JOBS) program, which provided education, training, and work experience activities. However, many individuals were exempt from these requirements (because of age, illness or having a small child). Under TANF, states require adults heading an assistance unit to perform some type of work-related activity. Work programs vary widely from state to state in terms of who must work, how much work is required, and what activities are considered work. The first rule defines the minimum number of hours a recipient must participate in work-related activities. The hours requirements vary from a mere effort to find a job up to full-time employment. The second rule defines whether the work requirement applies after several months of benefit receipt or by the time of application or approval. The third rule indicates whether there is a time limit of benefit receipt if a parent fails to work at least 20 hours per week in a regular job after a certain number of assistance months. The fourth rule counts the number of work exemptions due to, for example, pregnancy, disability or caring for a young child or elderly person.

Sanctions. Under AFDC, sanctions were sometimes imposed if the adult heading an assistance unit did not comply with the JOBS program; in that case, the adult portion of the benefit was not paid out to the assistance unit. Under TANF, states now require household heads to perform some type of work-related activity sooner or later. If a benefit unit does not comply with these requirements, states can impose drastic sanctions. The first policy rule characterizes the initial sanction if a benefit unit fails to comply with the work requirements for the first time. The initial sanction varies from reduction of 25% or less to a suspension of the full family benefit. The other three rules characterize the severity of the worst sanction that can be imposed. The second rule defines the severity of the worst sanction varying from less than 25% of the benefit to a suspension of the full family benefit and even case closure. The third rule defines the duration of the worst sanction ranging from until the unit complies with the requirement to a permanent suspension of the family benefit. The final rule defines whether a unit has to reapply (or not) after the worst sanction has been imposed.

Work-requirement time limits. Under AFDC, families were entitled to receive benefits as long as they met the eligibility requirements. Under TANF, many states imposed both intermittent and lifetime time limits. The first rule characterizes the number of months an assistance unit can receive benefits without interruption (many states impose 24 months). The second rule defines how benefits are reduced when the assistance unit reaches the intermittent time limit. The loss of benefits might be just for the adult members or for the entire assistance unit. The third rule defines the duration of the lifetime limit ranging from no time limit to only 24 months. If states wish to extend benefits beyond the federal time lifetime limit of 60 months, they have to use their own state funds to finance it. The final rule defines whether the state allows for any extensions to the lifetime limit or not.

B.2 Politics and ideology measures

Governor characteristics. We collect biographies of U.S. governors from the website of the National Governor's Association and data provided by David J. Andersen from the Eagleton Institute of Politics at Rutgers University. We code information on the age and educational degree of the governor. To measure a governor's competence whether an experiment is called for, we calculate her political experience prior to becoming governor. Specifically, the variable measures the number of years between the first electoral office (such as member of the State Senate or Attorney General) and the inauguration as governor. We then construct a dummy variable equal to one if a governor has prior political experience below the median and zero otherwise. Information on gubernatorial election results (especially vote margins) and term limits are taken from List and Sturm (2006); and Leip (2012).

Ideology. Voter and governor ideology is based on ideology ratings of the state's congressional delegation, the American for Democratic Action (ADA) rating and the AFL/CIO's Committee on Political Education (COPE) rating (Berry et al., 1998). Berry et al. (1998) assign an ideology rating to the citizens of each congressional district by using an average of the score of the congressional member and his or her election opponent weighted by the number of votes the score received. Zero denotes the most conservative and 100 the most liberal. They then generate a state-wide measure by averaging over all congressional districts. The measure of governor ideology is constructed by assigning to the governor the ratings of the members of Congress from their party. Updates of these

ideology data are available at <http://rcfording.wordpress.com/state-ideology-data/>. Alternatively, we use the Cook Partisan Voting Index to measure voter ideology obtained from <https://cookpolitical.com/pvi-map-and-district-list>.

State Legislature. Data for the composition of the state legislature, state competitiveness, the party of the governor and indicators for a divided government are obtained from Klarner (2003) and updated from <http://www.indstate.edu/polsci/klarnerpolitics.htm>. Polarization in the state legislature is calculated as $|\text{democratic seat share} - 50\%|$ for the state senate and house respectively. The divided government indicator is equal to one if the governor belongs to a different party than the majority of legislators in either the state senate or the state house. To measure state competitiveness, we use the Holbrook-Van Dunk index calculated from district-level returns to state legislative elections (Holbrook and Van Dunk, 1993).

B.3 State Demographics and Other Controls

Demographics. Population size, the number of Blacks and the age structure are taken from the United States Census Bureau (2011). The size of the immigrant population refers to the number of legal immigrants admitted by state of intended residence and is taken from Fang and Keane (2004) for 1970 to 2002 and updated using the Yearbook of Immigration Statistics (U.S. Department of Homeland Security, 2011). Personal income per capita is taken from the website of the Bureau of Economic Analysis. The unemployment rate for 1960 to 1998 is taken from the website of the Bureau of Labor Statistics. Our measure of income inequality is the ratio of the 90th over the 10th percentile of total household income calculated from the March Current Population Survey (Center for Economic and Policy Research, 2012). All income variables are deflated by the urban consumer price index with base year 2002. Unmarried birth refers to the % of births to unmarried women per 1,000 unmarried women aged 15-44 years by state of residence. For the years 1992 to 2003, the data are available from Table 8.3 in the TANF Annual Reports to Congress (U.S. Department of Health and Human Services, Administration for Children and Families, 2009). For earlier and later years, we obtain the data from the Center for Disease Control and Prevention (CDC) and the Vital Statistics.

State AFDC and TANF spending. To measure the potential gains from experimentation, we use state-level AFDC and TANF spending. These data come from Paul Ehmann at the U.S. Census Bureau. All fiscal variables are then converted into real measures using the urban consumer price index (with years 1982-1984=100) from the Bureau of Labor Statistics.

Spillover effects across States. To analyze spillover effects across states, we define a “neighbor” along three different dimensions. The first variable (geographic neighbors) codes the extent of experimentation for all states sharing a common border with the current state. The second variable (similar population sizes) calculates experimentation in states with a population size similar to that of the current state. For the calculation, we use the following ten bands based on population size in 1978: (CA NY TX PA IL), (OH MI FL NJ MA), (NC IN GA VA MO), (WI TN MD LA MN), (WA AL KY CT SC), (IA OK CO AZ OR), (MS KS AR WV NE), (UT NM ME RI HI), (ID NH MT NV SD) and (ND DE VT WY AK). The third measure (ideologically aligned governor neighbors) takes the average number of experiments of the states sharing a common border with the current state if the governor belongs to the same party as the present governor of the current state.

Table 1: Summary Statistics

	Mean	Std. Dev.
# Policy Experiments Waiver Period	0.20	0.58
# Policy Experiments TANF Period	0.50	0.91
# Policy Reversals	0.11	0.35
Governor's Past Vote Margin	9.84	9.38
Governor Lane Duck	0.27	0.44
Governor Age	53.59	7.90
Republican Governor	0.46	0.50
Governor Ideology	53.41	17.60
Citizens Ideology	47.68	16.69
Cook Partisan Voting Index (PVI)	-1.15	9.46
Divided Government	0.52	0.50
Democratic Seat Share Upper House	0.57	0.19
Democratic Seat Share Lower House	0.57	0.18
Polarization Nominate Upper House	0.16	0.12
Polarization Nominate Lower House	0.15	0.12
Competitiveness Index	41.14	12.64
AFDC/TANF Spending (in US\$ millions)	506.33	993.58
Growth Rate in AFDC/TANF Spending	-0.02	0.26
Population (/1000)	4.74	5.28
% Population 65+	11.83	2.35
%Population 5-17	19.58	4.30
% Black Population	9.80	9.73
% Immigrant Population	1.62	1.99
% Unmarried Births	28.79	10.37
Income Inequality (P90/P10)	7.97	1.69
Geographic Neighbors' Experiments (t-1)	0.29	0.45
Pop. Size Neighbors' Experiments (t-1)	0.25	0.40
Ideological Neighbors' Experiment (t-1)	0.08	0.18

Notes: For details on the welfare experimentation and reversal measures, see the data appendix. Governor lame duck is equal to one if the governor cannot run for reelection. The past electoral support measures the winner's distance in votes to the runner-up in the last gubernatorial election. The ideology measure for the government is calculated from Berry et al. 1998 and ranges from zero for most conservative to 100 for most liberal, see the data appendix for details. The citizen ideology measure is also taken from Berry et al. (1998). Divided Government is equal to one if the party of the governor is different from the party of the majority of legislators in either the state's lower or upper house. The polarization variables are calculated as difference in the party's Nominate scores. The competitiveness index is the vote share of winning candidates in state legislative elections over the current and past 3 years (Holbrook and van Dunk, 1993). Experiments by geographic neighbors refer to the average number of experiments for adjacent states. Experiments by population size neighbors refer to the average number of experiments for states of similar population size. Experiments by ideological neighbors count the average number of experiments in adjacent states that have a governor of the same party.

Table 2: Policy Experimentation during the AFDC and TANF Period

	Main Effect		Differential Effect		Control for	
	Electoral Support (1)	Electoral Support (2)	for Lane Ducks (3)	for Lane Ducks (4)	Potential Savings (5)	Potential Savings (6)
Electoral Support Last Gubernatorial Election	-0.005** (0.002)	-0.005* (0.003)	-0.007*** (0.003)	-0.008** (0.003)	-0.006** (0.003)	-0.007** (0.003)
Lame Duck * Electoral Support			0.008* (0.004)	0.010** (0.005)	0.008 (0.005)	0.011** (0.006)
Lame Duck			-0.048 (0.068)	-0.061 (0.072)	-0.046 (0.077)	-0.080 (0.083)
Past AFDC/TANF Spending					0.000*** (0.000)	0.001*** (0.000)
State Characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Governor Age	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State-specific Linear Trends	No	Yes	No	Yes	No	Yes
Observations	1,376	1,376	1,376	1,376	1,286	1,286
R-squared	0.265	0.289	0.267	0.292	0.276	0.302

Notes : The dependent variable in all specifications is the number of policy experiments in a given state and year over the period 1978-2007. Columns (1) and (2) include the governor's vote margin in the last gubernatorial election (measured as the distance to the runner-up in percent). Columns (3) and (4) allow the effect of the vote margin to differ for governors who cannot be reelected ("lame ducks"). Columns (5) and (6) also control for lagged AFDC/TANF spending (measured in real US\$ million) as a proxy for the potential savings from experimentation. All specifications control for year and state fixed effects as well as governor age, the size of the state population and its age structure. Even columns also control for state-specific linear trends. Standard errors clustered at the state electoral cycle are shown in parentheses. *** p<0.01, ** p<0.05 and *p<0.1.

Table 3: Electoral Incentives and Policy Experimentation (Robustness)

	<u>Cubic</u>	<u>Decade</u>	<u>Competitiveness of Election</u>			<u>Nonlinear Elec. Support</u>		<u>Alt. Definitions</u>	<u>Lame Duck</u>	<u>Alt. Potential Savings</u>	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Electoral Support Last Gubernatorial Election	-0.007** (0.003)	-0.009** (0.004)	-0.008** (0.003)	-0.007** (0.003)	-0.007** (0.004)	-0.010 (0.007)		-0.007** (0.003)	-0.008* (0.005)	-0.009*** (0.004)	-0.007** (0.003)
Lame Duck * Electoral Support	0.011** (0.006)	0.011* (0.006)	0.012** (0.006)	0.011** (0.006)	0.012** (0.006)			0.012** (0.006)	0.012* (0.006)	0.011* (0.006)	0.011* (0.006)
Lame Duck	-0.080 (0.083)	-0.065 (0.087)	-0.094 (0.084)	-0.077 (0.083)	-0.093 (0.085)	0.035 (0.052)	0.040 (0.051)	-0.088 (0.086)	-0.101 (0.089)	-0.071 (0.082)	-0.082 (0.085)
Past AFDC/TANF Spending	0.001*** (0.000)	0.000** (0.000)	0.001*** (0.000)	0.000*** (0.000)	0.000*** (0.000)	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)		0.001*** (0.000)
Governor Change past Electoral Cycles			-0.032 (0.033)		-0.034 (0.034)						
Competitiveness of Elections				0.008 (0.006)	0.009 (0.006)						
Electoral Support Squared						0.0002 (0.000)					
Lame Duck * Electoral Support Squared											
Dummy (25th Perc. <Popularity< 75th Perc.)							-0.0465 (0.055)				
Dummy (Popularity > 75th Perc.)							-0.1044* (0.057)				
Growth in Past AFDC/Tanf Spending										0.156** (0.075)	
Observations	1,286	1,286	1,286	1,277	1,277	1,286	1,286	1,229	846	1,236	1,286
R-squared	0.302	0.333	0.302	0.302	0.303	0.300	0.299	0.300	0.378	0.299	0.303

Notes: The dependent variable in all specifications is the number of policy experiments in a given state and year over the period 1978-2007. The main independent variables and controls are the same as in Table 2 (see notes for details). Column (1) adds state-specific quadratic and cubic trends, column (2) state-specific decade dummies. Columns (3) and (5) include an indicator whether a state has changed its governor over the past two electoral cycles, while columns (4) and (5) include the Holbrook and Van Dunk index of competitiveness in state elections. Column (6) includes a quadratic term in electoral support (but no interaction with lame duck), while column (7) splits electoral support into low (reference), medium (75th-25th percentile) and high (above 75th percentile) support (but omits interaction effects with lame duck). Column (8) drops states with a one-term limit for state governors; column (9) keeps only states who have a term limit for state governors. Column (10) includes the growth rate in past AFDC/TANF spending and column (11) adds the share of Black population, the share of immigrants, the share of unmarried births and the 90-10 income ratio. Standard errors clustered by state electoral cycle are shown in parentheses. *** p<0.01, ** p<0.05 and *p<0.1.

Table 4: Governor Ideology and Policy Experimentation

	<u>Governor Ideology</u>		<u>Baseline with Ideology</u>	
	(1)	(2)	(3)	(4)
Electoral Support Last Gubernatorial Election			-0.007** (0.004)	-0.006* (0.004)
Lame Duck * Electoral Support			0.011* (0.006)	0.010* (0.006)
Lame Duck			-0.094 (0.084)	-0.077 (0.088)
Past AFDC/TANF Spending			0.001*** (0.000)	0.001*** (0.000)
Republican Governor	0.079** (0.037)		0.087** (0.040)	
Ideology Governor		-0.003** (0.001)		-0.003* (0.001)
Observations	1,387	1,269	1,265	1,148
R-squared	0.291	0.309	0.305	0.324

Notes: The dependent variable in all specifications is the number of policy experiments in a given state and year between 1978 and 2007. All main variables are defined as before (see notes to Table 2 for details). The ideology measure for the government in columns (2) and (4) is calculated from Berry et al. (1998) and ranges from zero for most conservative to 100 for most liberal. All specifications control for year and state fixed effects, state-specific linear trends and the same state characteristics as in Table 2. Standard errors clustered by state electoral cycle are shown in parentheses. *** p<0.01, ** p<0.05 and *p<0.1.

Table 5: Voter Preferences, the Demand for Welfare and Policy Experimentation

	<u>Voter Ideology</u>		<u>Preferences for Redistribution</u>	
	(1)	(2)	(3)	(4)
Electoral Support Last Gubernatorial Election	-0.007** (0.003)	-0.007** (0.003)	-0.007** (0.003)	-0.007** (0.003)
Lame Duck * Electoral Support	0.011** (0.006)	0.011** (0.006)	0.011** (0.006)	0.011** (0.006)
Lame Duck	-0.080 (0.083)	-0.082 (0.082)	-0.082 (0.085)	-0.083 (0.084)
Past AFDC/TANF Spending	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)
Citizen Ideology (Berry et al. 1998)	-0.001 (0.004)			
Cook Partisan Voting Index (PVI)		-0.008 (0.007)		
Preferences for Redistribution			0.054 (0.046)	0.099** (0.049)
Past Spending*Preferences			-0.000 (0.000)	-0.000 (0.000)
Observations	1,286	1,286	1,286	1,286
R-squared	0.302	0.302	0.303	0.304

Notes: The dependent variable in all specifications is the number of policy experiments in a given state and year between 1978 and 2007. All main variables are defined as before (see notes to Table 2 for details). The citizen ideology measure in column (1) is taken from Berry et al. (1998); see the data appendix for further details. The Cook Partisan Voting Index (PVI) in column (2) indicates the strength of voter support for the Democratic candidate (positive values) or the Republican candidate (negative values) in the last presidential election relative to the national average. Preferences for redistribution in column (3) is an indicator equal to one if respondents in a state have stronger preferences for redistribution than respondents in the median state (based on a General Social Survey question on a 5-point Likert scale). The preference measure in column (4) is also an indicator if respondents are more likely to say that their government spends too little on welfare relative to the median state (based on a 3-point scale "too little", "about right" or "too much"). For both measures, higher values indicate a preference for more redistribution. In columns (3)-(4), both measures are also interacted with past welfare spending. All specifications control for state and year fixed effects, state-specific linear trends and the same control variables as in Table 2. Standard errors clustered by state electoral cycle are shown in parentheses. *** p<0.01, ** p<0.05 and *p<0.1.

Table 6: Composition of Legislature and Policy Experimentation

	(1)	(2)	(3)	(4)
Electoral Support Last Gubernatorial Election	-0.008** (0.003)	-0.007** (0.003)	-0.007** (0.004)	-0.007** (0.004)
Lame Duck * Electoral Support	0.012** (0.006)	0.012** (0.006)	0.012** (0.006)	0.013** (0.006)
Lame Duck	-0.082 (0.083)	-0.081 (0.083)	-0.100 (0.084)	-0.103 (0.084)
Past AFDC/TANF Spending	0.001*** (0.000)	0.001*** (0.000)	0.000*** (0.000)	0.000*** (0.000)
Dem. Seat Share Upper House	-0.493* (0.277)			-0.466 (0.293)
Dem. Seat Share Lower House	-0.101 (0.450)			-0.154 (0.460)
Polarization Nominate Score House		-0.208 (0.337)		0.033 (0.353)
Polarization Nominate Score Senate		-0.288 (0.473)		-0.345 (0.489)
Divided Government			0.100** (0.040)	0.091** (0.043)
Observations	1,286	1,286	1,269	1,269
R-squared	0.303	0.302	0.306	0.308

Notes : The dependent variable in all specifications is the number of policy experiments in a given state and year between 1978 and 2007. All main variables are defined as before (see notes to Table 2 for details). Divided Government is equal to one if the party of the governor is different from the party of the majority of legislators in either the state's lower or upper house. The polarization variables are calculated as the difference in NOMINATE scores in the House and Senate. All specifications control for state and year fixed effects, state-specific linear trends and the same state characteristics as in Table 2. Standard errors clustered by state electoral cycle are shown in parentheses. *** p<0.01, ** p<0.05 and *p<0.1.

Table 7: Spillovers between States and Policy Experimentation

	(1)	(2)	(3)	(4)
Electoral Support Last Gubernatorial Election	-0.007** (0.003)	-0.007** (0.003)	-0.007** (0.003)	-0.007** (0.003)
Lame Duck * Electoral Support	0.011** (0.006)	0.011* (0.006)	0.011* (0.006)	0.011* (0.006)
Lame Duck	-0.076 (0.083)	-0.073 (0.083)	-0.081 (0.082)	-0.069 (0.083)
Past AFDC/TANF Spending	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)
Experiments by Similar-Sized States	0.110 (0.085)			0.125 (0.084)
Experiments by Geographic Neighbors		-0.197*** (0.072)		-0.202** (0.085)
Experiments by Neighbors with Aligned Governor			-0.195 (0.138)	-0.023 (0.166)
Observations	1,286	1,286	1,286	1,286
R-squared	0.304	0.306	0.303	0.309

Notes: The dependent variable in all specifications is the number of policy experiments in a given state and year between 1978 and 2007. All main variables are defined as before (see notes to Tables 2 and 3 for details). Experiments by similar sized states refer to the average number of experiments in states with similar population size (where states are divided into ten groups of similar population size), while geographic neighbors refer to the average number of experiments in adjacent states (sharing a boundary). An ideological neighbor is defined as a neighboring state (sharing a boundary) that has a governor from the same party. The variables for policy experiments in neighboring states are lagged by one year. All specifications control for year and state fixed effects as well as state-specific trends. We further include governor age, the population size and age structure of the state. Standard errors clustered by state electoral cycle are shown in parentheses. *** p<0.01, ** p<0.05 and *p<0.1.

Table 8: Policy Reversals in the TANF Period

	Policy Reversal			
	(1)	(2)	(3)	(4)
Past TANF Spending	-0.000* (0.000)	-0.000* (0.000)	-0.000 (0.000)	-0.000* (0.000)
Lame Duck		-0.000 (0.046)	-0.001 (0.047)	
Republican Governor			-0.075** (0.028)	
Electoral Support Last Gubernatorial Election				0.001 (0.003)
Observations	561	561	549	548
R-squared	0.150	0.150	0.158	0.150

Notes: The dependent variable in columns (1)-(4) is the number of policy reversals in a given state and year during the TANF period (1996-2007). Column (1) includes past state TANF spending (measured in US\$ million) as a measure of the potential gains from experimentation. Column (2) adds an indicator whether the governor cannot be reelected ("lame duck"). Column (3) includes whether the governor is a Republican or Democrat. Finally, column (4) includes the electoral support in the past gubernatorial election. All specifications control for state and year fixed effects as well as state-specific linear trends. Standard errors clustered at the state electoral cycle are shown in parentheses. *** p<0.01, ** p<0.05 and *p<0.1.

Figure 1: Policy Experimentation and Past Vote Margins by Lame Duck Governors

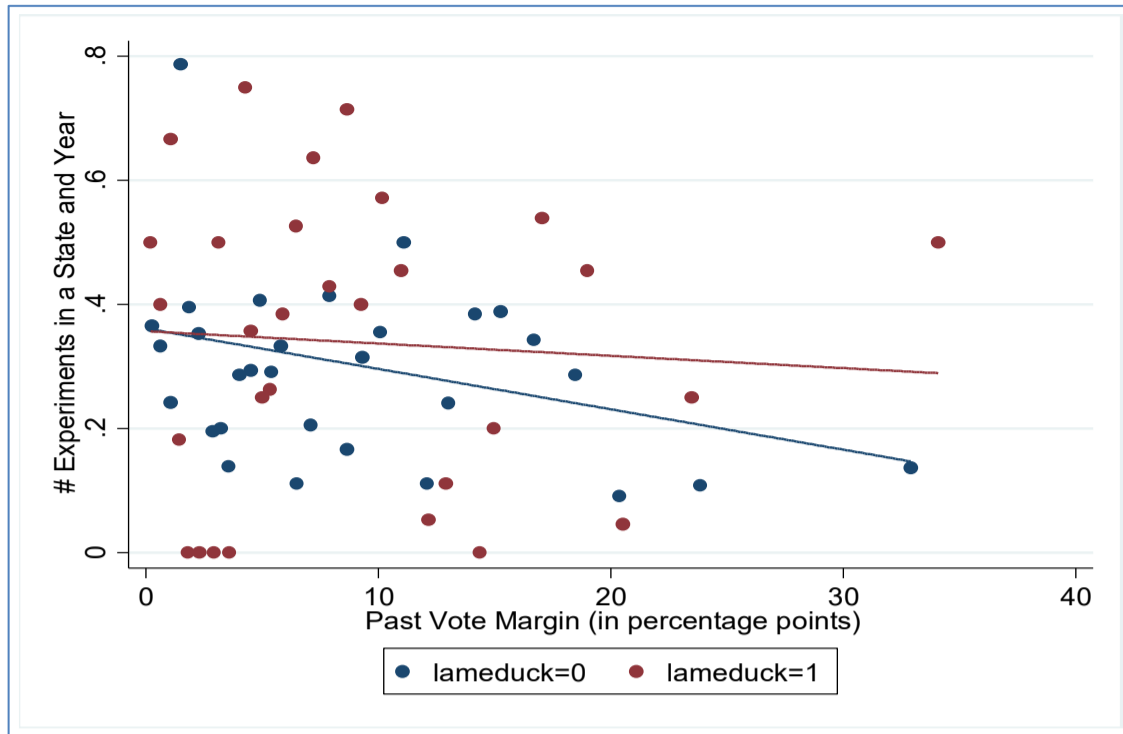


Figure 2: Policy Experimentation and Past Vote Margin by Governor Party

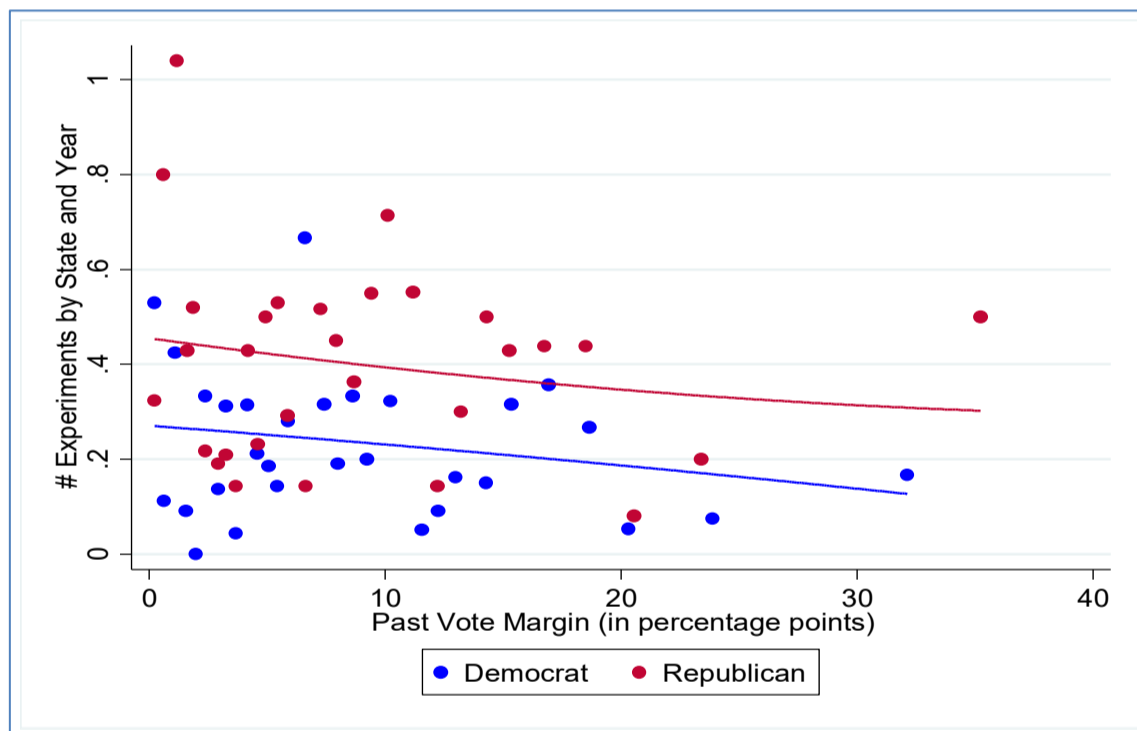


Table A1: Coding of Policy Experimentation and Reversals during the TANF Period (1996-2007)

Policy Rule	Description of Rule	Experimentation = 1	Reversal = 1
<u>Family Cap (1)</u>	Benefits do not increase if an additional child is born in family while receiving benefits.	State adopts family cap before or after 1996.	State abolishes family cap in 1997 or later.
<u>Work Requirements (3)</u>			
Hours Requirement	Minimum # of hours a recipient must participate in work-related activities.	Increase in minimum number of working hours required of the recipient.	State reduces or abolishes minimum number of working hours required of the recipient.
Work upon Enrollment	Work requirements apply at application, approval or upon benefit receipt (or only later).	State requires work requirement to be fulfilled at application, approval or upon benefit receipt.	State allows work requirement to be fulfilled later.
Time Limit to Work	Work at least 20 hours per week in an unsubsidized job after a certain period of benefit receipt.	State adopts a time limit to fulfill the requirement of working in a unsubsidized job for at least 20 hours.	State abolishes time limit to fulfill the requirement of working in an unsubsidized job for at least 20 hours.
<u>Time Limits (3)</u>			
Duration of Lifetime Limit	Maximum # of months an assistance unit can receive benefits over the lifetime.	Prior to 1997, if a time limit is adopted; after 1997, when a state adopts a time limit stricter than the 60 months federal time limit. In 1997, if a state adopts a time limit other than the federal time limit.	If a state abolishes the time limit or increases # of months.
Intermittent Time Limit	# months time an assistance units can receive benefits without interruption.	State reduces intermittent time limit.	State increases intermittent time limit.
Benefit Reduction after Intermittent Time Limit	How much benefits are reduced (adult portion or benefit unit) when assistance units hits the intermittent time limit.	State reduces benefits more after unit hits intermittent time limit.	State reduces benefits less after unit hits intermittent time limit.
<u>Sanctions (4)</u>			
Severity of Worst Sanction	How much benefits are reduced under worst sanction for non-compliance with work requirements.	State imposes a more severe penalty for non-compliance.	State reduces severity of worst sanction for non-compliance.
Duration of Worst Sanction	Length of most severe sanction for not complying with work requirements.	State increases duration of worst sanction.	State decreases duration of worst sanction.
Reapply	Whether the unit has to reapply after worst sanction for non-compliance has been imposed.	State adopts requirement that unit has to reapply after worst sanction has been imposed.	State abolishes requirement that unit has to reapply after worst sanction has been imposed.
Severity of Initial Sanction	Whether initial sanction for non-compliance is removal of full family benefit.	State adopts that full family benefit is removed as initial sanction.	State abolishes rule that full family benefit is removed as initial sanction.

Source : Welfare Rules Database (2015)

Table A2: Alternative Estimators of Policy Experimentation

	<u>Poisson</u> (1)	<u>Probit</u> (2)	<u>AB Estimator</u> (3)
Vote Margin Last Gubernatorial Election	-0.030** (0.012)	-0.002 (0.002)	-0.008*** (0.003)
Lame Duck * Gubernatorial Vote Margin	0.042** (0.017)	0.003 (0.003)	0.009* (0.005)
Lame Duck	-0.383* (0.203)	-0.015 (0.035)	-0.046 (0.065)
Past AFDC/TANF Spending	0.001* (0.001)	0.000** (0.000)	0.000 (0.000)
Experimentation (t-1)			-0.071** (0.028)
Observations	1,286	1,193	1,286
p-value of AR(2) residuals			0.440
p-value of AR(3) residuals			0.307

Notes : The table reports estimates from a Poisson model (in column (1)), a probit model (in column (2)) and the Arellano-Bond dynamic panel data estimator including a lagged dependent variable to account for unobserved heterogeneity (in column (3)). The AB estimator uses all included variables lagged 2 to 4 years as potential instruments. The dependent variable is the number of policy experiments in columns (1)-(3) and an indicator whether a state implements an experiment in a given year in column (2). All specifications include year fixed effects and the same state characteristics as in Table 2. The p-value tests the null hypothesis that the residuals are serially correlated of the specific degree: AR(2) or AR(3). *** p<0.01, ** p<0.05 and * p<0.1.

Table A3: Governor Heterogeneity in Competence and Experimentation

	<u>Less than a</u>		<u>Little Prior</u>	
	<u>Master's Degree</u>		<u>Political Experience</u>	
	(1)	(2)	(3)	(4)
Electoral Support Last Gubernatorial Election	-0.007** (0.003)	-0.008** (0.003)	-0.007** (0.003)	-0.007** (0.003)
Lame Duck * Electoral Support	0.011** (0.006)	0.011** (0.006)	0.011** (0.006)	0.012** (0.006)
Lame Duck	-0.080 (0.083)	-0.050 (0.108)	-0.080 (0.084)	-0.147 (0.100)
Past AFDC/TANF Spending	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)
Governor with Low Education/Experience	0.062 (0.051)	0.075 (0.055)	-0.025 (0.057)	-0.066 (0.067)
Lame Duck * Low Education / Experience		-0.060 (0.095)		0.150 (0.092)
Observations	1,286	1,286	1,284	1,284
R-squared	0.303	0.303	0.302	0.304

Notes : The dependent variable in all specifications is the number of policy experiments in a given state and year over the period 1978-2007. Columns (1) and (2) use an indicator whether a governor has low education, i.a. less than a Master's degree. Columns (3) and (4) include an indicator whether the governor has little prior political experience (below the median). Columns (2) and (4) interact both measures of governor ability with an indicator whether the governor is a lame duck. The other independent variables are defined as in previous tables. All specifications control for year and state fixed effects, state-specific trends and the same controls as in Table 2. Standard errors clustered by state electoral cycle are shown in parentheses. *** p<0.01, ** p<0.05 and *p<0.1.

Table A4: Descriptives on Policy Reversals

	# Experiments (1)	# Reversals (2)	% Same Elec. Cycle (3)	Mean # Years (4)	% Same Party (5)	% Same Governor (6)
<u>Family Cap</u>	22	5	0%	10.2	20%	0%
<u>Work Requirements</u>						
Hours Requirement	126	14	21%	5.4	50%	30%
Work upon Enrollment	43	2	0%	6.5	100%	50%
Time Limit to Work	3	2	0%	6	100%	50%
<u>Time Limits</u>						
Duration of Lifetime Limit	18	6	17%	3.7	100%	80%
Intermittent Time Limit	20	10	30%	5.4	50%	40%
Benefit Reduction after Intermittent Time Limit	53	3	33%	7.7	100%	33%
<u>Sanctions</u>						
Severity of Worst Sanction	55	1	0%	2	100%	100%
Duration of Worst Sanction	44	38	44%	4.3	78%	50%
Reapply	20	6	0%	6	67%	50%
Severity of Initial Sanction	22	0				

Notes : The table reports for the 11 policy rules (and 87 policy reversals in total) the percentage of reversals that occur within the same electoral cycle as the experiment; the mean number of years between experiment and reversal; the share of reversals that are implemented by the same party (Republican or Democrat) as the initial experiment; and the percentage of reversals that are implemented by the same governor as the experiment. The number of experiments of the duration of the worst sanction is lower than the number of reversals because many states initially applied the federal guidelines when adopting their state TANF policies. Later on, many states deviated from these rules which accounts for the relatively high number of reversals over our sample period.

Figure A1: Evolution of Welfare Waivers, 1978-1996




Figure A2: Evolution of Policy Experiments and Reversals, 1996-2007

