

The Wrong Kind of Learning: The Case of Subnational Recall Diffusion

Mariana Alvarado
University of Geneva

June 20, 2025

Abstract

The diffusion literature has traditionally considered learning as a mechanism through which successful policies spread. This paper builds on recent advances in this literature to show that under certain conditions learning can lead to the diffusion of normatively questionable practices. It uses subnational recall elections in Peru —the most intensive user of the recall worldwide— to present the first empirical evidence consistent with perverse learning and outlining the conditions under which it can occur. In particular, it shows that recalls diffuse spatially through a learning process in which information about the political viability of this institution is transmitted to the local politicians who are responsible for its repeated activation. These politicians learn to (ab)use participatory institutions to re-do elections regardless of voter’s preferences.

While long-running, the diffusion literature has witnessed an important re-emergence and transformation over the past twenty years. Early work in this tradition focused mainly on the diffusion of policies, understood as the result of a fact-based assessment of their policy effects by policymakers (Gilardi and Wasserfallen, 2019). More recent research has led to a broadening of the scope of diffusion on practically all fronts: diffusion is no longer limited to policies, driven exclusively by policymakers, resulting only from fact-based assessments, or based solely on policy effects. Political strategies, platforms and frames have also been found to diffuse (Graham et al., 2013; Van Hauwaert, 2019), as a result of the actions of parties, political elites more broadly and even citizens (Böhmelt et al., 2016; Pacheco, 2012; Martini and Walter, 2024; Walter, 2021), through the use of cognitive shortcuts and biases (Weyland, 2005; Meseguer, 2005), and often on the basis not of policy but of political effects (Gilardi, 2010). In terms of the mechanisms, while there has been a shift regarding how learning comes about (through the use of heuristics rather than rational bayesian learning), when it comes to what is learned, the assumption remains that learning is a way for beneficial or successful policies/strategies to spread (Bamert et al., 2015; Martini and Walter, 2024). This paper contributes to the ongoing expansion in the diffusion literature by presenting evidence of perverse learning and outlining the conditions under which this type of diffusion can occur.

Learning has traditionally been seen in the diffusion literature as a mechanism through which successful policies spread. Indeed, much of the early work on diffusion focused on the cross-state diffusion of policies in the U.S., and assumed that successful policies spread through a combination of policy innovation and careful evaluation (Karch, 2007). Recently, researchers have highlighted that political actors may learn not only about policy but also about political effects, and have noted that this may lead to the spread of bad policies (Gilardi and Wasserfallen, 2019). However, no empirical cases of this type of diffusion have been studied. In this paper, I present one such case of perverse learning: the learning of political strategies that may benefit individual politicians but do so at the cost of collective well-being. In so doing I focus on the diffusion of political strategies rather than policies and discuss why this distinction is important.

The case in question is the diffusion of subnational recall elections in Peru, the country considered to be the most intensive user of the recall worldwide (Welp, 2015). Recalls are a type of direct democracy mechanism that seeks to remove elected officials from office before the end of their term. Recurrent use of the recall to remove district mayors in Peru has been universally considered by local analysts as prejudicial to the functioning of democracy. In this context, recalls not only promote polarization, disrupt municipal governance and bolster citizen distrust of politicians but have also been found to reduce the quality of candidates (Artiles et al., 2021; Welp, 2013; Tuesta Soldevilla, 2014b).

The Peruvian setting provides a large number of observations (1634 districts) and a relatively long time series (1996-2014), allowing me to exploit variation in the presence

of recall elections within districts over time, to uncover the effect of neighbors' behavior and the dynamics of political competition on the probability a recall takes place. This constitutes the first rigorous analysis of the within-country diffusion of recalls. Results indicate that recalls do diffuse spatially, through a process of learning in which information about the political viability of this institution is transmitted to the local politicians who are responsible for its repeated activation. Despite optimistic views regarding the potential for direct democracy mechanisms to improve citizen engagement and satisfaction (Matsusaka, 2020; Whitehead, 2020), my results highlight the risk that these mechanisms may be coopted and abused by self-serving politicians, ultimately contributing to citizen apathy and distrust. And while it is difficult to generalize from a single country study, there are reasons to believe similar dynamics are at play in (very) different contexts.

The paper makes significant contributions to both the diffusion and the direct democracy literatures. In terms of the former, it illustrates the potential and outlines the conditions for negative or perverse diffusion dynamics. Indeed, while the possibility that negative practices may diffuse has sometimes been acknowledged, the diffusion literature has tended to see learning as a process through which efficient or successful policies expand, with insufficient consideration of how success is defined (and by whom). Moreover, by focusing on the diffusion of negative practices, this paper deals directly with the normative implications of diffusion, an aspect that has received little attention in the literature.

In terms of direct democracy, despite growing interest in the recall in recent years it is still the least studied of these mechanisms (Whitehead, 2018). The limited literature there is on it is characterized by (i) an ongoing (and long-running) debate regarding its theoretical benefits and drawbacks, along with (ii) a mostly descriptive analysis of its use in particular cases or countries. And while there is some discussion regarding potential diffusion in the adoption of recall regulations, the possibility that recall *use* might diffuse spatially has heretofore not been seriously considered.¹ This paper thus provides the first systematic analysis of the diffusion of recall practices, highlighting the importance of learning and outlining how it occurs. Much attention has been paid to the internal determinants of recall use in the literature but, as I show below, diffusion is also very important and can determine not only whether it gets used, but also how. In a scenario in which use of the recall is often associated with deteriorating political outcomes (Whitehead, 2018), a better understanding of the dynamics underlying its use is particularly important. Moreover, it adds new evidence to the literature on the within-country diffusion of direct-democratic practices, which has so far been largely U.S.-centered. As the use of these practices continues to expand among developing countries, this evidence

¹This is somewhat surprising considering that research on the diffusion of policies via the initiative has a long history, at least in the U.S. (see for example Seljan and Weller (2011)). However, part of the reason why there are no studies of diffusion in recall use in the U.S. seems to be the absence of a centralized database registering all recall elections taking place at the local level (Spivak, 2020).

becomes increasingly relevant.

The paper proceeds as follows. Section 1 introduces the diffusion literature, summarizing its main findings and highlighting how this paper contributes to it. Section 2 presents the institution of the recall, the theoretical debates regarding its potential effects and existing empirical findings. Section 3 describes the functioning and use of the recall in Peru, as well as existing research on its determinants and consequences. Section 4 introduces my hypotheses, focused on explaining how recalls diffuse, what it is that diffuses and who the actors driving this diffusion are. Section 5 presents my empirical strategy and data. Section 6 presents and discusses the results of my analyses. Section 7 describes robustness tests and section 8 concludes with a discussion of my findings and their implications.

1 The Politics of Diffusion

Diffusion can be defined as “any process where prior adoption of a trait or practice in a population alters the probability of adoption for the remaining non-adopters” (Strang, 1991, p. 325). It therefore implies that the behavior of individual units (typically countries, states or other subnational units) is affected by the behavior of other units. The way in which they are affected may vary—they may copy them, learn from them, or attempt to distinguish themselves from them—but the key point is that their behavior is a response or a result of the prior behavior of other units. Diffusion has typically been studied between neighboring units, but while this is the most common scenario, it is not necessary: countries may react to the behavior of their trade partners—or competitors—even if they are far away. The challenge then, is to distinguish true from spurious diffusion: a pattern of successive adoptions emerging from fully independent decisions as units respond in a similar way to similar conditions (Braun and Gilardi, 2006).

The diffusion literature has posited multiple mechanisms that might explain why the behavior of one unit may influence that of another, including emulation, learning, competition and coercion.² Emulation (sometimes also called imitation) happens when units simply copy what others have done, typically larger or more prestigious units that they aspire to resemble. Learning is considered to occur when successful policies are copied, as the experiences of others leads to a change in one’s beliefs regarding the potential effects of policy adoption (Meseguer, 2005). Accordingly, learning is usually evaluated in terms of whether it leads to the adoption of better or more efficient decisions (see for example Ellison and Fudenberg (1993) or Volden et al. (2008)). Competition happens when units are forced to adapt their behavior to what other units have done in

²There is some discussion as to whether coercion should be considered a diffusion mechanism (Maggetti and Gilardi, 2016). Some authors also consider additional mechanisms, such as common norms and taken-for-grantedness (Braun and Gilardi, 2006).

an effort to maintain their relative attractiveness. They can do so by adopting the same policy or a complementary one. And coercion happens when units are pressured into adopting a certain policy through the imposition of either costs or rewards.³

The study of diffusion in political science has almost exclusively focused on a particular form of diffusion: policy diffusion, which occurs when “policy choices in one unit are influenced by policy choices in other units” (Maggetti and Gilardi, 2016, p. 89). Indeed, the field developed largely as a result of early efforts to understand the cross-state diffusion of policies within the U.S. (Karch, 2007). However, in recent decades the study of cross-country diffusion has gained prominence in both comparative politics and international relations. As this literature has expanded, it has also gained in complexity. Indeed, as mentioned above, the policy diffusion literature has gone from adopting a technocratic view according to which policies spread because policymakers evaluate the policy implications of the actions of other units, to acknowledging the importance of a broader set of actors and dynamics. Not only policies but all kinds of things have been found to diffuse: institutions, political platforms and strategies, ideational frameworks, to name a few (Graham et al., 2013; Gilardi, 2012). This has two important consequences. The first is that policymakers are not the main actors driving diffusion processes anymore, but so are political parties, political elites more broadly (including policy opposers) and even citizens (Pollert and Mooney, 2022; Pacheco, 2012; Walter, 2021). The second is that when it comes to learning, not only policy effects are learned, but so are the political effects not only of policies, but of all kinds of political actions. This means that political parties may learn from the successful platforms of parties in other countries (Böhmelt et al., 2016), and that citizens may learn from the consequences of political decisions abroad (Walter, 2021). Moreover, the manner of learning has also evolved: while it was initially assumed that learning was a rational, fact-based process, a growing body of research documents the use of heuristics and the prevalence of cognitive biases (Weyland, 2005; Meseguer, 2005; Bamert et al., 2015).

This paper builds on and expands this growing literature by bringing these different aspects of diffusion together and outlining the conditions under which perverse forms of learning may take place. Indeed, prior research has highlighted the fact that diffusion is an intrinsically political process in which policymakers care not only about the policy effects of policies, but also about their political effects. One possible consequence of this, it is claimed, is that bad policies may spread if policymakers care more about the political effects of policies than about their policy effects (Gilardi and Wasserfallen, 2019). However, the fact that not only policies spread, and that diffusion is not only driven by policymakers, means that situations in which political actors care more (if not exclusively) about the political effects of their actions are likely to be quite prevalent. In fact, while

³For a detailed account of each of these mechanisms see Braun and Gilardi (2006) or Shipan and Volden (2008).

policymakers may be forced to at least consider the policy consequences of their actions, this is not the case for other types of political actors (such as office-seeking politicians). Under these conditions, all kinds of bad practices may diffuse. Perverse learning thus happens when politicians care mainly about winning office, as this narrow focus leads them to adopt strategies that will further their selfish interests regardless of the common good.⁴ While this is a general statement that I cannot test here and should be examined by future research, in what follows I offer an example of this logic at work.

In the sections that follow, I document an instance of perverse learning in which political entrepreneurs learn to abuse a participatory institution to the detriment of citizen trust in the political system. This has important consequences for the normative interpretation of learning by highlighting the fact that under certain conditions learning can lead to worse social outcomes as a result of politicians maximizing their short-term utility. It is worth noting that in addition to Gilardi and Wasserfallen (2019), Shipan and Volden (2012) also acknowledge that policy diffusion is not always beneficial. However, they focus on cases in which the wrong lessons are drawn from others' experiences either because they are not adapted to the new context or because information about effectiveness was not collected. This is different from the argument I am making here: I would argue politicians learn the right lesson (recalls are a useful strategy for re-doing elections), the problem is that —as in any social dilemma— their private interests are at odds with collective interests and thus lead to worse social outcomes. Shipan and Volden (2012) on the other hand do not consider the possibility that politicians may define success in terms of their narrow interests rather than the collective good.⁵ Finally, Dobbin et al. (2007) are one of the few to recognize that the aggregation of individual choices in the context of policy learning may not be socially optimal.

As such, this paper responds to demands that research on diffusion should expand its focus beyond policy adoptions and include a broader set of political motivations (Gilardi and Wasserfallen, 2019). For one thing, it does not study the diffusion of a policy but rather of a political strategy: the use of recalls. This is important in terms of how to model the analysis empirically, as we will see below. Consequently, it is not interested in the behavior or decision-making process of policymakers or bureaucrats but rather of political elites. This is also of consequence as it affects how success is defined.

⁴It is entirely possible for politicians to engage in motivated reasoning and convince themselves that their success is also what is best for society, but this does not change the fact that they care mainly about furthering their self-interest.

⁵Pollert and Mooney (2022) accurately note that the problem with using success to identify learning is that what counts as success is criterion-related and the criterion is often left implicit in the literature.

2 The Recall

The recall is a direct democracy mechanism that allows voters to remove an elected official from office before the end of her term via a special election. The recall stands apart from all other direct democracy mechanisms (referendums, plebiscites, popular initiatives) in that it involves a decision regarding a person instead of an issue. It also has two important differences with respect to impeachment procedures: decisions are made by the people instead of parliament, and it does not need to be based on legal grounds (although in some countries a justification is required) (Beramendi et al., 2008).

Since the third wave of democratisation, the recall has been adopted by multiple new democracies throughout the world. In fact, the majority of its adopters are new or unconsolidated democracies, presumably in an attempt to legitimate new democratic regimes by increasing citizen oversight of elected officials and promoting citizen participation (ONPE, 2010).⁶

Regarding its effects, specialists have highlighted a number of potential benefits of the recall, such as increased accountability and control, better communication between voters and politicians and a reduction in political alienation and apathy (Cronin, 1989; Welp, 2014; Twomey, 2011). On the other hand, some drawbacks have also been noted: it distorts legislators' decision-making process by turning them into delegates, it is disruptive and destabilizing, it may accentuate political conflict and polarization, and deepen voters' cynicism and distrust of politicians (Cronin, 1989; Welp, 2014; Twomey, 2011; Siegel, 2015; Welp and Whitehead, 2020a).

In practice, assessing the extent to which the hypothesized effects of the recall have materialized has proven difficult. One aspect that hampers comparative analyses is the fact that procedural requirements vary between countries (and often also states), and play an important role in determining how and when the recall is actually used.⁷ These high levels of variation in institutional design have meant that research on the actual use of the recall is scarce and has remained mostly limited to descriptive case studies (Qvortrup, 2011).

A considerable portion of the existing research has focused on the experience in the U.S., leading to the erroneous impression that “the recall is predominantly an American phenomenon” (Qvortrup, 2011, p. 163). Findings in this literature are quite general, focusing on what determines how frequently recalls are used and what its effects are on

⁶Direct recalls —those that are both initiated and decided by the people— are allowed at the subnational or national levels in countries such as: Bolivia, Ecuador, Venezuela, Peru, Colombia, Argentina, Mexico, Cuba, Poland, Slovakia, Belarus, Kyrgyzstan, Ethiopia, Gambia, Nigeria, Kenya, Kiribati, South Korea, Taiwan and Palau (Beramendi et al., 2008).

⁷The most important variations concern the offices subjected to the recall, the time period during which a recall can be called, the causes for calling it, the magnitude of citizen support required to initiate the process, the quorum required for the recall election to be valid, the majority required for the recall to pass, and the procedure for replacing a recalled authority.

turnout, incumbent power, legislative representation and the influence of special interest groups (Bowler and Cain, 2004; Weinstein, 2005; Qvortrup, 2011; Siegel, 2015; Garrett, 2004; Amar, 2004).

Recently, the proliferation of recall elections outside the U.S. has led to the development of novel efforts to understand both the conditions under which it is used and its effects, in a broader set of countries. This work recognizes that recalls can be associated with negative outcomes and highlights the importance of comparative research in identifying the conditions under which positive effects will materialize. Interestingly, while it considers the possibility that the introduction of the recall may be driven by diffusion, its growing use is attributed to purely internal causes: mainly, growing dissatisfaction with elected representatives (Welp and Whitehead, 2020a).

The experience with the recall in Latin America, a region where it is quite widespread, often used and—in some cases—highly visible, has also been examined. Indeed, this region is considered “the leading laboratory for experimentation with MDDs [direct democracy mechanisms] in general and with recall procedures in particular” (Welp and Whitehead, 2020b, p. 176). However, existing research has mostly consisted of descriptive country studies, detailing the context of introduction of the recall, its procedural features and the number of times it has been activated, and offering plausible explanations for its level of use. Recent comparative approaches have attempted to explain the relative frequency of its activation on the basis of internal variables such as citizen dissatisfaction, institutional design, political party dynamics and the behavior of Electoral Management Bodies (Welp and Castellanos, 2020). The possibility that recalls may diffuse spatially or that there could be learning effects at play has not been considered. However, given its increasing levels of use, and growing awareness that it may not always contribute to strengthening democratic systems (ONPE, 2010; Welp and Whitehead, 2020a), understand how use of the recall spreads is also important. To that end, I turn to the Peruvian case.

3 Recall Elections in Peru

Recalls in Peru were originally regulated in the 1994 Law of Participation and Citizen Control Rights (Law 26300), which allowed for them to be initiated against all members of the executive or legislative branches of subnational governments (i.e., at the district, province and regional levels).⁸ In light of the importance of institutional design when it comes to understanding how recalls are used, in this section I will describe them at some length.

During my period of study, 1996-2014, recalls could take place after an official’s first

⁸Districts are the smallest level of subnational government in Peru. There are currently 1874 districts embedded in 196 provinces, which are embedded in 24 regions.

year in office and before their last, out of a 4 year term. In order to trigger a recall any citizen residing in the targeted official's district could buy a "recall kit"⁹ from the electoral management body (ONPE) by filling in a form stating the name of the official they wanted to recall and the reason, which did not need to be proven. This citizen was then known as the promoter of the recall and was responsible for gathering the signatures of at least 25% of registered voters in the official's district (up to a maximum of 400,000 signatures) in support of the recall. If sufficient signatures were gathered and verified, a recall election took place on a date predetermined by the electoral management body. In this election, citizens had to cast one vote (Yes/No) for every official on the ballot (mayor and/or council members). In order for the recall election to be valid, at least 50% of registered voters in the district had to participate and voting was mandatory. For the recall to pass, it had to be supported by an absolute majority (i.e. 50% + 1) of the votes. Finally, if the recall passed and less than a third of the council members were recalled, the mayor (if recalled) was replaced by the deputy mayor (the council member who was second in the party list) and any recalled council members were replaced by their substitutes. If more than a third of the council was recalled, new elections were called to replace all recalled officials.

Since their regulation in 1994 every subnational government term has witnessed a recall election. Not only that, but the share of districts with recall elections increased continuously during the period of study, from under 4% in the 1996-1998 term to almost 24% in the 2011-2014 term (see figure 1). In total, 5902 officials were subjected to recall votes during this period, with their numbers growing every term.¹⁰ In all, 98.5% of the officials subjected to a recall have been members of district-level local governments, with 747 different district governments (45% of the total) experiencing recall elections but only 9 provinces, and no regional governments.


These high levels of use call for an explanation. Part of it can undoubtedly be found in the features of electoral institutions and the dynamics of political competition, which in combination with the recall's procedural design make it both attractive and undemanding. At the institutional level, closed list plurality elections (in which the party with the most votes wins both the mayor's office and at least 50% +1 of the seats in the municipal council), meant mayors could win with a relatively small percentage of votes and yet have an absolute majority in the council. At the political level, the recall took place in a context of extreme party fluidity, as expressed in high levels of volatility, fragmentation and personalism (Tanaka, 1998; Roberts and Wibbels, 1999). In practice, parties were (and continue to be) little more than electoral platforms politicians use to run for office, with low levels of ideological or organizational cohesion (ONPE, 2010), and short life spans.¹¹

⁹Consisting mainly of the forms needed to gather signatures, with a cost of approximately \$35.

¹⁰See figure A.1 in the SI for the number of officials subjected to recalls each term.

¹¹This situation was made worse by the 2003 Law of Political Parties, which lowered the requirements to register new parties and authorized the creation of subnational political organizations with even lower

Figure 1: Number of districts with recalls by term, Peru 1996-2014



Data source: ONPE.

Finally, the process of decentralization initiated in 2002 increased the competences and resources of district governments, making them more attractive and contributing to the dispersion of votes.¹² As a result, in many districts mayors were elected by a small margin and with fewer votes than the number of signatures required to initiate a recall (i.e., less than 25% of able voters) (Remy, 2005).

These features have all been mentioned in the numerous qualitative studies seeking to explain why, where and to what effect the recall is used in Peru (Tuesta Soldevilla, 2014a,b; Welp, 2013, 2015; Quintanilla, 2013). Reasons advanced for why the recall is so intensively used typically include: high levels of citizen dissatisfaction, weak parties and an unconsolidated party system, and low procedural thresholds. Regarding where recalls take place, answers include: in small districts, in districts with large budgets, in conflictive districts, and in districts where mayors win with a small percentage of votes or

requisites. As a result, political competition at the district level can take place between national parties, regional movements and local political organisations. National parties are those that can compete in all elections throughout the country and are the most stable over time, while regional movements can compete only in elections within the region in which they were created. Local political organisations can exist at the provincial or district levels, the former can compete in all elections within the province of their creation and the latter can only compete in the districts in which they were created. These types of organizations are highly volatile and new ones tend to appear and disappear for each election cycle.

¹²In fact, between 1995 and 2009, transfers from the central government to local —provincial and district— governments increased tenfold, going from a little over one billion New Soles, to over 10 billion New Soles (ONPE, 2010).

by a small margin. In terms of their effects, there is almost universal consensus that these are negative.¹³ Most studies highlight that recalls have a negative effect on competition, distract officials from their tasks and contribute to citizens’ distrust in politicians.

In a 2015 article, Welp argues that in a context of low party system institutionalization “in which neither political leaders nor their parties have any guarantee of survival beyond the next election, recall procedures create unexpected incentives for politicians to use them as an extension to “normal” electoral competition” (Welp, 2015, p. 11). She thus claims that the incentive to bring about a new election through the use of recall procedures, combined with unstable patterns of political competition, facilitates the formation of an alliance of electoral losers in the hope of triggering a do-over election. Unfortunately, she does not attempt to test this hypothesis. However, this article raises an important point: the fact that local political organizations lose their registration (and thus their capacity to run for election) after every election means that local politicians have very high discount rates: they may be unwilling to wait patiently until the next election because there is no guarantee that their political organization will still exist or that they themselves will be candidates.

Two more recent articles use a quantitative approach to shed more light on the matter. In the first, Holland and Incio study the determinants of recalls between 2002 and 2015 and argue that it is losing politicians who are organizing recall referenda (Holland and Incio, 2019). They show that recall organizers were candidates in the previous election in 18% of cases and that recalls are more likely to occur when politicians win by narrow margins or are female. In the second, Artiles et al. (2021) argue that the political use of recalls has a negative effect on candidate selection. They show that in districts where mayors were recalled by a small margin candidates in the next round of elections were of lower quality—in terms of education and prior experience—than in districts where recalls failed by a small margin.

Overall, there is broad agreement in this literature on a number of stylized facts regarding recall elections in Peru. First, that they are not generally initiated by citizens to remove corrupt or unresponsive politicians but rather by losing candidates in an attempt to re-do municipal elections.¹⁴ Second, this is seen as a response to the incentives created by a context of high fragmentation, in which winners win with low vote shares and by low margins. The former makes it easy to gather the signatures required to initiate the recall process among the large majority that did not vote for the winning candidate, while the latter is expected to affect losers’ calculus regarding their possibilities in a

¹³Perhaps the only work that defends the use of recalls, while still acknowledging the negative effects of their use in Peru, is Quintanilla (2013).

¹⁴The fact that an overwhelming majority of recall processes target—in addition to the mayor—at least 1/3 of council members (despite their low visibility) is further evidence that their goal is to trigger new elections. See figure A.1 for the share of districts with recalls in which this condition is observed in each government term.

new election. Moreover, high levels of party fluidity means politicians have short time horizons: in most cases they will have to re-register their political organization before the next election, making the prospect of waiting unattractive. All of this results in high levels of instability, under-performing municipal governments, high levels of polarization and ultimately high voter dissatisfaction and low levels of politician quality.¹⁵

It is noteworthy that despite having attracted so much attention, the possibility that recalls may diffuse as candidates in one district imitate or learn from what happens in neighboring districts has not been seriously considered. Rather, the working assumption in this body of work seems to be that politicians in each district will independently decide to organize recall elections in response to similar conditions (close races, high fragmentation, small districts). Given the overwhelming evidence that recalls are prejudicial to democratic functioning in Peruvian districts, this paper seeks to expand upon prior contributions by asking: is there diffusion in the spread of recall elections in Peru? If so, was learning the main mechanism, and what was being learned?

Figure 2, which shows the term in which each district experienced its first recall election, provides a first indication that spatial diffusion may be at play. Indeed, the existence of spatial autocorrelation in the use of recalls is confirmed by a positive and significant Moran's I (except for the first round of recall elections in the 1996-1998 term).¹⁶

In what follows I will try to explain the mechanism driving this spatial pattern of diffusion. To do so I focus on the period between 1996 and 2014, which is when the bulk of recall elections took place. In 2015, after a very contentious and disruptive attempt to recall the mayor of the capital city of Lima, recall procedures were reformed with the intention of making the recall less attractive. The reform eliminated the possibility of new elections by stipulating that recalled mayors would be replaced by the first substitute council member in their party list, and postponed any recalls until the third year out of a four-year term.¹⁷ In the years that have elapsed since this reform less than 200 officials have been subjected to recall votes, compared to 5,902 in the prior 15 years.¹⁸

¹⁵The disruptive effect of recalls on municipal performance can easily be inferred from the fact that districts in which new elections are triggered by a recall end up having three different municipal governments in the span of four years: the original municipal government, a provisional government after the recall, and the newly elected government after new elections take place.

¹⁶Moran's I is a measure of the spatial autocorrelation in a dataset. It measures the correlation in both location and values simultaneously. Intuitively, it calculates deviations from the mean for each observation and multiplies deviation values for neighboring features to obtain cross-products that will be positive if neighboring values are both smaller/larger than the mean. Its numerator is made up of the sum of these cross products and is normalized by the variance so that index values fall between -1 (perfect dispersion) and +1 (perfect clustering). A significant value means we can reject the null of complete spatial randomness, and a positive value indicates similar values—in this case recalls—are clustered spatially.

¹⁷The reform also eliminated the cap of 400,000 signatures to initiate a recall (which was only binding for Lima) and required recall promoters to disclose their campaign contributions and expenses.

¹⁸Indeed, in the 2015-2018 term only 27 districts experienced recalls and in the 2019-2022 term this number went down to 13 districts.




Figure 2: Districts by term with first recall election

4 Theory and Hypotheses

The empirical strategy described below starts from the assumption that prior research is right in placing the responsibility for the activation of recalls in the hands of losing politicians (as opposed to disgruntled citizens). In section 7 I will present further evidence in support of this assumption. Note however that while who is responsible for recalls is of no consequence when it comes to identifying diffusion patterns, it does matter for our interpretation of the mechanisms and their normative implications, as will become obvious below.

The diffusion literature has traditionally studied the spatial dissemination of policies by looking at the behavior of neighbors. Geographic proximity is expected to promote diffusion through multiple potential avenues: by facilitating the development of communication networks through which information about policies can spread; via overlapping media markets alerting to the existence of policies in nearby states; and because officials may be most inclined to use nearby states as models based on their cultural or demographic similarity (Karch, 2007).¹⁹ However, as access to communication and transportation has expanded, spatial proximity has lost some of its relevance (Karch, 2007; Shipan and Volden, 2012). Here I nonetheless retain the focus on neighbors for two key reasons. The first is that even if information about what happens all over the country—or even in other countries—is available, politicians are still likely to find what happens in districts that are similar to theirs more informative. Indeed, this is consistent with sociological research showing people draw lessons from members of their networks (Dobbin et al., 2007). The second is that in the Peruvian districts that I study access to the internet—and to information more generally—, cannot be taken for granted, particularly during the period of observation.

As is always the case when studying diffusion, the null hypothesis is that the patterns we observe are a result of spurious diffusion. This would mean residents of neighboring districts independently decided to conduct recalls in reaction to the same set of incentives. As mentioned above, this seems to be the assumption underlying most of the existing literature. If it were so, we would expect to observe recalls propagating spatially (again, assuming proximity proxies for similarity) as the existence of a recall in one district is correlated with the existence of a recall in a nearby district. However, neither the number of neighbors experiencing recalls, nor the results of recalls in nearby districts should matter, given that districts are activating recalls independently of their neighbors. For clarity, it is worth spelling out this null hypothesis.

Hypothesis 0. *If recalls propagate independently, the existence of recalls in neighboring districts will increase the probability of a recall taking place in the following term, but*

¹⁹One notable exception is the study of competition in international trade, where structural equivalence is much more relevant than geographic location.

neither the number of neighbors with recalls nor the result of their recalls should have an effect.

If, on the other hand, there is diffusion, the question then is: what is the mechanism driving diffusion? In the context of subnational recall elections neither coercion nor competition make sense. Two potential mechanisms are thus worth considering: simple imitation or learning. The key distinguishing aspect is that imitation “is not related to the objective consequences of a policy” (Gilardi, 2016, p. 10), while learning is. Imitation implies copying what others have done because of who has done it, in an effort to be like that other or to conform to one’s normative environment. In our case, it would mean that politicians in district A observe a recall happening in district B and naively decide to imitate it. Learning on the other hand “means being influenced by successful policies” (Maggetti and Gilardi, 2016). The standard test of learning is thus that only successful policies are adopted (Dobbin et al., 2007; Shipan and Volden, 2008; Maggetti and Gilardi, 2016; Gilardi, 2016). However, as has been noted elsewhere, learning also implies a negative effect of unsuccessful ones (Volden et al., 2008; Lee and Strang, 2006).

Hypothesis 1. *If learning is the mechanism driving the diffusion of recall elections, the probability that a district has a recall will increase with the share of successful recalls among its neighbors in the prior term, and decrease with the share of unsuccessful ones.*

The definition of success when it comes to recall elections may not be straightforward. Is it simply replacing the current mayor? If so, does it matter who replaces him? Or is it triggering a new election? According to existing research, the goal of the politicians promoting recalls is to prompt new elections with the hope of performing better than in the last one. Success therefore means triggering new elections by removing the mayor and at least a third of the council.

I focus on the effect of the share of successful/failed recalls rather than the mere occurrence of a successful/failed recall among a districts’ neighbors following the common assumption that the opportunity to learn is increasing in the number of instances of an action (Shipan and Volden, 2008; Boehmke and Witmer, 2004). Moreover, the more consistent the evidence, the more likely politicians are to update their beliefs (Dobbin et al., 2007).

As noted by Braun and Gilardi (2006), expected utility theory can help clarify the decision processes behind diffusion. They argue that decision-makers’ utility is determined by both the expected payoffs of the new policy and its effectiveness in delivering the expected outcomes ($EU = \text{Payoffs} \times \text{Effectiveness}$). In our case, we have office-seeking politicians who expect a certain payoff from participating in new elections (a function of their perceived probability of success) and have a given prior regarding the effectiveness of recalls as an instrument for triggering new elections. The observation of recall processes in nearby districts provides them with new information that allows them to update

their beliefs (i.e., learn) regarding recall effectiveness. Observing successful recalls would therefore increase the perceived effectiveness of recalls and their expected utility from promoting them, while observing failed recalls would have the opposite effect. Note that the relative size of payoffs from participating in a new election is not affected by the experiences of neighbors. However, we might imagine that the smaller the margin of victory in the original election (which they lost), the higher politicians' perceived probability of success in a new election. And depending on their "prior beliefs, policymakers may be more or less sensitive to the cues coming from other countries" (Gilardi, 2010, p. 651). In essence, the fact that their prior regarding their expected utility from a new election is higher means they are more likely to update their beliefs after observing a successful recall. This allows us to postulate a new expectation:

Hypothesis 2. *If politicians are driven by the desire to win a future election, the positive effect of observing successful recalls in neighboring districts will be larger in districts where the incumbent's margin of victory was smaller.*

These hypotheses reflect the main strategies suggested for the identification of learning in the policy diffusion literature: a positive effect of successful policies, a negative effect of ineffectiveness, and the identification of relevant conditional effects (including on the basis of adopters' prior beliefs or relevant characteristics of the adopting and/or observed unit) (Volden et al., 2008; Gilardi, 2010; Shipan and Volden, 2008).

It is worth noting that none of this requires the assumption that politicians are perfectly rational Bayesians. Indeed, these expectations are consistent with the more likely possibility of bounded rationality.²⁰ As noted by Shipan and Volden (2012), effective learning is time-intensive and takes a high-degree of skill, making it unlikely among time-pressed policymakers with limited staff, not to mention losing candidates. These politicians are thus expected to update their beliefs on the basis of new information but at the same time engage in well-known cognitive shortcuts. This is part of the reason why they would update based on what happens in neighboring districts rather than use full information about what happens in all districts (availability heuristic) and expect neighbors' success to matter for theirs (representativeness heuristic) (Dobbin et al., 2007).

5 Empirical Strategy and Data

I use data on district-level recall elections in Peru between 1996 (the starting year of the first term with recalls) and 2014 to test my hypotheses. Using data collected from a number of different official sources I have constructed a balanced panel data set in which the unit of analysis is the district-subnational government term. I therefore have 1632

²⁰Bounded rationality is in fact not only more likely but has been documented in the context of policy diffusion in Latin America and elsewhere (Weyland, 2005; Dobbin et al., 2007).

districts²¹ observed five times: once for each of the 1996-1998, 1999-2002, 2003-2006, 2007-2010 and 2011-2014 terms.

Data on the officials subjected to a recall vote and the recall's results, by district and term, come from the electoral management body (ONPE), and was used to construct my dependent variable: a dummy for the occurrence of a recall election in a district-term. This data is also used to construct my main independent variables: a dummy for the occurrence of a recall in neighboring districts in the previous period, and the proportion of neighbors with (successful or failed) recall elections in the previous period. GIS data from the Ministry of the Environment was used to identify neighboring districts, defined as contiguous districts, and construct a contiguity matrix.²² Data on the parties competing in each election, the distribution of votes and the identity of the winner in each district also come from ONPE and were used to calculate the margin of victory, and the effective number of electoral parties (ENEP).²³ I also collected data on mayor's prior experience by scraping the National Jury of Elections' webpage, which offers information on all elected officials, including the history of elective offices they have held. With this information I constructed a variable recording the number of times the current mayor previously held different types of elected office at the provincial or district levels.²⁴

The fact that I don't study policy adoption but rather the use of recalls has implications for my empirical strategy. The policy diffusion literature has typically used event history analysis, an approach based on the assumption that events (policy adoptions) can only occur once, with units dropping out of the database once adoption occurs. In my case, districts can conduct recalls every government term, requiring an altogether different modeling strategy. I therefore use the following linear probability model to test my hypotheses by running fixed effects panel regressions:²⁵

$$Recall_{dt} = \beta PropSuccesful_{-dt-1} + \gamma PropRecalls_{-dt-1} + \lambda_d + \eta_t + \epsilon_d \quad (1)$$

The dependent variable is an indicator for the occurrence of a recall election in district d at time t . The coefficient of interest is β , which captures the effect of the proportion

²¹In 1996 there were 1829 districts in Peru. However, districts that are provincial capitals do not have a district government, as its functions are taken over by the provincial government. Thus, of the 1829 districts, 195 provincial capitals and 2 islands are excluded, leaving 1632 observations. Since 2002 around 40 new districts have been created, but these are not included to maintain a balanced data set. Moreover, there have not been any recalls in these districts.

²²For robustness, the (straight line) distance to the centroid of the nearest district with a recall in the previous period was also used as a measure of proximity.

²³In districts in which elections were annulled, data from the supplementary elections was used. The annulment of elections usually results either from social unrest that prevented votes from being counted, or from the fact that more than two thirds of the ballots were either blank or null.

²⁴Regional experience was not considered because only one mayor had ever been elected to a regional-level position.

²⁵In light of the fact that OLS estimates are expected to be reliable as long as the model specification is correct, the interpretability of OLS is preferred to logit or probit specifications, despite the risk of non-sensical predictions.

of neighboring districts (not d) with successful/failed recalls in the previous term ($t - 1$).²⁶ The model also includes a control for the overall proportion of neighbors with recalls in the prior term ($PropRecalls_{-dt-1}$). This allows us to isolate the effect of recall outcomes (success/failure), as captured by β , from that of their frequency, which will be captured by γ .²⁷ λ_d are district fixed effects, which control for time-invariant district-level characteristics (e.g., slow-moving socio-economic characteristics) that may affect the probability of a recall taking place. η_t are time fixed effects, controlling for the influence of period-specific shocks that may affect the probability of a recall in all districts at the same time (e.g., the overall level of recalls in a given term). Cluster robust standard errors (clustered at the district level) are used to account for repeated observations.

This model is expected to be a quite consistent estimator, as it assesses how the probability of a recall changes over time for a given district as the behavior of its neighbors changes. It therefore eliminates the need to account for district-level time-invariant covariates. As for time-varying district-level characteristics, the policy diffusion literature broadly acknowledges the need to include controls for potential internal determinants of the outcome (Shipan and Volden, 2012, 2008; Franzese Jr and Hays, 2008). I therefore run additional models in which I include controls for time-varying district-level characteristics that may affect the probability of a recall taking place. These include two indicator variables for whether the district mayor comes from the same party as the provincial mayor or the regional president, the ENEP in the most recent election²⁸, a measure of mayor quality (as proxied by experience in elected office) and the cumulative number of past recall elections.²⁹ The same party dummies are intended to account for the possibility that higher-level authorities may benefit or hurt mayors depending on their party affiliation, thus affecting the likelihood that a recall will take place. Moreover, sharing the same party as higher-level officials may affect the level of ability needed to be elected (i.e., candidates can compensate for less ability with political support from these authorities and their parties). The number of parties is included as it will determine the dynamics of competition by affecting both politicians' incentives to promote a recall (through its effect on vote shares) and the quality of elected representatives. Mayor quality controls for the possibility that recalls are used to remove incompetent mayors. The number of recall elections that have taken place in district d before t seeks to control for the negative effect of past recalls on the candidate pool. Furthermore, experience with the occurrence

²⁶Note that according to Franzese Jr and Hays (2008) the inclusion of time-lagged spatial lags in the model prevents any simultaneity bias.

²⁷To clarify, the main independent variable counts the number of districts experiencing successful/failed recalls as a share of neighboring districts *with recalls* (as opposed to of all neighboring districts). The variable counting the proportion of neighbors with recalls thus captures the effect of the number of neighbors with recalls while the independent variable isolates the effect of observing particular recall outcomes.

²⁸I used Laakso and Taagepera's formula for the effective number of electoral parties (1979).

²⁹Recall events are counted here, regardless of the number of district officials on the recall ballot.

of recalls in the past may also affect the probability that recalls will take place by lowering informational costs and/or increasing citizens' opposition to recalls.

Models with controls are run on data for the period between 2002 and 2014. I thus exclude the first two terms with recall elections (1996-1998 and 1999-2002) with the purpose of maintaining institutional features (and their effects) as constant as possible. These two recall elections took place before (i) the Decentralisation Law was passed, which means the structure and competences of subnational governments was different; (ii) the Law of Political Parties was passed, which means the type of political organisations in competition (and the requirements to create them) were different and; (iii) the Municipal Elections Law was reformed, which means the electoral rules governing municipal elections were different (namely, there was a run-off election if no party got at least 20% of the vote). These institutional differences are expected to have a non-negligible effect on the dynamics of local political competition (and therefore on politicians' incentives), raising issues about their comparability with later periods. Indeed, as a result of these reforms, while in the 1995 and 1998 elections the number of national political parties competing at the district level was 2 and 5, respectively, in 2002 this number jumped up to 14.

To test hypothesis 2 I use the same baseline model but add an interaction between the variable measuring the proportion of successful neighbor recalls in $t - 1$ and a dummy identifying districts d in which the incumbent mayor was elected with a margin of victory smaller than 5 percentage points.³⁰ Again, these results are restricted to the 2002-2014 period.

6 Results

Table 1 presents the results of the tests of hypothesis 1. Model 1 presents initial evidence against the null hypothesis of independent use of the recall by showing that the mere occurrence of a recall in a neighboring district has no effect on the probability of a recall. This suggests that similar districts are not independently conducting recalls in reaction to similar conditions. Moreover, it is worth noting that the assumption of independent activation implicit in the existing literature would require the conditions determining the use of the recall —close elections— to expand in tandem with recalls during my period of study. However, as shown in figure A.2, this has not been the case. Indeed, while the number of districts with recalls has monotonically increased each term, the number of districts with close elections remained stable (or decreased) starting in 2003. All of this suggests diffusion is at play. Model 2 presents further supporting evidence by showing that the probability of a recall increases with the proportion of neighbors who had recalls in the prior period. This type of variable has often been used in the study of diffusion.

³⁰In the robustness section I show that results are robust to alternative measures of close elections.

However, as noted by Maggetti and Gilardi (2016) it is much too general to identify any particular mechanism.

Models 3 and 4 zero in on the possibility of learning as the mechanism driving diffusion by focusing on the effect of successful and unsuccessful neighbor recalls. As predicted by hypothesis 1, model 3 shows that the probability of experiencing a recall increases with the share of neighbors experiencing successful recalls. As noted above, a recall is considered successful if it leads to the removal of a mayor and triggers a new election, and unsuccessful if the mayor is not removed.³¹ Similarly, failed neighbor recalls reduce the probability that a recall will take place. The significance and direction of these results does not change when adding controls, as shown in table A.4. Coefficients do change (they become bigger), but at least some of this is due to differences in the sample as a result of excluding the 1999-2002 term (see table A.3 for results excluding controls for the 2002-2014 period). Based on the full model with controls, even after accounting for the effect of the overall number of neighbors experiencing recalls, a unit increase in the proportion of neighbors with a successful recall is expected to increase the probability of a recall by 7 percentage points. Failed recalls are expected to reduce the probability of a recall by a similar amount. Given that the average probability of a recall in the sample is 14%, the diffusion effects found are substantial.

³¹In theory, recalls will only trigger a new election if over one third of the municipal council is also recalled. In practice, the vast majority of recalls targeted both the mayor and the required number of councilors, as shown in figure A.1.

Table 1: Is there learning?

	Probability of a Recall			
	(1)	(2)	(3)	(4)
At least one neighbor had recall	0.019 [0.012]			
Proportion of neighbors with recall		0.078* [0.043]	0.043 [0.046]	0.139** [0.056]
Proportion of neighbors with successful recall			0.044** [0.020]	
Proportion of neighbors with failed recalls				-0.033* [0.017]
Observations	6,528	6,528	6,528	6,528
Districts	1632	1632	1632	1632
R-squared	0.374	0.375	0.375	0.375
Term Fixed Effects	YES	YES	YES	YES

NOTES. Results from OLS panel regressions with district fixed effects. Cluster robust standard errors (clustered by district) are reported in brackets. The dependent variable is a binary indicator coded as 1 for districts in which a recall election took place and 0 for those in which it did not.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

According to Seljan and Weller (2011) the driving force in diffusion processes is the flow of information. In that sense, it is worth considering exactly what is being learned. Two possibilities have been raised regarding the content of the information that is communicated: (i) the first one focuses on policy information regarding the existence or availability of a given policy; (ii) the second one focuses on political information regarding the political viability, attractiveness and potential (ensuing) electoral benefits associated with a given policy (Seljan and Weller, 2011). One way of interpreting the results in table 1 is thus that the proportion of neighbors with recalls variable (regardless of outcomes) captures the effect of information transmission regarding the availability of the recall as a participatory mechanism. On the other hand, the positive and negative effects of successful and failed recalls (respectively) capture the transmission of political information regarding the value of the recall as an instrument for triggering new elections.

Table 2: Conditional effects

	Probability of a Recall	
	(1)	(2)
Proportion of neighbors with successful recall	0.040 [0.030]	0.030 [0.026]
Close election	0.055*** [0.013]	0.026** [0.012]
Proportion successful $\times$ close election	0.090* [0.047]	0.104** [0.042]
Observations	4,896	4,896
Districts	1632	1632
R-squared	0.459	0.573
Term Fixed Effects	YES	YES
Controls	NO	YES

NOTES. Results from OLS panel regressions with district fixed effects. Cluster robust standard errors (clustered by district) are reported in brackets. The dependent variable is a binary indicator coded as 1 for districts in which a recall election took place and 0 for those in which it did not. Close election is a dummy identifying terms in which the incumbent mayor was elected with a margin of victory under 5 percentage points. Model 2 includes controls for proportion of neighbors with recalls, ENEP, same party as provincial mayor, same party as regional president, mayor experience and prior recalls.


*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Moving on to hypothesis 2, table 2 presents the results of models including an interaction between the share of successful neighbor recalls in the prior term and a dummy identifying district-terms in which the incumbent mayor was elected in a close election.³² In both models (with and without controls), the positive interaction term indicates that the effect of successful neighbor recalls is larger in districts with mayors elected in close elections. This is consistent with our expectation that when the incumbent was elected by a small margin of victory losing politicians in the district have a higher perceived probability of success in a re-do election. These politicians are thus more likely to positively update their beliefs regarding the effectiveness of recalls upon observing the success of their peers in neighboring districts. The positive and significant coefficient on the main effect of the close election variable indicates that recalls are more likely in the presence of close elections, even in the absence of successful neighbor recalls. However, as shown

³²Here we consider elections as close when the margin of victory is under 5%. However, as shown in section 7, results are robust to different operationalizations.

in figure 3, successful neighbor recalls have an insignificant effect in the absence of close elections.

Figure 3: Marginal effect of the proportion of successful neighbor recalls on the probability of a recall, 2002-2014



Note: Figure graphs marginal effects from model 2 in table 2 (with controls).

These results are consistent with losing politicians driving the diffusion of the recall by learning about their political viability from their neighbors and using them when conditions are favourable. They show that office-seeking politicians learn to behave strategically and abuse participatory institutions to increase their chances of gaining office, to the detriment of governmental efficiency, citizen trust in politicians and institutions, and candidate quality. It is worth noting that while prior research had shown that political parties (in a very different setting: Switzerland) use direct democracy mechanisms strategically to improve their electoral chances (Leemann, 2015), the possibility of diffusion was not previously considered.

7 Robustness Tests

This section discusses the robustness of my findings to alternative definitions of the dependent and independent variables, and further probes the validity of my main assumption: that it is losing politicians who are driving the diffusion of recalls.

Above, I have presented evidence consistent with the spatial diffusion of recall elections through a learning mechanism. Based on prior research, I built on the assumption that the actors driving the diffusion of recall elections were losing politicians seeking to re-do elections. The evidence presented in support of hypothesis 2 is consistent with this assumption. However, it may still be the case that citizens are also learning about the recall and its political value and are therefore also responsible for initiating recalls in an attempt to remove unsatisfactory politicians. To probe this possibility, I examined whether recalls are associated with observable measures of mayor quality by testing whether incapable, unpopular or inexperienced mayors are more likely to be recalled. This analysis is included in SI section A.3. Results are actually the opposite of what we would find if citizens were recalling authorities that are unpopular or unable to manage social conflicts in their jurisdiction. Likewise, at least in the 2011-2014 term (the only one for which budget execution data is available), recalls were not associated with mayors' incapacity to execute their districts' budget. All of these results lend credence to the assumption that recalls are not used by citizens to get rid of unsatisfactory officials.

In terms of the dependent variable, I re-estimated all models using the dummy indicating the occurrence of a recall vote targeted specifically at the mayor instead of any recall vote. Results were unchanged.

For the spatial models, table A.5 in the SI presents the results of a series of regressions using distance to the nearest district with a recall in the previous period as a proxy for proximity (or similarity), instead of contiguity. Maximum distance in this table refers to the maximum distance to the nearest district with a recall, which can be the full maximum distance (645 kms.) or a threshold set at the 95th and 90th percentiles (107 and 77 kms., respectively). I find that distance only has the expected negative effect in districts within the 90th percentile of distance, which means recalls in other districts will only have an effect if they are close enough.

In terms of the tests of H2, table A.6 in the SI presents the results of a series of replications of model 2 in table 2 using alternative specifications for close elections. Model 1 uses a dummy identifying district-terms in which the margin of victory was under 8%, model 2 uses 10% as the critical margin of victory for close elections and model 3 uses a continuous measure of the margin of victory, truncated at 30% or over. The main results remain unchanged: the effect of successful neighbor recalls is larger when the incumbent was elected in a close election. Figure A.3 shows that the marginal effect of the proportion of neighbor recalls is positive when the margin of victory is small and decreases in magnitude as this margin grows, ultimately losing significance around a margin of 15%.

8 Discussion

The evidence presented here indicates that in Peru recalls diffuse through a process of learning, in which politicians learn about the use of the recall from their neighbors (either literally or in the sense of proximate). In particular, they learn not only about the availability of recalls as a potential tool in their political arsenal, but more importantly, about the recall’s political viability and payoffs. In contrast, there is no indication that recalls are activated in response to citizen demands. Moreover, qualitative research has consistently argued that recalls in Peru are disruptive, destabilizing and contribute to citizens’ distrust of politicians (Welp, 2013, 2015; Tuesta Soldevilla, 2014b). Empirical research has further shown that they lower the quality of the pool of political candidates in subsequent elections (Artiles et al., 2021). The diffusion of recalls in Peru can therefore be described as a case of perverse learning, in which politicians learn to abuse participatory institutions to further their political interests at the expense of the common good.

It may be argued that the Peruvian case is an exceptional one —as expressed in its particularly high levels of use—, and we therefore cannot draw any general (or generalizable) lessons from it. However, there are good reasons to believe that at least some of the dynamics uncovered here are also present in other, very different, contexts. Indeed, at the subnational level in the U.S. there are multiple reports that (i) the number of politically motivated recalls is expanding (Siegel, 2015; Spivak, 2020); (ii) that they are used to re-do elections in cases in which voters have been unable to elect candidates who mirror their preferences, or as a result of politicians’ desire to capture an office, or with the goal of flipping a closely divided state legislature (Weinstein, 2005; Spivak, 2020)³³; (iii) that they have a polarizing effect on local politics and may prove harmful to the long term stability of government, partly because they allow special interest groups to exercise disproportionate control over state politics relative to constituents (Siegel, 2015; Amar, 2004); (iv) and that a potential reason for their dramatic increase is that there is more awareness of their availability, with the success of recent recalls making them more appealing (Siegel, 2015; Spivak, 2020). Similar dynamics have been observed elsewhere. In Colombia, for example, it has also been found that a majority of recalls are triggered by losing politicians, ultimately eroding democratic governance (Welp and Milanese, 2018). In Central and Eastern Europe, recalls are similarly described as instruments of partisan account setting (Misoiu, 2020). More generally, a recent book on the recall summarizes the comparative evidence by noting that it is often used by political actors to direct their fire against rival parties (Welp and Whitehead, 2020a). The extent to which there is learning involved in these additional cases is an open question for future research.

One possible take-away from the evidence presented here could be to reiterate the

³³As noted by Bowler (2004) “the recall skews incentives of poor losers by providing them with a means of revenge” (Bowler, 2004, p. 206).

well-known fact that institutional design plays a key role in determining how institutions are used. However, its importance in determining how —and how often— recalls in particular, and direct democracy mechanisms in general, are used, has already been highlighted elsewhere.³⁴ What I wanted to show here is the process by which actors can *learn* to abuse these institutions, when institutional design allows.

Subnational governments are often considered to function as laboratories in which policy innovations can be tested and, if successful, replicated through a learning process. This paper challenges the assumption that learning will necessarily lead to efficiency gains and shows that it can lend itself to perverse dynamics. In the case presented here, politicians learn to (ab)use participatory institutions in a way that furthers their interests —or at least is perceived by them to do so— at the cost of government stability, effectiveness and very likely, citizen trust. Future research should continue to investigate the conditions under which different kinds of actors are likely to engage in perverse forms of learning.

³⁴For example, Gerber notes for the initiative that specific design features will determine “how the incentives that elected officials face are structured and whether or not these incentives influence behavior as intended” (Gerber, 1996, p. 125). See also Bowler and Cain (2004) for the recall.

References

- Vikram David Amar. Adventures in direct democracy: The top ten constitutional lessons from the california recall experience. *California Law Review*, 92(3):927–958, 2004.
- Miriam Artiles, Lukas Kleine-Rueschkamp, and Gianmarco León-Ciliotta. Accountability, political capture, and selection into politics: Evidence from peruvian municipalities. *Review of Economics and Statistics*, 103(2):397–411, 2021.
- Justus Bamert, Fabrizio Gilardi, and Fabio Wasserfallen. Learning and the diffusion of regime contention in the arab spring. *Research & Politics*, 2(3):2053168015593306, 2015.
- Virginia Beramendi, Andrew Ellis, Bruno Kaufman, Miriam Kornblith, Larry LeDuc, Paddy McGuire, Theo Schiller, and Palle Svensson. *Direct democracy*. IDEA, International Institute for Democracy and Electoral Assistance, 2008.
- Frederick J Boehmke and Richard Witmer. Disentangling diffusion: The effects of social learning and economic competition on state policy innovation and expansion. *Political Research Quarterly*, 57(1):39–51, 2004.
- Tobias Böhmelt, Lawrence Ezrow, Roni Lehrer, and Hugh Ward. Party policy diffusion. *American Political Science Review*, 110(2):397–410, 2016.
- Shaun Bowler. Recall and representation arnold schwarzenegger meets edmund burke. *Representation*, 40(3):200–212, 2004.
- Shaun Bowler and Bruce Cain. Introduction—recalling the recall: Reflections on california’s recent political adventure. *Political Science and Politics*, 37(01):7–10, 2004.
- Dietmar Braun and Fabrizio Gilardi. Taking ‘galton’s problem’ seriously: Towards a theory of policy diffusion. *Journal of theoretical politics*, 18(3):298–322, 2006.
- Thomas E Cronin. *Direct democracy: The politics of initiative, referendum, and recall*. Harvard University Press, 1989.
- Frank Dobbin, Beth Simmons, and Geoffrey Garrett. The global diffusion of public policies: Social construction, coercion, competition, or learning? *Annu. Rev. Sociol.*, 33:449–472, 2007.
- Glenn Ellison and Drew Fudenberg. Rules of thumb for social learning. *Journal of political Economy*, 101(4):612–643, 1993.
- Robert J Franzese Jr and Jude C Hays. Interdependence in comparative politics: Substance, theory, empirics, substance. *Comparative Political Studies*, 41(4-5):742–780, 2008.

- Elizabeth Garrett. Democracy in the wake of the california recall. *University of Pennsylvania Law Review*, 153(1):239–284, 2004.
- Elisabeth R Gerber. Legislative response to the threat of popular initiatives. *American Journal of Political Science*, pages 99–128, 1996.
- Fabrizio Gilardi. Who learns from what in policy diffusion processes? *American journal of political science*, 54(3):650–666, 2010.
- Fabrizio Gilardi. Transnational diffusion: Norms, ideas, and policies. *Handbook of international relations*, 2:453–477, 2012.
- Fabrizio Gilardi. Four ways we can improve policy diffusion research. *State Politics & Policy Quarterly*, 16(1):8–21, 2016.
- Fabrizio Gilardi and Fabio Wasserfallen. The politics of policy diffusion. *European Journal of Political Research*, 58(4):1245–1256, 2019.
- Erin R Graham, Charles R Shipan, and Craig Volden. The diffusion of policy diffusion research in political science. *British journal of political science*, 43(3):673–701, 2013.
- Alisha C Holland and José Incio. Imperfect recall: the politics of subnational office removals. *Comparative Political Studies*, 52(5):777–805, 2019.
- Andrew Karch. Emerging issues and future directions in state policy diffusion research. *State Politics & Policy Quarterly*, 7(1):54–80, 2007.
- Markku Laakso and Rein Taagepera. The “effective” number of parties: A measure with application to west europe. *Comparative political studies*, 12(1):3, 1979.
- Chang Kil Lee and David Strang. The international diffusion of public-sector downsizing: Network emulation and theory-driven learning. *International organization*, 60(4):883–909, 2006.
- Lucas Leemann. Political conflict and direct democracy: Explaining initiative use 1920–2011. *Swiss Political Science Review*, 21(4):596–616, 2015.
- Martino Maggetti and Fabrizio Gilardi. Problems (and solutions) in the measurement of policy diffusion mechanisms. *Journal of Public Policy*, 36(1):87–107, 2016.
- Marco Martini and Stefanie Walter. Learning from precedent: how the british brexit experience shapes nationalist rhetoric outside the uk. *Journal of European public policy*, 31(5):1231–1258, 2024.
- John G Matsusaka. *Let the people rule: How direct democracy can meet the populist challenge*. Princeton University Press, 2020.

- Covadonga Meseguer. Policy learning, policy diffusion, and the making of a new order. *The Annals of the American Academy of Political and Social Science*, 598(1):67–82, 2005.
- Sergiu Miscoiu. Recall practices in central and eastern europe: From citizen accountability to partisan account settling. In Yanina Welp and Laurence Whitehead, editors, *The politics of recall elections*, pages 143–157. Springer, 2020.
- ONPE. *Consultas populares de revocatoria 1997-2009 y nuevas elecciones municipales 2005-2010*. ONPE, 2010.
- Julianna Pacheco. The social contagion model: Exploring the role of public opinion on the diffusion of antismoking legislation across the american states. *The Journal of Politics*, 74(1):187–202, 2012.
- Isaac Pollert and Christopher Z Mooney. Substantive and political learning among the us states: Abortion policy diffusion, 1993–2016. *State Politics & Policy Quarterly*, 22(3):320–343, 2022.
- Alfredo Quintanilla. El debate sobre la revocatoria y las reformas de la ley 26300. *Perú Hoy. Susurros desde Babel*, 2013.
- Matt Qvortrup. Hasta la vista: a comparative institutionalist analysis of the recall. *Representation*, 47(2):161–170, 2011.
- María Isabel Remy. Los gobiernos locales en el Perú: entre el entusiasmo democrático y el deterioro de la representación política. In Victor Vich, editor, *El estado está de vuelta: desigualdad, diversidad y democracia*, pages 111–36. Instituto de Estudios Peruanos, 2005.
- Kenneth M Roberts and Erik Wibbels. Party systems and electoral volatility in latin america: a test of economic, institutional, and structural explanations. *American Political Science Review*, pages 575–590, 1999.
- Ellen C Seljan and Nicholas Weller. Diffusion in direct democracy: The effect of political information on proposals for tax and expenditure limits in the us states. *State Politics & Policy Quarterly*, 11(3):348–368, 2011.
- Charles R Shipan and Craig Volden. The mechanisms of policy diffusion. *American journal of political science*, 52(4):840–857, 2008.
- Charles R Shipan and Craig Volden. Policy diffusion: Seven lessons for scholars and practitioners. *Public Administration Review*, 72(6):788–796, 2012.

- Zachary J Siegel. Recall me maybe: The corrosive effect of recall elections on state legislative politics. *U. Colo. L. Rev.*, 86:307, 2015.
- Joshua Spivak. Recall elections in the us: Its long past and uncertain future. In Yanina Welp and Laurence Whitehead, editors, *The Politics of Recall Elections*, pages 73–93. Springer, 2020.
- David Strang. Adding social structure to diffusion models: An event history framework. *Sociological Methods & Research*, 19(3):324–353, 1991.
- Martín Tanaka. *Los espejismos de la democracia: el colapso de un sistema de partidos en el Perú, 1980-1995, en perspectiva comparada*. Instituto de Estudios Peruanos, 1998.
- Fernando Tuesta Soldevilla. Perú: Entre la participación y la gobernabilidad local. In Yanina Welp and Uwe Serdült, editors, *La dosis hace el veneno: Análisis de la revocatoria del mandato en Suiza, Estados Unidos y América Latina*, pages 7–30. Consejo Nacional Electoral, 2014a.
- Fernando Tuesta Soldevilla. Las revocatorias en el Perú: entre la participación y la gobernabilidad local. In Fernando Tuesta Soldevilla, editor, *Una onda expansiva. Las revocatorias en el Perú y América Latina*, pages 45–66. Jurado Nacional de Elecciones, 2014b.
- Anne Twomey. Second thoughts—recall elections for members of parliament. *Sydney Law School Research Paper*, 11/55, September 2011.
- Steven M Van Hauwaert. On far right parties, master frames and trans-national diffusion: understanding far right party development in western europe. *Comparative European Politics*, 17:132–154, 2019.
- Craig Volden, Michael M Ting, and Daniel P Carpenter. A formal model of learning and policy diffusion. *American Political Science Review*, 102(3):319–332, 2008.
- Stefanie Walter. Brexit domino? the political contagion effects of voter-endorsed withdrawals from international institutions. *Comparative Political Studies*, 54(13):2382–2415, 2021.
- Rachel Weinstein. You’re fired!, the voters’ version of “the apprentice”: An analysis of local recall elections in california. *S. Cal. Interdisc. LJ*, 15:131, 2005.
- Yanina Welp. ¿por qué Perú? análisis de la revocatoria del mandato en perspectiva comparada. *Elecciones*, 12(13):53–77, 2013.

- Yanina Welp. De venenos y fármacos. la regulación y prácticas de las revocatoria del mandato en suiza y las américas. In Yanina Welp and Uwe Serdült, editors, *La dosis hace el veneno: Análisis de la revocatoria del mandato en Suiza, Estados Unidos y América Latina*. Consejo Nacional Electoral, 2014.
- Yanina Welp. Recall referendums in peruvian municipalities: a political weapon for bad losers or an instrument of accountability? *Democratization*, pages 1–21, 2015.
- Yanina Welp and Ana Sofía Castellanos. Understanding the use of recall referendums: Evidence from ecuador. *International Political Science Review*, 41(3):335–348, 2020.
- Yanina Welp and Juan Pablo Milanese. Playing by the rules of the game: partisan use of recall referendums in colombia. *Democratization*, 25(8):1379–1396, 2018.
- Yanina Welp and Laurence Whitehead. Recall: democratic advance, safety valve or risky adventure? In Yanina Welp and Laurence Whitehead, editors, *The Politics of Recall Elections*, pages 9–27. Springer, 2020a.
- Yanina Welp and Laurence Whitehead. The recall revival and its mixed implications for democracy: Evidence from latin america. In Yanina Welp and Laurence Whitehead, editors, *The Politics of Recall Elections*, pages 159–178. Springer, 2020b.
- Kurt Weyland. Theories of policy diffusion lessons from latin american pension reform. *World politics*, 57(2):262–295, 2005.
- Laurence Whitehead. The recall of elected officeholders the growing incidence of a venerable, but overlooked, democratic institution. *Democratization*, 25(8):1341–1357, 2018.
- Laurence Whitehead. On reconciling recall with representation. In Yanina Welp and Laurence Whitehead, editors, *The Politics of Recall Elections*, pages 229–247. Springer, 2020.

A Supplementary Information

A.1 Descriptive Statistics

Table A.1: Descriptive Statistics: Main variables, 1996-2014

Variable	Observations	Mean	Standard Deviation	Minimum	Maximum
Dummy for district with recall	8,160	0.139	0.346	0	1
At least one neighbor had recall	6,528	0.394	0.489	0	1
Proportion of neighbors with recall	6,528	0.0984	0.149	0	1
Proportion of neighbors with successful recall	6,528	0.0277	0.0779	0	0.800
All neighbors with failed recalls	6,528	0.256	0.437	0	1
Distance to nearest district with recall	6,528	40.43	44.94	1.707	645.4
Number of districts	1,632	1,632	1,632	1,632	1,632

Table A.2: Descriptive Statistics: Additional Variables, 2002-2014


Variable	Observations	Mean	Standard Deviation	Minimum	Maximum
Mayor subjected to recall in district	4,902	0.180	0.385	0	1
At least 1/3 of council members subjected to recall	4,902	0.156	0.363	0	1
Experience: any elected position	4,902	0.636	0.899	0	7
Non-electoral conflicts	3,268	0.0652	0.247	0	1
Electoral conflicts	3,268	0.00887	0.0938	0	1
Conflicts against district authorities	3,268	0.00306	0.0552	0	1
% Annual budget executed	1,634	73.95	15.68	11.42	99.50
Margin of victory	4,902	9.177	9.093	0	100
ENEP	4,902	4.650	1.632	1	14.04
Same party as provincial mayor	4,902	0.248	0.432	0	1
Same party as regional president	4,902	0.168	0.374	0	1
N of prior recalls in district	4,902	0.290	0.577	0	4
% Illiterate	4,902	13.95	8.970	0.220	50.41
% with at least 2 UBN	4,902	25.08	15.37	0	92.41
Infant mortality rate	4,902	23.88	8.082	9.100	53.90
% Rural	4,902	53.92	30.36	0	98.69
District budget per capita (original)	1,634	741.2	830.2	50.34	14,284
District budget per capita (amended)	1,634	1,538	2,523	66.09	62,766

Tables A.1 and A.2 present summary statistics. It is worth noting that during the 2002-2014 period almost 20% of district-government terms in Peru experienced recalls. The majority of these districts subjected at least one third of the council to the recall—the

lower bound for triggering a new election—, with their overall distribution in the data reaching almost 16%. The number of previous recalls in a district has a maximum value of 4, indicating that by 2011 some districts (3, to be precise) had experienced a recall every term since 1996. Another aspect that stands out is the great variation in budget execution across districts, ranging from under 12% to almost 100% of the annual budget. Similarly, the margin of victory can be as small as 0%, which is the case in districts in which there was a tie and the winner was decided by a coin toss.

A.2 Additional Results

Figure A.1: Recalls by Subnational Government Term



Data source: ONPE. Dotted line indicates the share of districts with recalls in which new elections could be triggered (at least 1/3 of council members are subjected to the recall).

Figure A.2: Districts with close elections and recalls, 1998-2014




Table A.3: Is there learning? Restricted sample 2002-2014

	Probability of a Recall			
	(1)	(2)	(3)	(4)
At least one neighbor had recall	0.000 [0.015]			
Proportion of neighbors with recall		-0.001 [0.052]	-0.054 [0.055]	0.082 [0.067]
Proportion of neighbors with successful recall			0.075*** [0.026]	
Proportion of neighbors with failed recalls				-0.045** [0.020]
Observations	4,896	4,896	4,896	4,896
Districts	1632	1632	1632	1632
R-squared	0.452	0.452	0.454	0.453
Term Fixed Effects	Yes	YES	YES	YES

NOTES. Results from OLS panel regressions with district fixed effects. Cluster robust standard errors (clustered by district) are reported in brackets. The dependent variable is a binary indicator coded as 1 for districts in which a recall election took place and 0 for those in which it did not.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.4: Is there learning? Controls added, 2002-2014

	Probability of a Recall			
	(1)	(2)	(3)	(4)
At least one neighbor had recall	0.020 [0.013]			
Proportion of neighbors with recall		0.140*** [0.046]	0.091* [0.049]	0.265*** [0.059]
Proportion of neighbors with successful recall			0.070*** [0.022]	
Proportion of neighbors with failed recalls				-0.067*** [0.018]
ENEP	0.020*** [0.005]	0.020*** [0.005]	0.020*** [0.005]	0.020*** [0.005]
Same party as provincial mayor	-0.011 [0.013]	-0.011 [0.013]	-0.011 [0.013]	-0.010 [0.013]
Same party as regional president	0.009 [0.015]	0.009 [0.015]	0.009 [0.015]	0.011 [0.015]
N of prior recalls	-0.515*** [0.022]	-0.521*** [0.022]	-0.520*** [0.022]	-0.524*** [0.022]
Mayor experience	0.003 [0.007]	0.003 [0.007]	0.003 [0.007]	0.003 [0.007]
Observations	4,896	4,896	4,896	4,896
Districts	1632	1632	1632	1632
R-squared	0.568	0.569	0.570	0.571
Term Fixed Effects	Yes	YES	YES	YES

NOTES. Results from OLS panel regressions with district fixed effects. Cluster robust standard errors (clustered by district) are reported in brackets. The dependent variable is a binary indicator coded as 1 for districts in which a recall election took place and 0 for those in which it did not. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.5: Robustness test: distance to nearest recall

	Probability of a Recall		
	(1)	(2)	(3)
Distance to nearest district with recall	-0.000 [0.000]	-0.000 [0.000]	-0.001* [0.000]
Observations	6,528	6,189	5,834
Districts	1,632	1,603	1,571
R-squared	0.374	0.385	0.388
Year Fixed Effects	YES	YES	YES
Maximum distance	Full	95%	90%

NOTES. Results from an OLS panel regression with district fixed effects. Cluster robust standard errors (clustered by district) are reported in brackets. The dependent variable is a binary indicator coded as 1 for districts in which a recall election took place and 0 for those in which it did not.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.


Table A.6: Robustness to different measures of close elections

	Margin <8%	Margin <10%	Margin of victory
Proportion of neighbors with successful recall	0.012 [0.029]	0.005 [0.031]	0.109*** [0.033]
Close election	0.033*** [0.013]	0.043*** [0.013]	
Proportion successful × close election	0.104** [0.041]	0.097** [0.039]	
Margin of victory			-0.004*** [0.001]
Proportion successful × margin of victory			-0.004* [0.002]
Observations	4,896	4,896	4,896
Districts	1632	1632	1632
R-squared	0.573	0.574	0.574
Term Fixed Effects	YES	YES	YES
Controls	YES	YES	YES

NOTES. Results from OLS panel regressions with district fixed effects. Cluster robust standard errors (clustered by district) are reported in brackets. The dependent variable is a binary indicator coded as 1 for districts in which a recall election took place and 0 for those in which it did not. Models include controls for proportion of neighbors with recalls, ENEP, same party as provincial mayor, same party as regional president, mayor experience and number of prior recalls. Model 3 uses a continuous measure of the margin of victory truncated at 30 or over.

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

Figure A.3: Marginal effect of the proportion of successful neighbor recalls on the probability of a recall, 2002-2014



Note: Figure graphs marginal effects from model 3 in table A.6.

A.3 Exploring citizens' role

Citizens' role in driving recall diffusion is tested using the following model:

$$Recall_{dt} = \beta X_{dt} + Z'_{dt}\gamma + \lambda_d + \eta_t + \epsilon_d \quad (2)$$

This model is analogous to my baseline model in that it estimates the effect of the independent variables (X_{dt}) on the probability of a recall taking place by leveraging variation within districts, over time. Time-invariant district-level features and exogenous time-shocks are again controlled for using fixed effects. Plus, this model adds a vector of covariates (Z'_{dt}) for relevant time-varying district-level characteristics that may affect both the independent variables and the probability of a recall taking place. Like in our main analyses, these include two indicator variables for whether the district mayor comes from the same party as the provincial mayor or the regional president, the ENEP in the most recent election and the cumulative number of past recall elections. For the same reasons as above, I use data for the period between 2002 and 2014.

I examine whether recalls take place when politicians are inexperienced, incapable, or unpopular.³⁵ Experience is assessed using a variable that records the number of times

³⁵Strictly speaking this test should use as dependent variable a dummy for the occurrence of a recall

the current mayor previously held different types of elected office. If a mayor's experience does not affect the probability of a recall (or it is not the least experienced mayors that are subjected to the recall), its coefficient is expected to be zero (or positive). Data on mayors' experience comes from the National Jury of Elections (JNE), whose webpage offers information on all elected officials, including the history of elective offices they have held. On the basis of information collected by scraping this website I constructed ratio variables measuring different types of previous elected positions held by mayors (executive and legislative positions at the provincial or district levels).

To assess mayor's popularity I use a dummy variable indicating the presence of conflicts caused by dissatisfaction with district authorities. If recalls are not associated with this form of dissatisfaction with authorities (or recalled authorities are not the least popular ones), its coefficient is also expected to be zero (or negative). Data on social unrest comes from the Office of the Ombudsman's monthly report on social conflicts. This document, available since March 2004, documents the occurrence of social conflicts throughout the country, including a description of the location, causes and actors involved in the conflict, on the basis of reports from decentralised state agencies, NGOs, the media and participants' reports.³⁶ Using these reports I hand coded dummies indicating the existence of conflicts caused by dissatisfaction with district authorities. I only considered conflicts that took place during the first year of the district government's term (when recall processes cannot be initiated) to exclude conflicts that could result from the recall campaign itself. Since conflict data is only available since 2004, these analyses only include two terms (2007-2010 and 2011-2014).

Capacity is assessed on the basis of two aspects: mayor's political ability to manage social conflicts in their jurisdictions and their managerial ability to execute the municipality's budget. To test the former I include dummies for the presence of two types of conflicts in the district: non electoral and electoral conflicts. If recalls are not associated with social unrest in the district (or at least recalls do not take place in more conflictive districts), the coefficient on non-electoral conflicts is expected to be zero (or negative). The dummy for electoral conflicts is included separately to allow for the possibility that recalls are a way of solving conflicts caused by citizen's dissatisfaction with electoral results, in which case its coefficient would be positive. Data source and periods covered remain the same as above.

Finally, to test the role of managerial ability I use a variable measuring the percentage of progress in the execution of the municipal budget. Data on municipal government's

election targeted specifically against the mayor, rather than one indicating the occurrence of any recall election in district d at time t . In practice, the difference between these two dummies is small: out of 904 district-government terms with recall elections, only 20 of them (2%) did not include the mayor among the officials subjected to the recall.

³⁶Social conflicts are defined as contentious collective actions that involved at least one of the following: threats to people's life, integrity or health; damages to public or private property; limitations to free movement; obstruction of an official in the exercise of his duties; paralysation of public service provision.

budget management comes from the Ministry of Economy and Finance, which reports each district's percentage of progress in the execution of its annual budget since 2009.³⁷ On the basis of this information I constructed a variable recording progress in the execution of the municipal budget during the first year of the term, to avoid reverse causality due to the recall campaign's effect on the management of the municipal government. Because this information is only available since 2009, I only have data for the 2011-2014 term which means this analysis relies only on cross-sectional variation.³⁸ If the mayor's capacity to execute the budget is not associated with the occurrence of a recall (or at least it is not the governments that are less capable of executing their budget that are faced with a recall), the coefficient of interest is expected to be zero (or positive). Since I am now looking at the effect of variation in budget execution across districts, it is important to account for relevant differences between districts that may affect both budget execution and the occurrence of a recall. Thus, my vector of controls now includes additional covariates. To control for socio-economic differences across districts I include variables measuring (i) the percentage of rural population in the district; (ii) the percentage of illiterate adults in the district; (iii) the infant mortality rate in the district; and (iv) the percentage of population in the district with at least two unsatisfied basic needs.³⁹ These data come from the 2007 census, the only source available with data at the district level. Finally, I also include a variable measuring the per capita annual budget in each district during the first year of the term. This variable is intended to control, on the one hand, for an alternative political incentive to promote a recall, and on the other, for varying levels of administrative capacity that may affect budget execution. Since districts are embedded in provinces in some models I also include province fixed effects to avoid potential confounding due to province-level features that may affect the relationship between budget execution and the probability of a recall.

Table A.7 presents the results. Column 1 shows that mayors' prior experience in elected positions has no effect on the probability of a recall taking place. These results indicate that it is not the least experienced mayors that are being recalled.⁴⁰ Column 2 assesses the role of elected officials' popularity, as measured by the existence of social conflicts against district authorities. As we can see, districts with conflicts expressing opposition to district authorities during their first year in office are less likely to experience a recall. Column 3 examines the role of mayor's capacity as measured by the level of

³⁷The measure is constructed as the ratio of accrued expenses to the annual amended budget, expressed as a percentage.

³⁸Formally, the regression model I estimate is $Recall_{dp} = \beta ExecBudg_{dp} + Z'_{dp}\gamma + \lambda_p + \epsilon_d$.

³⁹The basic needs considered are: dwellings with inadequate physical characteristics, overcrowded dwellings, dwellings with no connection to a water network, households with at least one child aged 6 to 12 that does not attend school, households with high economic dependency (household head has incomplete primary education and more than three dependents).

⁴⁰I also ran models using different types of experience: at the district and provincial level, and in executive and legislative positions. Coefficients are consistently small and insignificant in all models.

social unrest in the first year of their term. It shows that while electoral conflicts have no effect on the probability of a recall, other social conflicts do. Districts with non-electoral conflicts are found to be less likely to have a recall. These results are the opposite of what we would find if citizens were recalling authorities that are unpopular or unable to manage social conflicts in their jurisdiction. In that sense, they are consistent with the idea that recalls are not used to get rid of incapable officials.

Table A.7: Who Drives Diffusion? Testing Citizens' Role

	Probability of a Recall		
	(1)	(2)	(3)
Experience: any elected position	0.003 [0.007]		
Conflicts against district authorities		-0.190* [0.101]	
Non-electoral conflicts			-0.093*** [0.033]
Electoral conflicts			-0.074 [0.092]
Same party as provincial mayor	-0.011 [0.013]	-0.021 [0.017]	-0.021 [0.017]
Same party as regional president	0.010 [0.015]	0.043** [0.019]	0.044** [0.019]
ENEP	0.019*** [0.005]	0.017** [0.007]	0.016** [0.007]
N of prior recalls	-0.514*** [0.022]	-0.813*** [0.031]	-0.816*** [0.031]
Observations	4,902	3,268	3,268
Districts	1634	1634	1634
R-squared	0.568	0.740	0.741
Term Fixed Effects	YES	YES	YES

NOTES. Results from OLS panel regressions with district fixed effects. Cluster robust standard errors (clustered by district) are reported in brackets. The dependent variable is a binary indicator coded as 1 for districts in which the mayor was subjected to a recall vote and 0 for those in which she was not.

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

Table A.8 presents the results of the pooled OLS regressions assessing the link between subnational governments' capacity in the management of the district's budget and the probability that a recall takes place.⁴¹ As expected, the full model shows no association between these variables once political covariates are accounted for.⁴² We can therefore conclude that at least in the 2011-2014 term, recalls were not associated with officials'

⁴¹This table reports 194 provinces because even though 195 provinces existed in Peru at this time, one of them -Purus- does not have any district-level governments.

⁴²The following alternative model specifications were also tested but did not significantly alter the results: (i) using the amended per capita annual budget instead of the original per capita annual budget; (ii) dropping the variable measuring illiteracy (it is correlated with the other socioeconomic controls and may be inflating the standard errors); and (iii) using regional fixed effects instead of provincial ones.

incapacity to execute the budget.

Table A.8: Who Drives Diffusion? Testing Citizens' Role II

Variables	Probability of a Recall		
	(1)	(2)	(3)
Percent annual budget executed	0.002** [0.001]	0.001 [0.001]	0.001 [0.001]
Percent rural	0.000 [0.001]	0.000 [0.001]	-0.000 [0.001]
Infant mortality rate	-0.001 [0.002]	-0.001 [0.002]	0.031** [0.015]
Percent with at least 2 UBN	0.000 [0.001]	-0.000 [0.001]	-0.001 [0.001]
Percent illiterate	-0.003 [0.002]	-0.002 [0.002]	-0.001 [0.003]
District budget per capita	0.000** [0.000]	0.000* [0.000]	0.000 [0.000]
Same party as provincial mayor		-0.046** [0.022]	-0.061** [0.025]
Same party as regional president		-0.001 [0.027]	-0.001 [0.030]
ENEP		-0.006 [0.007]	0.007 [0.008]
Previous recalls in district		0.119*** [0.017]	0.076*** [0.018]
Observations	1,634	1,634	1,634
Provinces	194	194	194
R-squared	0.011	0.054	0.213
Province Fixed Effects	NO	NO	YES

NOTES. Results from pooled OLS regressions. Cluster robust standard errors (clustered by province) are reported in brackets. The dependent variable is a binary indicator coded as 1 for districts in which a recall election took place and 0 for those in which there was no recall.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.