JOURNAL OF THEORETICAL POLITICS

Article

Simulating policy diffusion through learning: Reducing the risk of false positive conclusions

Journal of Theoretical Politics
2016, Vol. 28(3) 497–519

DOI: 10.1177/0951629815581461
jtp.sagepub.com
SAGE

Christian Adam
Ludwig-Maximilians-University of Munich, Germany

Abstract
This article uses agent-based computer simulation to investigate the dynamics of policy diffusion through learning. It compares these dynamics across state systems in which policy-makers possess different capabilities to learn about policy effectiveness: independent decision-makers focusing on own experiences vs. interdependent social learners relying heavily on experiences of others. The simulation can thus compare policy adoption patterns in the presence and absence of policy diffusion within a controlled setting. The simulation makes two propositions. First, it supports the existing critique that relying on the identification of policy clusters can lead researchers to draw false positive conclusions about the relevance of policy diffusion. Second, it suggests that relying on the identification of policy volatility under political stability minimizes this risk.

Keywords
Agent-based modeling; diffusion; policy learning; policy volatility; simulation

1. Introduction

While empirical research on policy diffusion predominantly concludes that policy decisions are interdependent and policy change often results from learning-induced diffusion processes (e.g. Gilardi, 2010; Meseguer, 2009; Simmons and Elkins, 2004), critical voices have come to doubt the validity of empirical evidence generally taken to support these conclusions. Volden et al. (2008) argue that the

Corresponding author:
Christian Adam, Ludwig-Maximilians-University of Munich, Oettingenstrasse 67, 80538 Munich, Germany.
Email: Christian.Adam@gsi.uni-muenchen.de

empirical patterns of policy adoptions taken to indicate policy diffusion could also have emerged through individual problem-solving by independent states. Essentially, this point of criticism suggests that conventional tests used to assess whether or not policies diffuse—such as spatial lag regression—can at times suggest policy diffusion when policy decisions are really independent. In statistics, this would be referred to as 'type-I error' or 'false positive result' (e.g. Cohen, 1977: 4–5): the null hypothesis of independent policy decisions is rejected although policy adoptions really were independent. By solely relying on the identification of policy clusters to draw conclusions about policy diffusion, we thus risk overestimating the relevance of policy diffusion under the impression of false positive results. Against this background, Volden et al. (2008) suggest investing theoretically in order to identify characteristics of 'behavior that emerges only in a cross-state learning environment' (Volden et al., 2008: 329). Such characteristics can be used by empirical research to better distinguish between interdependent processes of policy diffusion and isolated, independent policy decisions.

This article proposes that agent-based modeling and computerized simulation can be helpful tools supporting the development of innovative theoretical propositions about effects and characteristics of policy diffusion through learning. Therefore, a computerized simulation is presented that models the effects of different forms of learning (individual learning and social learning) and the effect of imperfect information on the dynamics of policy adoption patterns in a closed network of 49 states.[1] The simulation makes two propositions. First, it supports arguments pointing to the risk of false positive conclusions about policy diffusion when these conclusions rely on the identification of policy clusters. Second, it proposes that this risk can be minimized by focusing on the identification of policy volatility. Policy volatility refers to the frequency with which policy decisions are reversed. Policy decisions are considered volatile when initial policy decisions are reversed at least once (i.e. number of policy changes ≥ 2). The simulation suggests that policy volatility under political stability is a genuine feature of state systems consisting of interdependent social learners (policy diffusion).

The article first presents a short review of the relevant literature. Second, the simulation's theoretical foundations are presented. Third, the article outlines the basic structure of the simulation. Fourth, results are presented. Finally, the article concludes with a brief discussion of these results and their implications.

2. Learning, policy volatility, and computerized simulation

Since the 1970s, a vast literature on the role of learning in policy-making has emerged. This literature shares the assumption that policy-making is not simply the result of power politics and the interplay of material interests but also the attempt to solve problems objectively. Learning – defined as the acquisition of new information about the ability of specific courses of action to solve the problem at hand (Sabatier, 1988) – is key to solving problems. For some time, however, most of these efforts to analyze the role of learning in policy-making remained conceptual and only a few empirical attempts to isolate the role of learning in policy-making were undertaken (Bennett and Howlett, 1992). The theoretical literature on learning has accumulated a volume of different approaches to learning: Scholars distinguish between political learning (Heclo, 1974), governmental learning (Etheredge and Short, 1981), lesson-drawing (Rose, 1991), and social learning (Hall, 1993), to name just a few. Lately, this list has been complemented by the distinction between 'pure learning' and 'competitive learning' (Ward and John, 2013).

Ever since the call for more empirical work by Bennett and Howlett in 1992, a series of scholars have taken up the challenge to analyze policy diffusion through learning empirically. Because cognitive processes of learning by decision-makers about public policy are difficult to observe, these studies often rely on spatial patterns of policy adoption (e.g. Gilardi, 2010; Meseguer, 2009; Simmons and Elkins, 2004). While the observation of clustered policy adoption is taken as evidence for policy diffusion, critics argue that such clusters could just as well be the result of individual problem-solving by independent states dealing with similar problems (e.g. Volden et al., 2008). If we knew about additional indicators for policy diffusion through learning besides spatial clusters of policy adoption, we could reduce the risk to mistake independent for interdependent policy decisions.

In search for such additional features of interdependent states, the literature on complex systems—which specifically deals with phenomena emerging in systems characterized by interaction and interdependency—seems to be a promising source of inspiration (Axelrod and Cohen, 2000; Miller and Page, 2007; Vicsek, 2002). States engaging in policy learning can be interpreted to form a complex social system in which individual states seek information about policy effectiveness and where the experiences and perceptions of one state influences the experiences and perceptions elsewhere. The interactions which state A has with state B are thus not independent from the interactions that state B previously had with state C. Agent-based computer simulation has proved to be an effective tool for analyzing behavior of such complex systems. It is used to investigate the dynamics of processes resulting from complex social interactions. It allows models to contain heterogeneous agents and does not restrict the researcher to create models with closed-form, analytic solutions (Miller and Page, 2007). The advantage of this approach over conventional approaches to formal modeling is that it is designed to enable researchers to focus on the characteristics of process instead of equilibrium and to assess shifts dynamically and not through comparative statics. This way, agent-based computer simulation has successfully inspired theoretical debates in different disciplines, such as sociology or economics (Axelrod, 1997; Strang and Macy, 2001; Vriend, 2000). In the context of political science, this approach has been fruitfully introduced to research on party competition (Kollman et al., 1992), electoral systems (Clough, 2007), and conflict research (Cederman, 2003; Lustick et al., 2004). Against this background, it is surprising that, besides two notable exceptions (Luyet, 2009; Rapaport et al., 2009), there are—to my knowledge—no attempts to apply this tool to the study of policy diffusion.

The ability of agent-based modeling to focus on process instead of equilibrium also allows the analysis to focus on a concept hardly recognized and often only integrated implicitly in the analysis of policy diffusion: policy volatility. Policy volatility captures how often policy changes are adopted within a certain period of time. Naturally, this implies that adoption patterns need to be observed not only until a policy innovation is introduced for the first time, but also thereafter to capture potential subsequent policy decisions (i.e. the abolishment of the innovation and even potentially its re-introduction). Several empirical studies have recognized the importance of taking this into account. Simmons and Elkins (2004) allow for vacillations between liberal and protectionist economic policies in their analysis of policy diffusion. Boehmke and Witmer (2004) do not only keep track of policy innovation, but also capture subsequent policy expansion in their study of the diffusion of sub-national gambling policy. Similarly, in the context of hospital financing reforms, Gilardi et al. (2009) not only capture initial policy changes, but also subsequent reforms which happen after the initial reform. Policy volatility is a direct measure of how frequently states vacillate between different policy choices. And despite the fact that many policy decisions are highly persistent, other policy decisions are clearly volatile. Of course, policy change is often the result of ideological changes in government composition. Yet, governments do at times take back their own previous policy decisions and even re-introduce them after some time once again. One example for volatile policy decisions is the privatization of public utilities at the local level in Germany only to re-municipalize them again later on (Pigeon, 2012). Another example for policy volatility under political stability is the decision of the conservative Bavarian state government coalition in Germany to first introduce tuition fees for higher education only to abolish these fees again later (Die Zeit, 2013). Orthography reform in Germany can also be taken as an example of a volatile policy decision. In 1996, the conference of state education ministers,[2] including the ministers from Schleswig-Holstein, decided to introduce a large-scale reform of spelling in the German language. In 1998 Schleswig-Holstein passed legislation to undo the previously adopted reform, only to join the other state governments to re-introduce a similar reform in 2006.[3] These examples illustrate that policy-makers at times vacillate between different policy options.

The analytical leverage that the concept of policy volatility can develop is evident in the context of theorizing another mechanism of policy diffusion; namely, competition. It is the recurring change of the status quo policy which is at the heart of the prominent theoretical proposition that policy interdependence between different policy-makers based on competition creates a ‘race to the bottom’ (Gilardi, 2013: 462–463). Against this background, it seems worthwhile to dedicate more attention to the concept of policy volatility in the context of other diffusion mechanisms; such as learning. Consequently, this article moves beyond the conventional assessment of policy adoption patterns in terms of spatial clusters to assess also these patterns in terms of policy volatility.

3. Theoretical foundation of the simulation

The simulation uses the formal model of policy change and diffusion by Braun and Gilardi (2006) as theoretical foundation. This section summarizes how their model integrates learning – among other diffusion mechanisms – into a coherent framework of policy change and diffusion. To model policy change, the authors formulate a situation in which decision-makers are confronted with a choice between two policies i with $i \varepsilon [sq;other]$. Policy $i = sq$ refers to the policy currently in place (status quo), while policy $i = other$ refers to the alternative (other) policy which could potentially be adopted. Essentially, Braun and Gilardi argue that policy change is a matter of comparison between the utility received from the current policy and the expected utility from an alternative policy. The utility associated with a certain policy is then a function of its ability to generate votes V_i and its policy-related payoffs resulting from ideological considerations P_i. Against this background, the authors assume that decision-makers' utility from policy i take the form

$$U_i = wV_i + (1 - w)P_i \tag{1}$$

The factor w weighs the importance of vote-payoffs V_i and ideological policy-payoffs P_i, respectively. To determine whether to change from the current policy to an alternative policy, decision-makers compare the current utility against the expected utility of the alternative policy. To construct this decision, Braun and Gilardi add additional parameters. First, they allow for unsuccessful attempts to adopt the new policy by introducing the probability of a successful policy change p. Second, they include reform costs associated with policy change C. Third, they introduce policy effectiveness E_i to the equation. This implies that, in the model, politicians do not only care about votes and ideology, but also about a policy's ability to solve certain problems. Against this background, the authors expect change whenever the expected utility of attempting policy change is greater than the utility from maintaining the status quo. Change happens if

$$EU(change) > U(status\,quo) \tag{2}$$

$$pE_{other}U_{other} + (1 - p)E_{sq}U_{sq} - C > E_{sq}U_{sq} \tag{3}$$

$$E_{other}U_{other} - E_{sq}U_{sq} > \frac{C}{p} \tag{4}$$

Learning—understood as the acquisition of new information about the effectiveness of policy alternatives to solve a certain problem—affects this policy change equation (4) via the factor E_i. Other things equal, a high effectiveness of the alternative policy E_{other} and a low effectiveness of status quo policy E_{sq} make the attempt to change policies more likely.

The simulation presented here simplifies the work of Braun and Gilardi in three respects. First, it introduces a modest reinterpretation of the condition for policy change above. In line with Braun and Gilardi, the right-hand-side of the equation, $\frac{C}{p}$, can be interpreted as the threshold for change T^C. For high costs of change and low probabilities of success, the threshold for change increases. It decreases for low costs of reform and high probabilities of success. Second, the simulation interprets the left-hand-side of equation (4) as the utility from change U^C. The utility from change is only positive if the payoffs from the alternative policy outweigh the payoffs of the status quo. For high levels of effectiveness and a high utility associated with vote-payoffs and policy-payoffs from the alternative policy, this utility from change increases. It decreases with high effectiveness, high vote-payoffs, and high policy-payoffs from the status quo policy. Hence, equation (4) can be reformulated into

$$U^C_{k;t} > T^C_k \tag{5}$$

We thus expect policy change if the utility from change U^C of state k at time t is greater than its threshold for change $T^C_{k;t}$ at time t. To avoid a deterministic process, in a deviation from Braun and Gilardi, the simulation translates this condition for policy change into a probabilistic function of policy change. This is accomplished by introducing a stochastic function of logistic type.[4] This results in the following equation (6) reflecting the probability of change

$$Prob(change) = \frac{1}{1 + \left(\frac{U^C_{k;t}}{T^C_{k;t}}\right)^{-g}} \tag{6}$$

Equation (6) implies that the probability of policy change approaches 1 as the utility from change at time t, $U^C_{k;t}$, grows in relation to the threshold for change $T^C_{k;t}$. As the utility from change decreases in relation to the threshold for change, the probability of change approaches 0.

3.1. Two scenarios of learning

The simulation defines two scenarios which differ only with respect to the capabilities of states to learn about policy effectiveness. Learning is conceptualized in the simulation as states' ability to memorize and compare past experiences. Specifically, learning is implemented with the help of two components: comparing mean policy effectiveness and observing the trend of policy effectiveness.

Scenario 1: individual learning – no diffusion. In the first scenario, states are constructed as individual learners. This means that states are able to memorize and compare policy effectiveness, but they only rely on their own experiences. In consequence, states which have experience with only one policy can only take information about the effectiveness of this policy into account. The effectiveness of the other policy remains unclear. Therefore, policy change in this scenario is influenced by learning but is not driven by diffusion. The utility that states receive from changing to the other policy takes the form of

$$U^C_{k;t}(individual\,learning) = \frac{1}{E^k_{sq;t}} * \frac{mean(E^k_{other})}{mean(E^k_{sq})} * \frac{E^k_{sq;t-1}}{E^k_{sq;t}} \tag{7}$$

Here, $E^k_{sq;t}$ represents the effectiveness of the status quo policy. Any payoff related to ideological or electoral concerns is not considered explicitly. Since this article is interested in the effects of learning about policy effectiveness, these payoffs are kept constant over policy alternatives and time. States thus evaluate the effectiveness of the current status quo policy to solve the problem at hand, which states weigh with two additional factors.

The first of these factors, $\frac{mean(E^k_{other})}{mean(E^k_{sq})}$, reflects states' ability to compare. Specifically, this factor tries to incorporate states' ability to compare the mean effectiveness they received from the alternative policy, E^k_{other} (if they already had adopted this alternative policy for at least one time period), with the mean effectiveness they received from their current status quo policy, E^k_{sq}, in past periods. If states have experienced a mean effectiveness of the alternative policy greater than the mean effectiveness of the status quo policy, this factor is greater than 1 and the utility from change increases. If states have experienced a greater mean effectiveness of the status quo policy than of the other policy, the utility from change decreases.

The second factor, $\frac{E^k_{sq;t-1}}{E^k_{sq;t}}$, reflects states' ability to observe the trend of policy effectiveness. This should incorporate the intuition that states are less likely to give up their status quo policy if this policy shows a positive trend, i.e. if its effectiveness in this period is higher than in the previous period. In this case: $\frac{E^k_{sq;t-1}}{E^k_{sq;t}} < 1$, which reduces the utility from change. Should states have experience with only one policy or only have experience with the status quo policy in one period, the respective factor (mean and/or trend) defaults to 1.

Scenario 2: social learning—diffusion. In a second scenario, policy change is modeled as a process of policy diffusion through social learning. This means (again) that states have the ability to compare mean policy effectiveness and to observe the trend of policy effectiveness. However, in this scenario, states not only consider their own experience but also mainly focus on the experience of other states in the state system; their neighbor states. In this sense, the term 'social learning' simply means learning from others. In consequence, the equation for the utility from change essentially takes the same form as in the previous scenario. State k does not only consider the effectiveness of its own status quo policy at time t ($E^k_{sq;t}$), but also the experiences of neighboring states N that currently have state k's status quo policy ($E^N_{sq;t}$) or the other policy alternative (E^N_{other})

$$U^C_{k;t} = \frac{1}{E^k_{sq;t}} * \left(\frac{mean(E^N_{other})}{mean(E^N_{sq})} * \frac{E^N_{sq;t-1}}{E^N_{sq;t}} \right)^d \tag{8}$$

To ensure that policy decisions in this scenario are the result of diffusion, the outside experience of neighbors is given more weight than states' individual experiences. This is accomplished by adding the power d, with $d > 1$, to the equation. This way, a positive experience of neighboring countries with a state's status quo policy $\frac{mean(E^N_{other})}{mean(E^N_{sq})} < 1$ and $\frac{E^N_{sq;t-1}}{E^N_{sq;t}} < 1$ further reduces the utility from change. In turn, a negative experience results in an ever higher utility from change.

4. Structure of the simulation

The simulation defines a grid of 7×7 squares where each square represents an agent, i.e. a state. This way, the simulation includes 49 states approximating the number of continental US states often used in the context of spatial lag regressions. Since states are arranged within a square, each state has eight neighbors. These can be interpreted to be geographic neighbors, cultural neighbors, or economically related states. Only states located at the fringes of the square have fewer neighbors. At every time step in the simulation, each state chooses between two policies. This means that states consider whether to stick to their status quo policy or to change to the alternative policy. One policy is constructed to be more effective than the other policy in solving a problem common to all states. Initially, each state is assigned the less effective policy. Two factors are essential for the decision to adopt an alternative policy in all scenarios of the simulation: the state's threshold for change ($T^C_{k;t}$) and the observed policy effectiveness ($E^k_{i;t}$).

To make states heterogeneous, thresholds for change $\left(T^C_{k;t}\right)$ are assigned to each state by drawing random numbers from a uniform distribution at the beginning of the simulation ($t=0$). This way, the simulation tries to account for the fact that political constraints and institutional frictions are not the same for all states in the real world.

One run in the simulation comprises 10 time steps. This could be interpreted as one run representing governments making a decision of maintaining or changing the status quo policy once every year for a 10-year period in office or twice every year for a five-year period in office. Since policy learning is simulated under political stability, such a low number of time steps seems appropriate.[5]

To account for policy volatility, states' decisions to give up their status quo policy and adopt the policy alternative are modeled so that they are not necessarily definitive. In other words, states do not have to stick to their first policy decision forever. Instead, states can potentially change back and forth between the more and the less effective policy if they want to. Yet, the simulation accounts for the political costs that arise from such volatility. Correcting own policy decisions and correcting the corrections is rarely a politically opportune endeavor. While new information about policy effectiveness might call for such corrections, changing your mind about something while in office can be used for political attacks by opposition parties and the media and thus potentially carries costs. To implement this argument in the simulation, the threshold for policy change $\left(T^C_{k;t}\right)$ doubles after every policy change. Correcting a previous policy decision is thus twice as hard as making the previous decision. With every policy change by a certain government, the government will need a much higher perceived effectiveness of the alternative policy in order to be willing to change its status quo policy one more time.

One of the advantages of simulation studies is that they allow researchers to conduct a baseline run and then experimentally alter parameter values to see how these alterations affect the results within a controlled environment. This opportunity is used to see how the introduction of imperfect information influences the dynamics of policy adoption in both scenarios. Thus, a baseline run is conducted in which states are able to observe objectively the effectiveness of each policy alternative.

Table 1. Simulation parameters.

		Parameter name in simulation code	Parameter value	
			Baseline run	Experimental run
Threshold for change		threshold	random draws from a uniform distribution with minimum of 3 and a maximum of 6	random draws from uniform distribution with minimum of 3 and a maximum of 6
Factor of threshold increase after each policy change		incr	2	2
Time periods per simulation run		repeat.times	10	10
Policy payoff		P	random draws from normal distribution with minimum of 0.7 and a standard deviation of 0.2	random draws from normal distribution with minimum of 0.7 and a standard deviation of 0.2
Vote payoff		V	random draws from normal distribution with minimum of 0.7 and a standard deviation of 0.2	random draws from normal distribution with minimum of 0.7 and a standard deviation of 0.2
Weight		w	random draws from uniform distribution with minimum of 0 and a maximum of 1	random draws from uniform distribution with minimum of 0 and a maximum of 1
Percentage of states with less effective policy as initial policy		initial.perc	100%	100%
Diffusion parameter		d	2	2
Power in logistic function determining the probability of policy change		g	-5	-5
Number of states		n.states	49	49
Less effective policy	effectiveness	e0	1	1
	standard deviation	sde0	0	0.25
More effective policy	effectiveness	e1	2	2
	standard deviation	sde1	0	0.25

Subsequently, imperfect information is imposed accounting for the fact that assessing a policy's effectiveness is no trivial task. Instead, information regarding policy effectiveness is likely to contain 'noise' which obscures the 'true' policy effectiveness. Table 1 summarizes the parameter settings for the baseline run and the experimental run. In the baseline run, the policies' levels of effectiveness are fixed to 1 and 2. In the experimental run, I introduce uncertainty by drawing the respective policy effectiveness randomly from two normal distributions with mean 1 and 2 and a standard deviation of 0.25. In this sense, a higher standard deviation reflects noisier information regarding policy effectiveness. Each scenario (consisting of 10 time steps) is run 100 times to capture probabilistic variation.

This approach towards imperfect information is similar to the one chosen by Volden et al. (2008) in that a policy's true level of effectiveness is only realized with a certain probability. Yet, in contrast to the binary approach adopted by Volden et al. (2008), this article, with its continuous concept of policy effectiveness, includes imperfect information by introducing 'noise' as a continuous probabilistic component to states' perception leading states randomly to underestimate or overestimate policy effectiveness. Furthermore, the article diverges from the approach by Volden et al. (2008) in that the knowledge about imperfect information does not turn states into strategic learners that can decide to hold back on policy experimentation in a first period and profit from the experiences made elsewhere for policy choices in a second period. Instead, this article assumes states to be social learners. It adopts a decision-theoretic perspective assuming that an alternative policy is chosen whenever evidence about the superiority of this policy over the status quo policy is sufficient.

All relevant simulation parameters are displayed in Table 1. While the simulation is, of course, sensitive to the specification of values of key parameters, Table 4 in Appendix A shows that the results are robust for a broad range in parameter values.

5. Simulation results

Figure 1 displays the results of the simulation. It shows the aggregate adoption patterns of the four simulated scenarios which include policy adoption patterns in the individual learning scenario with perfect (square I) and imperfect (square II) information as well as patterns in the diffusion scenario with perfect (square III) and imperfect information (square IV). These plots capture the percentage of states in the simulation that possess the more effective policy at each of the 10 time steps in an individual run of the simulation. Such a run is repeated 100 times for each scenario. The limits of the grey areas in these plots indicate the minimum and maximum share of states that possess the more effective policy at each point in time. The black line in the middle of the grey areas captures the mean share of states with the more effective policy as average over all of the 100 runs of the simulation.

Figure 1 illustrates that temporal patterns of policy diffusion differ visibly between the scenarios. When states receive perfect information, the difference between individual learning and social learning (diffusion) is most apparent. In the individual learning scenario (square I) adoption patterns resemble a slow linear increase. On average, about 20% of the states have adopted the more effective

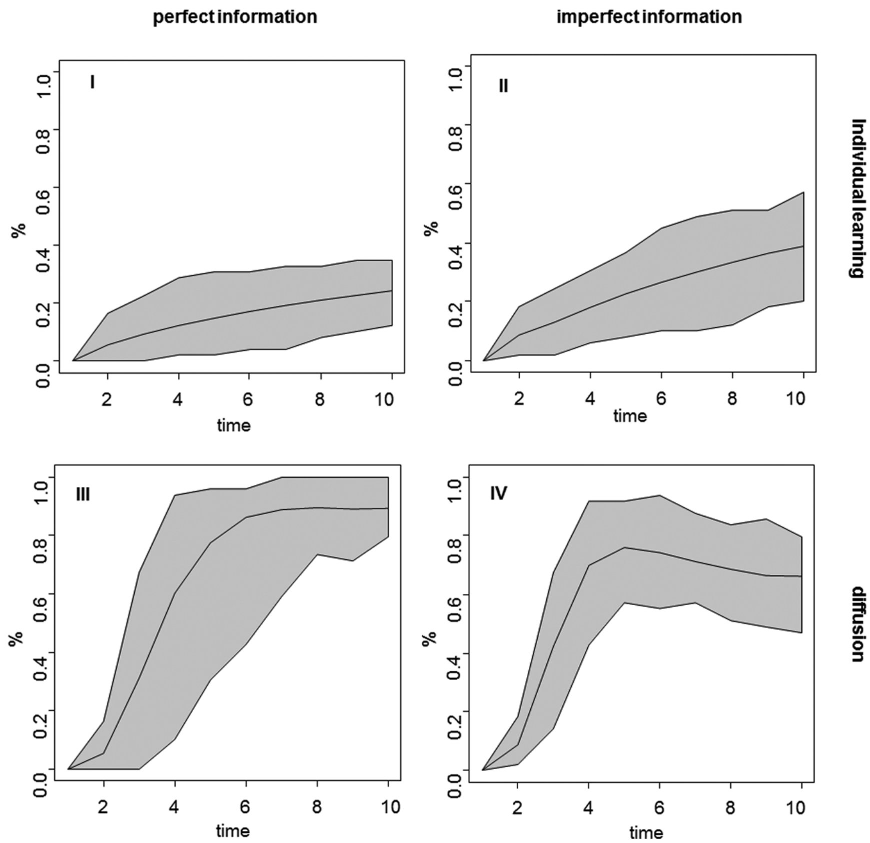

Figure 1. Spread of the more effective policy.

policy at the end of the 10 time steps. When states have the opportunity to take the experience of other states into account (square III), we observe a greater spread of the more effective policy as at the end of 10 time steps more than 90% of states have on average adopted the more effective policy. The spread is not only greater but also achieved relatively quickly.

The introduction of imperfect information changes the dynamics of policy adoption in both scenarios visibly. First, both scenarios move closer together in terms of how far the more effective policy has spread after 10 time steps. This is a result of a wider spread of the more effective policy in the individual learner scenario (square II) and a weaker spread in the diffusion scenario (square IV). This effect has two sources. First, the introduction of imperfect information enhances states' willingness to engage in policy experiments more quickly in both scenarios. This is most visible in the individual learner scenario, but also present in the diffusion scenario where the grey area is much narrower in the first couple of time steps under imperfect information than under perfect information because runs with a slower spread

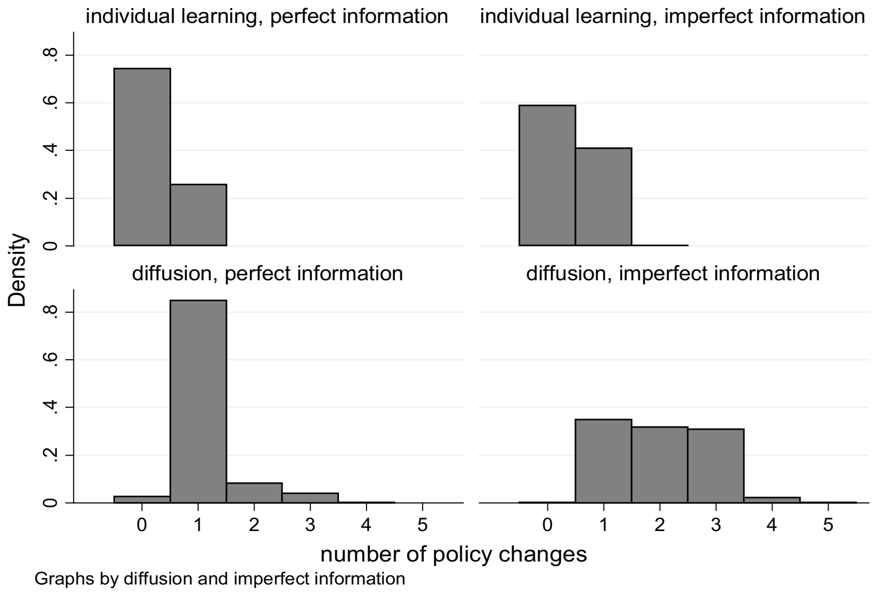

Figure 2. Policy volatility—frequency of policy change.

of policy adoptions are eliminated. As long as states only have information about the effectiveness of the initial policy, they only abandon this initial policy and experiment with the alternative policy when their perception of the effectiveness of the current policy is sufficiently low. Imperfect information sometimes leads states to underestimate (and overestimate) this level of policy effectiveness. Due to the underestimation of policy effectiveness, the threshold for change is reached in more states quicker under imperfect than under perfect information. Second, this effect is the result of an enhanced degree of policy volatility in both diffusion scenarios (see Figure 2). Social learners seem to be more likely to correct previous policy experiments than are individual learners. With imperfect information, this volatility in the social learning scenario is so enhanced that even the aggregate share of states with the more effective policy decreases slightly after a quick initial increase. This aspect of policy volatility will be discussed in the next section.

5.1. Patterns of policy volatility

One of the key results of the simulation is its suggestion that the presence of policy volatility is a genuine feature of the diffusion scenarios. Figure 2 captures the degree of policy volatility in terms of the number of policy changes per state including density estimates in the form of four histograms.

Policy volatility refers to the reversal of previous policy decisions. Specifically, I consider policies to be volatile when initial policy decisions are reversed at least once (i.e. number of policy changes ≥ 2). In this case, policy-makers oscillate

Table 2. Determinants of policy volatility.

	Poisson regression model	
	Model I	Model II
diffusion	1.55**	1.49**
	(0.02)	(0.03)
imperfect information	0.55**	0.47**
	(0.02)	(0.04)
diffusion X imperfect information		0.09*
		(0.04)
constant	−1.41**	−1.36**
	(0.02)	(0.03)
Pseudo R	0.20	0.20
N	19600	19600

Remarks: Coefficients represent unstandardized Poisson regression coefficients. Standard errors in brackets. *and ** reflect statistical significance on the 5% and 1% level. Data are generated by the simulation.

between the two policy options. Since this article is interested in patterns of policy diffusion under political stability, and assumes that reversing one's own previous policy decisions imposes political costs, it is not surprising that policy volatility is a rather rare phenomenon.

Yet, how rare it is, seems to crucially depend on whether policy decisions are taken independently by individual learners or interdependently by social learners. More specifically, under perfect information none of the 49 states constructed as individual learners changed its policy more than once within the period of 10 time steps over 100 runs of the simulation. States either never changed their initial policy or changed it only once. This picture hardly changes when imperfect information is introduced. While this generally enhances states' willingness to experiment with a policy alternative quicker—zero policy changes occur less frequently—imperfect information does not trigger policy volatility in the individual learner scenario (in only four out of 4900 possible occasions are two policy changes recorded). In contrast, when states are constructed as interdependent social learners, policy volatility becomes a relevant feature of policy dynamics. Even with perfect information about policy effectiveness, states change policies twice, three times, and even four times in over 12% of the time. In the scenario with imperfect information, states adopt two policy changes in about 32% and three policy changes in about 31% of the time. One consequence of this is that states with an even number of policy changes finish the simulation run with the same (less effective) policy with which they orginally started.

Table 2 complements these descriptive statistics of nominal frequencies by presenting results of a regression conducted on the basis of data from the simulation. The dependent variable in the regression is policy volatility in each state as captured by the frequency with which each state is observed to switch policies within 10 time steps in the observation. Due to the count data character of the dependent variable, i.e. the fact that it can only take values of non-negative integers, a Poisson regression model seems appropriate (Winkelmann, 2008). The number of observations is the result of 49 states in the simulation for which the number of policy switches are recorded over 10 times steps for 100 runs for each of four scenarios (49 states×4 scenarios×100 runs = 19,600 observations containing the number of policy changes over 10 times steps). These runs differ since there are probabilistic elements built into the simulation. For example, policy effectiveness at each point in time for each individual state is drawn randomly from a normal distribution (see Table 1). The data analyzed can thus be interpreted as a random sample of the general population of all possible simulation runs. The different scenarios are identified by dummy variables for *diffusion* and *imperfect information*. If the coefficients for these dummy variables were sufficiently small in comparison to estimated standard errors, the regression analysis would reject the argument that the degree of policy volatility varied systematically instead of randomly between the scenarios. It thus assesses the null hypothesis of no systematic difference in the degree of policy volatility between the different scenarios.

Regression results support the argument of systematic variance in the degree of policy volatility between the scenarios. Coefficients for the dummy variables *diffusion* and *imperfect information* are positive with standard errors sufficiently small to accept statistical significance on conventionally accepted levels. Both policy diffusion through social learning and the presence of imperfect information thus seem to enhance the frequency of policy changes. Furthermore, the positive coefficient for the variable *diffusion* in model II including the interaction term indicates the following: under perfect information (i.e. *imperfect information* = 0), this frequency is greater under policy diffusion than under the individual learning scenario. Furthermore, estimates for the interaction term between both dummy variables indicate that the effect of introducing imperfect information on the frequency of policy change is even greater when states learn from others (i.e. when *diffusion* = 1) than its effect when they learn from their own experiences (i.e. when *diffusion* = 0).

In the light of these results, policy volatility can be seen as a feature that is predominantly associated with interdependent decision-making. Furthermore, policy volatility is not the result of a process that sets in after the complete spread of the more effective policy (at time step 6 in Figure 1, square III illustrating diffusion under perfect information). Instead, the process leading up to the complete spread of the more effective policy is already characterized by volatility. After a complete spread of the more effective policy, volatility dissolves because its underlying cause dissolves. The susceptibility to policy experiments made elsewhere becomes irrelevant when policy experiments do not occur any longer.[6]

5.2. Spatial patterns

While the patterns of policy adoption under diffusion and individual learning differ visibly when analyzed in terms of aggregated temporal patterns (see Figure 1) and in terms of policy volatility (see Figure 2), this difference can become blurred when policy adoptions are analyzed in terms of their spatial patterns. Figure 3 illustrates the spatial patterns resulting from the simualtion after 10 time steps.[7]

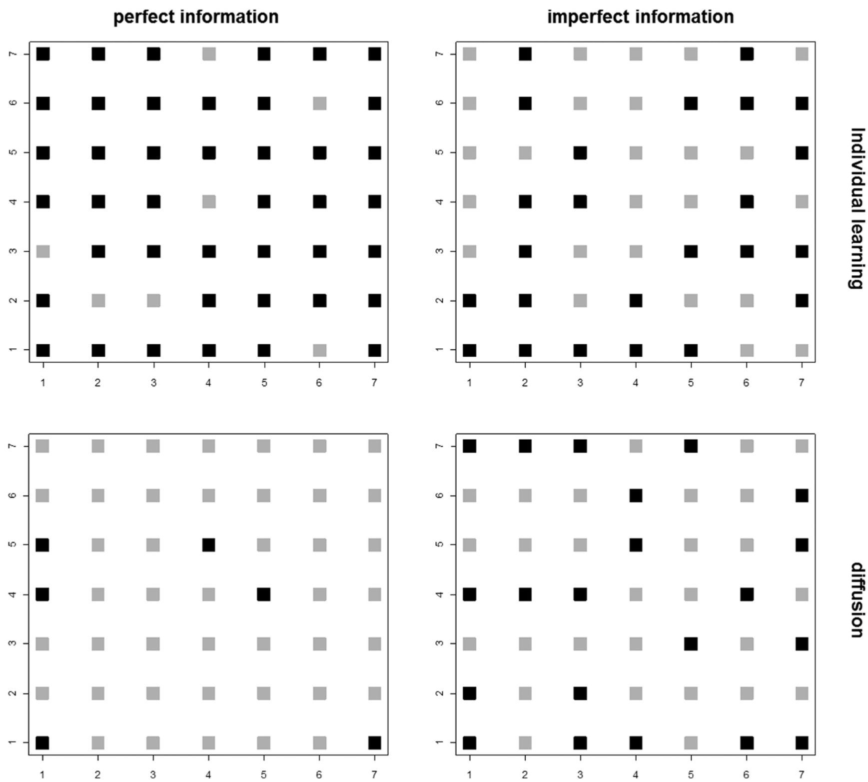

Figure 3. Spatial patterns.
Note: black squares indicate states holding the less effective policy after 10 runs; grey squares indicate states holding the more effective policy after 10 runs. The Figure represents the results of the final of 100 runs of each simulation.

Under perfect information, the policy maps resulting from the individual learning scenario and the diffusion scenario differ visibly. Under individual learning, the few states with the more effective policy (grey boxes) appear to be spread out randomly all over the map. In contrast, under diffusion the map seems to consist of one big cluster of states holding on to the more effective policy with just a few remaining isolated islands of states holding on to the less effective policy: under perfect information, the presence of policy clusters is thus clearly visible in the diffusion scenario and clearly absent in the individual learning scenario.

This picture changes substantially under imperfect information. With imperfect information, the shares of states holding on to the more effective policy under policy diffusion and individual learning move closer together (see Figure 1). This results in a situation in which the two sceanrios (diffusion and individual learning) are difficult to tell apart, as illustrated by Figure 3. While states with the more effective policy still seem to form spatial clusters in the diffusion scenario, these clusters also seem to be present in the individual learning scenario. Due to the wide

Table 3. Logistic regression with spatial lag.

	Model 1 Individual learning no bias	Model 2 Diffusion no bias	Model 3 Individual learning with bias	Model 4 Diffusion with bias
spatial lag	1.63	5.61***	2.22***	2.81***
	(1.75)	(0.84)	(0.74)	(0.46)
constant	−1.89***	−2.98***	−1.60***	−1.41***
	(0.47)	(0.43)	(0.28)	(0.24)
Pseudo R^2	0.008	0.45	0.04	0.10
N	490	490	490	490

Remarks: Unstandardized logistic regression coefficients. Robust standard errors for 49 state clusters in brackets. Data are generated by the simulation. *** reflects significance at 1% level.

spread of the more effective policy within a relatively small and closed network of states, the formation of spatial clusters is almost inevitable.

This graphical impression is underlined when conducting a simple, spatial lag regression that models the adoption of the more effective policy as a function of the adoption of this policy in a state's eight neighboring states (Table 3).

The data used for the analysis are generated by one run of the simulation over 10 time steps (490 observations = 49 states for 10 time steps). Again, this run can be interpreted to represent a randomly selected run from the population of all possible simulation runs due to the probabilistic elements that are built into the simulation. Since the dependent variable in this case is the policy currently held by each respective state, it only takes the values 0 (for the less effective policy) and 1 (for the more effective policy). Consequently, a logistic regression model is specified. It is able to assess whether the *spatial lag* variable—representing the share of a state's neighboring states holding on to the more effective policy—is independent of the probability that this state also features the more effective policy.

Under perfect information, this null hypothesis is rejected for the diffusion scenario (model 2) but not for the individual learner scenario (model 1). Under these conditions, one is thus able to correctly differentiate between diffusion and individual learning from statistical results based on a spatial lag variable. Since the true mechanism underlying policy adoption in these scenarios is known, it is evident that conclusions from these results yield true positive (model 2) and true negative (model 1) results.

Under imperfect information, regression results only allow for true positive conclusions about the nature of policy adoption when there is indeed diffusion (model). For the independent learner scenario, for which we know that policy decisions are independent, the analysis proposes to reject this null hypothesis of independent decision-making. Standard errors for the spatial lag coefficient are sufficiently small to impose this conclusion based on conventionally accepted levels of statistical signifance. While this leads to the conclusion that policy adoption is driven by policy diffusion, we know that this is a false positive conclusion. We know that the true underlying mechanism is one of independent individual learning.[8]

5.3. Propositions

Two propositions follow from the simulation results. First, the simulation suggests that the larger the share of states having adopted the more effective policy, the higher is the probability that any adoption pattern will look like a spatially clustered adoption pattern. When 80% of states in one's sample adopt a certain policy within a relatively short period of time, observing spatial clusters becomes almost inevitable—whether there is diffusion at play or not. This is the simulation's first proposition. The simulation is able to demonstrate the commonly expressed point of criticism concerning analyses relying on the identification of policy clusters to draw conclusions about the relevance of diffusion and interdependent policy decisions: these analyses can lead to false positive results (Volden et al., 2008).

> Proposition 1: Analyses relying on the identification of policy clusters to draw conclusions about policy diffusion are susceptible to false positive results when a certain policy innovation is adopted by a high share of states within a relatively short period of time.

Second, simulation results indicate that the concept of policy volatility can be helpful to differentiate between policy diffusion and independent decision-making. Of course, switching policies several times under political stability is generally rare. Yet, if it occurs, it is a strong indicator for interdependent policy-making. At first glance, this result might seem counter-intuitive. One might think that states that learn from the experience of their peers acquire more information, i.e. are able to work with more data points, and thus are quick to identify the superior policy and stick to it. The simulation suggests that this is only part of the story, however. In the system consisting of social learners the superior policy is indeed adopted by more states quicker. Yet, these states seem less capable of hanging on to this superior policy for long. This is because the exposure to the experience of others makes states reactive to policy experiments in other states. Change is more likely when states become aware of what seems to be a better policy solution for others. In consequence, states that learn from their peers are more regularly exposed to seemingly superior policy options and are therefore more likely to change policies. This volatility of policy decisions is caused not by the characteristics of individual states, but by their interaction. It is the interdependence of the state system which drives this policy volatility. If one state changes its status quo policy, it does so because it perceives this policy to have superior effectiveness. Yet, if this state perceives superior effectiveness from the alternative policy, then so will its neighboring states. This makes them more likely to switch to this policy alternative as a result of one of their neighbors switching which introduces positive feedback loops to the state system. That positive feedback loops induce instability is a well-known feature of complex systems (Miller and Page, 2007).

Furthermore, if a government introduces a policy innovation and subsequently decides to abolish the innovation and move back to the previous status quo, it might still be able to justify such a decision with evidence from national experiences regarding disappointing effectiveness of the policy innovation. A re-introduction of the policy innovation by the same government depends, however, on the existence of strong evidence about the effectiveness of the policy innovation found elsewhere.

In other words, only if the policy innovation is seen to work everywhere else will a policy innovation that was once abolished even have a chance of being re-introduced. States only looking at their own experience are by definition blind to such evidence, however.

Policy volatility is further enhanced by imperfect information as the overestimation and underestimation of policy effectiveness increases the willingness of individual states to engage in policy experiments. This is because critical thresholds for policy change are reached more often when policy effectiveness is overestimated and underestimated. Due to the interdependence of policy-making, this enhanced willingness to change the policy status quo translates into an enhanced level of policy volatility if information about policy effectiveness is imperfect. The instability of policy decisions, i.e. their volatility, is—according to the simulation—the result of interdependence of policy decisions. The introduction of misperceptions of policy effectiveness further destabilizes these decisions as it brings more activity (policy changes) to the state system.

> Proposition 2: As states are observed to reverse previous policy decisions more than once under political stability, the probability that this policy volatility is due to policy diffusion resulting from interdependent policy decisions by social learners increases substantially.

6. Discussion and conclusion

This article uses agent-based computerized simulation to model the dynamics of policy adoption. It does so in order to generate theoretical propositions about policy diffusion by indicating unique features of policy diffusion that can be exploited in empirical research designs. To do so, the article presents policy adoption patterns under four different scenarios that differ with respect to (a) states' ability to learn from the experience of other states and (b) the kind of information states receive about policy effectiveness (perfect vs imperfect). Controlled variation of these factors in the simulation allows for the creation of a computerized laboratory and the simulation of experimental conditions. Decisions about whether or not to change the current policy are simulated under political stability and under the consideration of political costs induced by the need to correct former policy decisions.

Based on the simulation results, the article formulates two propositions. First, relying on the identification of policy clusters can lead to false positive conclusions about the relevance of diffusion when a large share of one's sample adopts a certain policy within a relatively short period of time. Second, the concept of policy volatility can help to minimize the risk of such false positive conclusions. Both propositions yield empirically testable hypotheses.

These results can enrich diffusion research in at least two ways. First, the simulation suggests that policy volatility can be used as a complementary criterion for studies interested in the question of whether or not policies diffuse. If there is a non-zero number of states changing their policies more than once under political stability, then this is a strong indication for policy diffusion. For the sake of simplicity, the simulation has focused on policy volatility as vacillation between two policy options. Empirical studies could easily consider policy volatility in terms of states switching back and forth between several different policy options. Of course, the absence of policy volatility should not be mistaken for an indicator of the irrelevance of policy diffusion (many states in the diffusion scenario only change policies once). Policy decisions on smoking bans and gambling policy have been shown to depend strongly on decisions made elsewhere (e.g. Berry and Berry, 1990; Jensen, 2003; Shipan and Volden, 2008) without being volatile. In other words, solely relying on policy volatility will often lead to false negative conclusions about policy diffusion (type-II errors, to use the statistical analogy again). Yet, if states change policies frequently under political stability, then—according to the simulation results—these states are likely to attribute a lot of weight to the experiences with policy effectiveness made elsewhere. Therefore, despite this risk to produce false negative results, the tendency to avoid false positive results makes the criterion of policy volatility a valuable indicator able to complement evidence about policy clusters.

Second, policy volatility can be a helpful tool informing case selection of diffusion studies using comparative case studies and causal process-tracing to investigate underlying causes of interdependent policy-making. Researchers intending to use process-tracing to analyze the underlying cause of policy diffusion, i.e. answer the question of why certain states let external information strongly influence domestic policy decisions, need to know where to allocate their resources: in which states can the relevant outcome, i.e. the interdependence of a specific policy decision, most likely be traced back to its roots? The criterion of policy volatility can inform this decision. This can be illustrated with an example from nuclear energy policy. In order to see what causes a government to make its domestic nuclear energy policy strongly dependent on policy decisions and experiences made elsewhere, focusing on the volatile policy decisions of the German government promises to be a fruitful strategy. Specifically, the simulation would suggest that the flip-flop on nuclear energy of the German Chancellor, Angela Merkel, represented a most-likely case of interdependent decision-making. After her government coalition had ended the exit from nuclear energy—by passing legislation extending contracts with German suppliers of nuclear energy—the exit was quickly re-introduced by Merkel's government with the imposition of a 'moratorium' for German nuclear power plants after a disastrous earthquake triggered a nuclear catastrophe in Japan. Merkel claimed that this policy change was necessary after the events in Japan revealed new information about the potential dangers of nuclear power production (German Federal Government, 2011). This indicates that the volatility of the German policy decision was indeed strongly influenced by social learning and underlines how choosing diffusion cases on the criterion of policy volatility can be a fruitful strategy for researchers looking for most-likely diffusion cases because they want to investigate the process that links available information on policy effectiveness to the actual (interdependent) policy decisions.

In sum, this article hopes to inspire more research using agent-based computer simulation in the context of policy diffusion. It hopes to have demonstrated that simulations can be a valuable source of theoretical innovation in this area.

Appendix A

Any simulation is sensitive to changes in parameter values. Table 4 shows how sensitive the simulation presented in this article is to changes in parameter values. Simulation results can only be considered to be robust if they are realized for a range of parameter values. Table 4 indicates for which parameter values the simulation results are robust.

Table 4. Robustness of the simulation.

	Parameter name	Parameter value		All else equal, results are robust for:
		baseline run	experimental run	
diffusion parameter	d	2		$1 \leq d \leq 3$
number of states	n.states	49		≥ 25
mean performance of less effective policy	e0	1		≤ 1
mean performance of more effective policy	e1	2		≥ 1.1
power in logistic function determining the probability of policy change	g	-5	-5	$-2 \leq g \leq -9$
uncertainty / standard deviation of policy performance	sde0, sde1	0	0.25	≥ 0.25
factor of threshold increase after each policy change	incr	2	2	≤ 2
percentage of states with the less effective policy as initial policy	initial.perc	100%		$\geq 25\%$

Appendix B

In the main part of the article, states were observed for 10 time steps in which they could decide for either policy. The Figure 4 displays results for a longer time period. Here, states chose 100 times. Results for this higher number of time steps remain essentially the same. The incremental increase in the individual learning scenario results in a wider final spread of the more effective policy. While the S-curve pattern typical for diffusion processes is less evident than in Figure 1, due to the different scale of the x-axis, it is still present. Furthermore, policy volatility is still a unique feature of the diffusion scenarios. In fact, it becomes clear that, in the diffusion scenarios, relevant changes occur relatively early in the simulation. After these early developments equilibria are achieved. Interestingly, under imperfect information, the sub-optimal equilibrium is particularly evident as the more effective policy never manages to spread fully throughout the state system in the social learner scenario (square IV). Due to the high political costs of changing policies several times, some states remain stuck with the less effective policy.

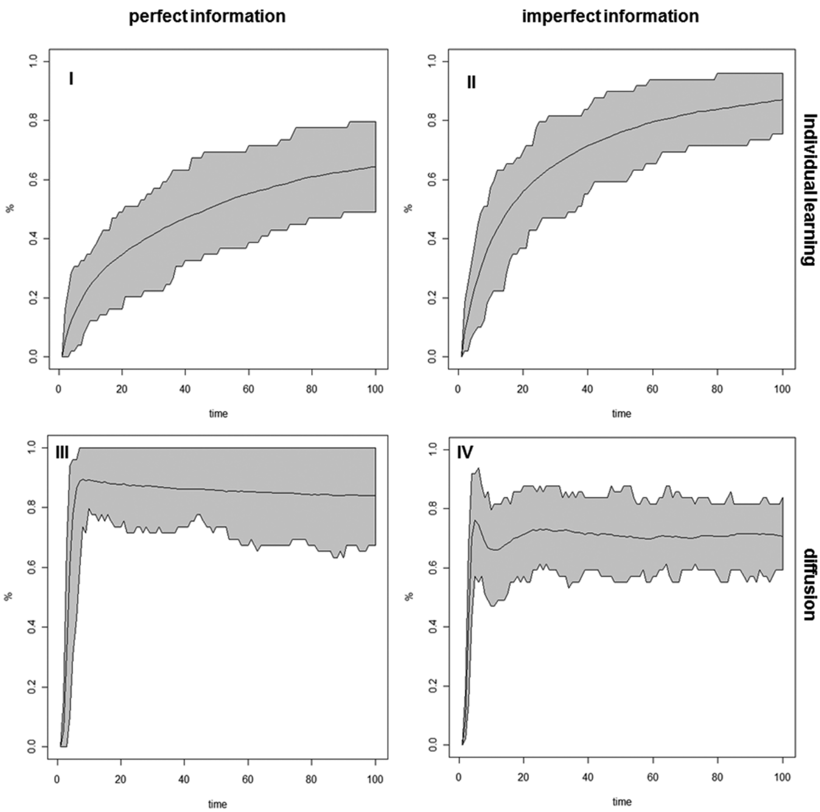

Figure 4. Spread of the more effective policy in a per cent of all states.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by a scholarship as per the State Law on Graduate Funding by the State of Baden-Wuerttemberg, Germany.

Notes

1. The simulation code will be available for replication of results from the author's personal website at: www.christianadam.org.
2. Kultusministerkonferenz.
3. This re-introduction was, however, less far-reaching in scope (Eisenberg, 2006).
4. Modeling social decisions as stochastic functions of logistic types reflects a probabilistic threshold model of decision-making and is common for agent-based models (e.g. Cederman, 2003). Change is more likely to occur when a certain threshold in the utility

of change is reached. The more the utility of change exceeds the threshold, the more likely change occurs.

5. The results remain robust, however, when including more time-steps per simulation run. In fact, the simulation's main message (as formulated in the discussion) becomes more accentuated for increasing numbers of time-steps per run (see Appendix B).
6. These results are evident when running the simulation with just six time steps and when running the simulation with the more effective policy as the initial policy. These results can be replicated with the replication code from the author's website.
7. While the simulation models connectivity between states through geography, the resulting map could just as well be interpreted as connectivity in terms of economic or cultural connectivity between states.
8. This result should not be mistaken for proof that spatial lag regression will necessarily produce false positive results. That is not the point. Other specifications of the statistical model might very well be able to distinguish correctly between the scenarios. Instead, the result underlines that spatial lag regression can lead to false positive conclusions.

References

Axelrod R (1997) *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton, NJ: Princeton University Press.

Axelrod R and Cohen MD (2000) *Harnessing Complexity*. New York, NY: The Free Press.

Bennett CJ and Howlett M (1992) The lessons of learning: reconciling theories of policy learning and policy change. *Policy Sciences* 25(3): 275–294.

Berry FS and Berry WD (1990) State lottery adoptions as policy innovations: an event history analysis. *American Political Science Review* 84(2): 395–415.

Boehmke FJ and Witmer R (2004) Disentangling diffusion: the effects of social learning and economic competition on state policy innovation and expansion. *Political Research Quarterly* 57(1): 39–51.

Braun D and Gilardi F (2006) Taking 'Galton's Problem' seriously: towards a theory of policy diffusion. *Journal of Theoretical Politics* 18(3): 298–322.

Cederman L-E (2003) Modeling the size of wars: from billiard balls to sandpiles. *American Political Science Review* 97(1): 135–150.

Clough E (2007) Strategic voting under conditions of uncertainty: a re-evaluation of Duverger's Law. *British Journal of Political Science* 37(2): 313–332.

Cohen J (1977) *Statistical Power Analysis for the Behavioural Sciences,*2nd edn. New Jersey: Lawrence Erlbaum Associates.

Die Zeit (2013) Landtag beschließt Ende der Studiengebühren in Bayern, Die Zeit. Available at: www.zeit.de/studium/hochschule/2013-04/studiengebuehren-bayern-abschaffung (accessed 24 July 2013).

Eisenberg P (ed.) (2006) *Niemand hat das letzte Wort: Sprache - Schrift - Orthographie, Valerio*. Göttingen, Germany: Wallenstein.

Etheredge L and Short J (1981) Thinking about government learning. *Journal of Management Studies* 20(1): 41–58.

German Federal Government (2011) German government imposes three-month moratorium on extension of nuclear power plant operations. Available at: www.bundesregierung.de/Content/EN/Artikel/_2011/03/2011-03-14-moratorium-kernkraft-deutschland_en.html (accessed 29 April 2013).

Gilardi F (2010) Who learns from what in policy diffusion processes? *American Journal of Political Science* 54(3): 650–666.

Gilardi F (2013) Transnational diffusion: norms, ideas, and policies. In: Carlsnaes W, Risse T and Simmons B (eds) *Handbook of International Relations*. Thousand Oaks: SAGE Publications.

Gilardi F, Flüglister K and Luyet S (2009) Learning from others: the diffusion of hospital financing reforms in OECD countries. *Comparative Political Studies* 42(4): 549–573.

Hall PA (1993) Policy paradigms, social learning and the state: the case of economic policymaking in Britain. *Comparative Politics* 25(3): 275–296.

Heclo H (1974) *Modern Social Politics in Britain and in Sweden*. New Haven, CT: Yale University Press.

Jensen JL (2003) Policy diffusion through institutional legitimation: state lotteries. *Journal of Public Administration Research and Theory* 13(4): 521–541.

Kollman K, Miller JH and Page SE (1992) Adaptive parties and spatial elections. *American Political Science Review* 86(4): 929–937.

Lustick IS, Miodownik D and Eidelson RJ (2004) Secessionism in multicultural states: does sharing power prevent or encourage it? *American Political Science Review* 98(2): 209–230.

Luyet S (2009) An agent-based approach of the international diffusion of welfare state reforms. Presented at the Annual Swiss Political Science Association Congress, St Gall, Switzerland.

Meseguer YC (2009) *Learning, Policy making, and Market Reform*. Cambridge: Cambridge University Press.

Miller JH and Page SE (2007) *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton, NJ: Princeton University Press.

Pigeon M (2012) *Remunicipalisation Putting Water Back into Public Hands*. Amsterdam, the Netherlands: Transnational Institute.

Rapaport O, Levi-Faur D and Miodownik D (2009) The puzzle of the diffusion of central-bank independence reforms: insights from an agent-based simulation. *Policy Studies Journal* 37(4): 695–716.

Rose R (1991) What is lesson-drawing? *Journal of Public Policy* 11(3): 3–30.

Sabatier PA (1988) An advocacy coalition framework of policy change and the role of policy-oriented learning therein. *Policy Sciences* 21(2/3): 129–168.

Shipan CR and Volden C (2008) The mechanisms of policy diffusion. *American Journal of Political Science* 52(4): 840–857.

Simmons B and Elkins Z (2004) The globalization of liberalization: policy diffusion in the international political economy. *American Political Science Review* 98(1): 171–189.

Strang D and Macy MW (2001) In search of excellence: fads, success stories, and adaptive emulation. *American Journal of Sociology* 107(1): 147–182.

Vicsek T (2002) Complexity: the bigger picture. *Science* 418: 131.

Volden C, Ting MM and Carpenter DP (2008) A formal model of learning and policy diffusion. *American Political Science Review* 102(3): 319–332.

Vriend NJ (2000) An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control* 24(1): 1–19.

Ward H and John P (2013) Competitive learning in yardstick competition: testing models of policy diffusion with performance data. *Political Science Research and Methods* 1(01): 3–25.

Winkelmann R (2008) *Econometric Analysis of Count Data*. Berlin, Germany: Springer.