UvA-DARE (Digital Academic Repository)

Learning dynamics and support for economic reforms: why good news can be bad

van Wijnbergen, S.J.G.; Willems, T.

DOI

[10.1093/wber/lhu005](https://doi.org/10.1093/wber/lhu005)

Publication date

2016

Document Version

Author accepted manuscript

Published in

The World Bank Economic Review

[Link to publication](#)

Citation for published version (APA):

van Wijnbergen, S. J. G., & Willems, T. (2016). Learning dynamics and support for economic reforms: why good news can be bad. *The World Bank Economic Review*, 30(1), 1-23.
<https://doi.org/10.1093/wber/lhu005>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Learning Dynamics and Support for Economic Reforms: Why Good News Can Be Bad

Sweder J. G. van Wijnbergen and Tim Willems*

Support for economic reforms has often shown puzzling dynamics: many reforms that began successfully lost public support. We show that learning dynamics can rationalize this paradox because the process of revealing reform outcomes is an example of sampling *without* replacement. We show that this concept challenges the conventional wisdom that one should begin by revealing reform winners. It may also lead to situations in which reforms that enjoy *both ex ante and ex post* majority support will still not come to completion. We use our framework to explain why gradual reforms worked well in China (where successes in Special Economic Zones facilitated further reform), whereas this was much less the case for Latin American and Central and Eastern European countries. JEL classification: D72, D83, P21

Why have gradual economic reforms worked out well for China, whereas this is much less the case for most Latin American and Central and Eastern European countries? How is it possible that so many of the reforms that began successfully while enjoying majority support subsequently lost this support, although there are also examples of reforms that did not begin well but nevertheless managed to maintain momentum among voters?

The most dramatic example of a reformist government that lost majority support in spite of strong economic performance is Slovakia in 2006. At that time, the *Wall Street Journal Europe* wrote,¹

“Imagine you’re the leader of a country where economic growth is running at 6.3%, your government has been praised by the World Bank as the best market reformer in the world [and] unemployment has fallen to a record low of 10.6% from around 20% in just four years. [. . .]”

* Sweder van Wijnbergen is a professor at the Department of Economics, University of Amsterdam, The Netherlands. He is also affiliated with the Tinbergen Institute. His e-mail address is s.j.g.vanwijnbergen@uva.nl. Tim Willems (corresponding author) is a research fellow at Nuffield College and the Department of Economics, University of Oxford, UK. He is also a member of the Centre for Macroeconomics. His e-mail address is tim.willems@economics.ox.ac.uk. The authors thank the editor (Andrew Foster), two anonymous referees, Philippe Aghion, Björn Brügemann, Tom Cunningham, Allan Drazen, Michal Horvath, Ruixue Jia, Matija Lozej, Torsten Persson, Dani Rodrik, Gérard Roland, and audiences at the EBRD and the 2013 SITE Conference in Stockholm for useful comments and discussions.

1. Robin Shepherd, “The Dzurinda Revolution”, *Wall Street Journal Europe*, June 12, 2006.

THE WORLD BANK ECONOMIC REVIEW, pp. 1–23

doi:10.1093/wber/lhu005

© The Author 2014. Published by Oxford University Press on behalf of the International Bank for Reconstruction and Development / THE WORLD BANK. All rights reserved.

For permissions, please e-mail: journals.permissions@oup.com.

With this record in mind, now consider that you face parliamentary elections this Saturday at which, unless the opinion polls change dramatically, you risk annihilation by a leftist opposition party with no experience of government and a policy agenda filled with populist rhetoric.

Welcome to the world of Mikuláš Dzurinda, prime minister of Slovakia, who for the past eight years has led what can reasonably claim to have been the most successful neo-liberal government of the 21st century so far.”

Despite his impressive reform successes, Dzurinda lost the 2006 elections to Robert Fico of the SMER party (a breakaway party from the successor to the original Communist Party of Slovakia), who reversed many of Dzurinda’s reforms.

With important reforms currently being implemented in many African and Southern European states, it is important to understand why such puzzling reversals can occur. In this paper, we focus on the interaction between learning from reform outcomes and the dynamics of public support for gradual economic reforms.²

We believe that learning processes play a key role in determining support for reforms. So far, however, the literature has remained relatively silent on this issue. Although there are many informal discussions of learning from reform outcomes, formal treatments are scarce.³

Even if everyone gains from efficient reforms in the long run, there will almost inevitably be losers during the transitional phase (for example, certain generations or professions). The model we construct captures the fact that reforms typically generate reform winners and reform losers, but (as emphasized in the seminal paper by [Fernandez and Rodrik \(1991\)](#)) these winners and losers cannot always be identified in advance. That is, there is individual uncertainty, which causes the reform to have uncertain distributional consequences. As the reform progresses over time, voters update their beliefs about whether they will end up in the winners’ group or in the group of losers. Because the full version of our model not only features individual uncertainty but also allows for aggregate uncertainty (which implies that voters are unsure about the exact share of the population that will benefit from reform), this paper can be seen as augmenting the Fernandez-Rodrik setup with aggregate uncertainty and learning dynamics.

The fact that existing reform measures affect the distribution from which future sampling will occur plays a key role in our analysis. Specifically, the process of revealing reform outcomes is an example of sampling *without* replacement. This implies that the revelation of reform winners deteriorates the quality

2. With the possible exception of price decontrol, all reforms are gradual (as opposed to “big bang”), if only because of implementation delays. As noted, for example, by [Gupta, Ham, and Svejnar \(2008\)](#), even reforms that were supposed to be “big bang” (such as the Balcerowicz reforms in Poland) were not completed instantaneously, thereby giving voters an opportunity to update their beliefs about the effects of the reform. In this sense, all reforms are gradual, but some reforms are “more gradual” than others.

3. Some exceptions are [van Wijnbergen \(1992\)](#), where voters learn about the effects of price reform), [Dewatripont and Roland \(1995\)](#), where the public uses early reform outcomes to learn about the expected outcome of later ones), [Veldkamp \(2009\)](#), where laid-off workers learn about their re-employment chances), [Strulovici \(2010\)](#), analyzing determinants of collective experimentation), and [Morrow and Carter \(2013\)](#), studying the impact of learning dynamics on the support for redistributive policies).

of the remaining pool, thereby making unreformed agents less eager to continue the reform process. We derive a condition under which these dynamics are so strong that they lead to the counterintuitive situation in which reform *successes* make the median voter begin *opposing* a reform he used to support. In these circumstances, the reforming government suddenly loses majority support although (or because) the reform is progressing in such a successful way. Consequently, even reforms that enjoy *both ex ante and ex post* majority support may not come to completion.

We emphasize that this phenomenon results from rational economic thinking and that it arises as soon as a reform is believed to generate losers whose identity is *ex ante* unknown - a feature that we see as being in accordance with many economic reforms in reality (see also [Fernandez and Rodrik 1991](#)). In addition, sampling without replacement continues to play a role when one adds aggregate uncertainty. In that setup, the revelation of winners may also lead to an upward revision of the expected aggregate number of reform beneficiaries (enhancing support for reform). However, we show that it is possible for this strategy to sow the seeds for its own destruction.

Although this learning mechanism applies to many reform types (such as land reform, the gradual abolition of subsidies/price controls, or the reduction in trade barriers), we often link it to privatization. Privatization is a good example of a reform in which learning dynamics may be important, and the choice between starting with “good” or “bad” companies arises consistently. Because the government is the incumbent owner of the firms that are to be privatized, it often has inside information on the profitability of these firms and on future policies that may benefit or harm them. This situation brings up the sequencing question for the government and the issue of learning for other agents (see, e.g., [Roland 2000](#), chapter 2).

On a more general level, this paper develops a theory of agents who are learning from realizations that are sampled without replacement. As a result, the sampling process itself affects the distribution from which future sampling will take place. From the so-called “Monty Hall problem”, we know that this situation can give rise to counterintuitive dynamics that are not easy to understand without a formal model. Building such a model may therefore make this paper of independent interest because it may have implications for many other economic problems (a topical example is the process of revealing the identity of good and bad banks in a financial crisis).

Our results question the political feasibility of the so-called “sectoral gradualist” way of privatization. This strategy has been advocated, for example, by [Kornai \(1990\)](#) through his plea for the “case-by-case” approach, and it has been applied to many countries in Latin America and Central and Eastern Europe as well as to the UK during its liberalization phase. This strategy implies that one sector (or firm) is reformed after the other (cf. [Berg and Blanchard \(1994, 53, 63\)](#)). However, as we argue below, following such a gradual, sequential approach triggers the “sampling without replacement” effect presented previously. This may explain why practitioners have experienced political difficulties with the case-by-case strategy. [Lipton and Sachs \(1990, 298\)](#) note that “in almost all

countries where privatizations have been attempted, there have been major political obstacles to the case-by-case approach". Boycko, Shleifer, and Vishny (1993, 148) state that reforms that proceed at a rather slow pace are likely to reach a deadlock. As we argue later in this paper, "spatial gradualism" (reforming one *region* after another, as China did by installing Special Economic Zones) can avoid the "sampling without replacement" problem. Thus, the mechanisms explored in this paper may help explain why gradual reform strategies have been more successful in China than in Latin America and Central and Eastern Europe.

The outline of this paper is as follows. We first describe various examples of reforms for which support dynamics have been counterintuitive. Next, we construct a learning model that provides an explanation for these puzzling dynamics. We then consider the question of why gradualism worked quite well for China, although this was not the case for most Latin American and Central and Eastern European countries. Finally, we conclude.

SUPPORT DYNAMICS FOR ECONOMIC REFORMS: A SHORT HISTORY

In addition to the case of Slovakia discussed in the introduction, there are many examples of economic reforms that lost support despite their initial success (and vice versa). Stokes (2001) provides a thorough analysis of support dynamics related to various reforms. In that volume, several authors examine the public's reactions to reforms in Spain, East Germany, Poland, Mexico, Peru, and Argentina. In her summary of the study, Stokes (2001, 25) notes that "[their] most startling result is that in every country people sometimes reacted to economic deterioration by supporting the government and its economic program more strongly. Conversely, they sometimes reacted to economic improvement with pessimism and opposition". Similar findings are reported by the econometric studies of Remmer (1991) and Tucker (2000), who analyzed data from 12 Latin American and five post-communist countries, respectively, and reported a negative causal effect of economic performance on support for incumbents.

Stokes (2001) provides various specific examples of these counterintuitive dynamics. For example, in all three Latin American countries studied (Mexico, Peru, and Argentina), economic expansion (measured by either wage or GDP growth) was followed by pessimism about the future and opposition to the reform program. Similarly, increased real wages in Poland did not generate support for the reforms but created agnosticism instead.

With respect to the latter case, Rodrik (1995, 404) expresses surprise as well. When discussing the return to power of the former Polish communist party in 1993, he writes, "Why this should be so is not so easy to understand. [. . .] By most standards, Poland must be judged a success case".⁴

4. Poland had a high unemployment rate at the time, but as Rodrik (1995, 405) notes, it is not clear whether that was to blame for the deadlock. The unemployed group is too small to be decisive in national elections, and it is not straightforward that their interests are best served by policies that slow reforms.

Regarding general experiences in Central and Eastern Europe, [Fidrmuc \(2000, 1491\)](#) notes that “the collapse of communism occurred amidst overwhelming popular support for fundamental economic and political reforms. However, only a few years later the pendulum swung back and the reformers were voted out”. For example, Slovenia faced great difficulties in its reform process though it already had quite a few positive experiences with market forces from the past ([Pleskovic and Sachs 1994](#)). Although the 1968 Hungarian reforms began successfully, they encountered difficulties in the mid-1970s when the country underwent periods of recentralization ([Qian and Xu 1993](#)). Similarly, after the second wave of reforms following the demise of communism, the reformist Hungarian government lost the 1994 elections, and the former communist party returned to power (as in Poland and later in Slovakia), a pattern that led [Kornai \(2000\)](#) to conclude that the gradual reform strategy may not be feasible from a political point of view.

Latin America offers examples of countries that have had similar experiences. Puzzled by this situation, [Tommasi and Velasco \(1996\)](#) ask, “Why did Venezuelans riot, twice attempt to overthrow and eventually impeach a president (Carlos Andrés Pérez) who in 1990–2 brought them an average growth rate of 7.8% (the highest in Latin America), while Peruvians massively re-elected Alberto Fujimori, under whose stewardship consumption dropped by 15.3% in 1990?” Similarly, [Iglesias \(1994, 497–8\)](#) notes, “In my country (Uruguay), which is growing by 11.5 percent, where unemployment and inflation are down, and where reserves are up, the popularity rating of the president is 12 percent. That’s why the administration lost its bid to privatize the telephone company”.⁵

A similar story holds with respect to India: notwithstanding the successes of the Indian liberalization policies adopted in the 1990s, India is currently struggling to implement new reforms and has turned into “a place that has fallen out of love with reform” (as stated in *The Economist*, 24 March 2012, 14). More generally, [Sachs and Warner \(1995\)](#) have documented how many countries slowed down (or even reversed) their liberalization policies in the 1960s and 1970s, though the economic performance under the more liberal regime was impressive.

These examples suggest that a successful beginning of a reform is by no means a sufficient condition for the reform to maintain majority support along the way. This observation is at odds with the conventional wisdom that a favorable start facilitates continuation.

5. In 1994, Luis Alberto Lacalle (of the Partido Nacional) was president of Uruguay. After taking office in 1990, he began significant economic reforms (in the sphere of both taxation and liberalization), but his initiatives later lost support (despite successes; cf. Iglesias’ quote). Subsequently, he lost the 1995 elections and was replaced by Julio María Sanguinetti of the rival Partido Colorado, who reversed many of Lacalle’s reforms.

At the other end of the spectrum are the gradual economic reforms in China. There, the government established Special Economic Zones in 1980, after which the economies of those regions began booming. In contrast to the experiences of Central and Eastern Europe and Latin America, the initial successes of these Chinese reforms led to increased support for further reforms in China (Litwack and Qian 1998; Qian, Roland, and Xu 1999, 2006).

In the next section, we develop a model that is able to rationalize the confusing support dynamics in Latin America and Central and Eastern Europe while simultaneously shedding light on the question of why the initial Chinese reform successes did not invoke such a paradoxical public response.

SUPPORT DYNAMICS FOR ECONOMIC REFORMS: A MODEL

In this section, we describe our model. Although we frame the model in the context of economic reforms, it can also be seen as a more general model of agents who are learning from realizations that are sampled without replacement.

Our model is dynamic and contains uncertainty at both the aggregate level and the individual level. The former makes agents in the model uncertain about the total number of reform winners, whereas the presence of the latter implies that the reform will have uncertain distributional consequences. Implementing reform takes time. Consequently, reforms are completed gradually (as in reality; recall footnote 2), and agents have the opportunity to update beliefs about their chances of benefiting from the reform as it progresses over time. This belief updating process lies at the heart of our paper. The model presented by Fernandez and Rodrik (1991) emerges as a special case of our framework without learning dynamics and without aggregate uncertainty.

To build intuition for the mechanisms at play and to set the stage for our full model, we first consider a setup in which there is only individual uncertainty but no uncertainty in the aggregate. That is, in the first subsection, individuals know *what fraction* of the population will gain from the reform, but ex ante they do not yet know *who* these winners will be.

Subsequently, we add aggregate uncertainty. Then, individuals are also not sure about *what fraction* of the population will benefit from reform. The existence of such uncertainty has often been named as a reason to reform gradually by starting with the revelation of reform winners (cf. Roland 1994, 1164). However, as we show, it is still possible that a successful beginning to a reform will sow the seeds for its own destruction in that setup.

Without Aggregate Uncertainty

To build intuition, we first illustrate our point in a setup without aggregate uncertainty. Time is discrete, the horizon is infinite, and there is a large number of risk-neutral voters aligned uniformly between 0 and 1, indexed by i . Each voter i

can be thought of as representing individuals associated with a particular firm or sector. We assume that voters are rational and forward looking.⁶

Voters are faced with a reform proposal R_γ , which is to replace the status quo. This proposal is assumed to leave everyone with a net present value payoff of 0. Reform R_γ , in contrast, is known to benefit a fraction $\gamma > 1/2$ of the population with certainty (yielding them a net present value payoff of $S > 0$). The losing fraction $(1 - \gamma)$ is assumed to receive a symmetric negative payoff of $-S$ (where the symmetry simplifies the algebra, without loss of generality).⁷

This implies that there is no aggregate uncertainty and because $\gamma > 1/2$, the reform is efficiency enhancing (according to the Kaldor-Hicks criterion) and would always be welcomed by a majority *ex post*.

However, the electorate faces individual uncertainty. In response to this uncertainty, voters form (potentially heterogeneous) beliefs about the effects of the reform on their personal well-being. We allow for belief heterogeneity in a discrete way: there is a fraction α_t that believes (or already knows) *ex ante* that it will belong to the group of reform winners, whereas a fraction β_t believes/knows that it will be among the losers (with $0 < \alpha_t, \beta_t < 1/2$). The remaining fraction $(1 - \alpha_t - \beta_t)$ (which we assume to share a common prior) does not know at time t whether it will gain or lose from reform. The members of that fraction will base its decision upon the expected value of the reform for them. If we sort all individuals (indexed by i) such that the γ *ex post* winners of the reform are located on the left of the interval and the $(1 - \gamma)$ losers are on the right, we obtain the configuration shown in Figure 1.⁸

Voters with $i < \alpha$ know that they are among the reform winners, whereas voters with $i > 1 - \beta$ (where the “one minus” follows from the fact that β is measured from the right) know that they are among the losers. At the beginning of the

6. Whether voters are forward looking or backward looking is somewhat debated. Although the early papers on this issue report that voters are myopic and backward looking (see, e.g., [Kramer 1971](#)), more recent studies tend to find that rational forward-looking behavior dominates (cf. [MacKuen, Erikson, and Stimson 1992](#); [Fidrmuc 2000](#)). Introducing retrospective voting would produce goodwill for the reforming government after early reform successes, somewhat similar to what arises when there is aggregate uncertainty (see Section II). As in that setup, the pace at which the various forces operate will then determine which effect eventually dominates.

7. With asymmetric payoffs, such as when winners obtain G and losers receive $-L$ with $G > L$, reforms with $\gamma < 1/2$ could also be welfare enhancing in the Kaldor-Hicks sense (and vice versa when $L > G$). This generalization would lead to a different set of reforms that can be implemented through majority voting but do not affect the mechanics of the “sampling without replacement” effect central to this paper because the “sampling without replacement” effect refers to the availability of winning places relative to the losing ones, not to the exact magnitude of the associated gain or loss. To see this, note that replacing S with G (respectively L) in equations (3) and (4) below does not change any of our results.

8. We rule out partial reform (e.g., reform the winners and keep the losers under the state wing) as a desirable outcome. Clearly, a formal analysis of such an outcome would require the introduction of interactions between reformed and unreformed parts of the economy and of costs associated with the use of public funds (keeping loss-making, government-owned firms in operation is costly to society). Incorporation of this point would detract from the clarity of our core message; thus, we do not introduce these obfuscating factors here. See [Murphy, Shleifer, and Vishny \(1992\)](#) for a model that addresses the problems related to partial reforms. In addition, [Dewatripont and Roland \(1992\)](#) show that partial reforms are typically time inconsistent and therefore unsustainable.

FIGURE 1. Graphical Illustration of the Setup

reform, α and β can be equal to 0, but this does not necessarily have to be the case; it is perfectly possible that some agents already operate under the new regime before the reform has started (for example, as a remnant of uncompleted past reform attempts) or that their identity is obvious up front. Agents between α and $1 - \beta$ are uncertain about their identity and do not know whether they will be a reform winner or a reform loser. Because the identity of more and more individuals is revealed as the reform progresses over time, α and β become time varying and thus obtain a time index. In contrast, γ is a time-invariant structural parameter characterizing the reform (with aggregate uncertainty, the public's *estimate* of γ can become time varying, though γ itself is fixed, which is what we allow for in the next section).

The expected value of the reform for uncertain individuals (i.e., those with $i \in (\alpha_t, 1 - \beta_t)$) equals

$$E_t\{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\} = (\gamma - \alpha_t) \cdot S_t + (1 - \gamma - \beta_t) \cdot (-S_t). \quad (1)$$

Individuals in that group follow the decision rule

$$\delta_t = \begin{cases} 1 & \text{if } E_t\{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\} > 0 \\ 0 & \text{if } E_t\{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\} \leq 0 \end{cases}, \quad (2)$$

where δ_t is a support indicator that takes the value 1 if the uncertain group votes in favor of the reform and zero otherwise. Because α_t and β_t are both smaller than $1/2$, the decisive median voter is located in this uncertain group.⁹

The expected value of this uncertain group (expressed by (1)) can be negative for a wide range of parameter combinations, thereby making all $(1 - \alpha_t - \beta_t)$ uncertain individuals oppose the reform package ex ante. Because $\alpha_t < 1/2$, this implies that the reform does not enjoy majority support up front, though it would be welcomed by a majority ex post (because $\gamma > 1/2$).¹⁰

9. We follow Fernandez and Rodrik (1991) and many others in assuming that a reform is more likely to be adopted if there is a larger number of individuals in favor of it, but we use the language of majority voting for concreteness. The model can, however, also be interpreted as describing support dynamics for reforms in non-democratic countries. Decision rule (2) could then be interpreted in terms of joining an anti-reform protest or not (in that case, it might be realistic to include a “protest cost”).

10. Here, one should note that the Fernandez-Rodrik model assumes that it is not possible to compensate the losers ex post. As noted by Messner and Polborn (2004, 118), this assumption is “standard (and often even implicit) in the literature”. Given the well-known difficulties that governments face in committing to future policies, this assumption may not be unrealistic.

As in [Fernandez and Rodrik \(1991\)](#), the presence of individual uncertainty can thus prevent efficiency-enhancing reforms from being implemented. In particular, there are currently ex post winners blocking the reform ex ante because they do not know that they will be among the ex post winners.

Because individual uncertainty lies at the core of the problem, one may think that reducing individual uncertainty by revealing winners (i.e., increasing α_t to $\alpha_{t+1} = \alpha_t + \Delta\alpha_{t+1}$, bringing it closer to γ) would make a yes vote more likely. This turns out *not* to be true, opening an interesting perspective on voter dynamics. To see this, consider how the expected value for uncertain individuals changes with α :

$$\frac{\partial E_t\{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\}}{\partial \alpha_t} = -S < 0. \quad (3)$$

Therefore, a decrease in individual uncertainty brought about by the revelation of additional winners makes individuals who remain uncertain more *negative* about their chances of gaining from the reform. The reason is that in the absence of aggregate uncertainty (which is added in the next section), increasing α_t to α_{t+1} implies that there are $\Delta\alpha_{t+1}$ fewer gaining places left for those who remain uncertain (because the revelation of reform outcomes is an example of sampling *without* replacement). This makes these uncertain individuals more pessimistic about their chances of ending up as reform winners (because the revelation of winners deteriorates the quality of the remaining pool). When the median voter is located within this uncertain group, he also becomes more pessimistic.

Revealing losers, in contrast, increases the expected value of the reform for those who remain uncertain:

$$\frac{\partial E_t\{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\}}{\partial \beta_t} = S > 0. \quad (4)$$

At this stage, one should note that there is a wide range of values for α and β where changes in uncertainty will not change the outcome of the vote. If the vote is initially “no”, then increases in α will only make $E_t\{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\}$ more negative, and the median voter will continue to oppose the package.

There is an intriguing possibility if the median voter initially supports the reform package. To see this, hold β_t constant at $\bar{\beta}$ for a moment¹¹ and let us investigate what happens if the government tries to complete the reform gradually by increasing α_t (i.e., revealing winners). If the increase in α is small enough, δ will remain 1, and the median voter continues to vote “yes”, pushing the overall vote in favor. However, because of the effect captured by (3), one can define a

11. The same argument applies, *mutatis mutandis*, to changes in β_t if the median voter initially opposes reform. However, in that case, the reform cannot be started along democratic lines.

critical value for α (call it α^*) such that if α rises above α^* , the median voter swings around, causing a rejection of the package.¹²

The critical value α^* is thus the point at which the median voter begins to oppose a reform that he used to support. Crossing it from below implies that the reform process stalls. Mathematically, α^* is defined by $E_t\{R_\gamma|i \in (\alpha^*, 1 - \bar{\beta})\}$; thus, from (1), we can derive

$$\alpha^* = 2\gamma + \bar{\beta} - 1. \quad (5)$$

More formally, we can now see that if the median voter initially favors the reform (i.e., $\delta_t = 1$), the total supporting fraction (given by $\Psi_t = \alpha_t + (1 - \alpha_t - \bar{\beta})\delta_t$) will remain constant at $1 - \bar{\beta}$ if α increases to an $\alpha_{t+1} < \alpha^*$.

In this case, the revelation of $\Delta\alpha_{t+1}$ additional winners does not make α cross the critical value α^* . When α is increased, there are more individuals supporting the reform (because they have now learned that they are reform winners), but the uncertain block (which also supports the reform in this case) shrinks one-for-one with the increase in α . On balance, total support for reform remains unaffected.

However, as more winners are revealed, α will eventually exceed α^* . If this happens when $\alpha^* < 1/2$ (a condition to which we return to in the next subsection), the median voter switches sides and begins opposing the reform package that he used to support. A sudden loss of majority support for the reforming government results.

This opens the possibility of a reform that starts well (individuals involved with reformed firms/sectors turn out to be better off), but as individual uncertainty continues to decrease, the “sampling without replacement” effect captured by (3) eventually causes the median voter to swing against the package. Thus, the model produces support dynamics that are very much like the practical experiences of many reformist governments (as noted above).

Hence, once one accounts for individual uncertainty and the “sampling without replacement” effect, the conventional sequencing wisdom that one should begin by reforming firms or sectors that are most likely to benefit from reform (to boost public support) is challenged.

This conventional wisdom, to which we adhered before analyzing a formal model, is expressed, for example, in Roland (1994, 1164), who writes that “if the best firms get privatized first [. . .] the likelihood of a successful economic performance will be higher. Initial economic successes for privatized firms will enhance support for privatization and build constituencies for further reforms”. Similarly, *The Economist* of March 24, 2012, writes about the opposition the Cuban reform process is currently experiencing and states that to increase public support for the reform process, “Raúl Castro urgently needs to create some winners” (p. 20). However, this line of reasoning only seems to consider aggregate uncertainty and overlooks individual uncertainty and the accompanying

12. The last part of this statement, of course, assumes that $\alpha^* < 1/2$, a condition to which we will return to in the next section.

“sampling without replacement” effect. By allowing for these elements, the present paper points out that Raúl Castro may very well decrease support for his reforms even further by revealing winners.

With Aggregate Uncertainty

To capture the reasoning underlying the aforementioned conventional wisdom, which relies upon the existence of aggregate uncertainty, we next investigate what happens when we add such uncertainty to the model. In that case, voters also do not know the true value of γ (the fraction of individuals who will benefit from the reform ex post) with certainty. Instead, the public has beliefs about γ . Let us use γ_t to indicate the beginning of period t estimate of γ . Any valuable information that becomes available during period t will lead to an updated estimate, γ_{t+1} (where updating occurs via the application of Bayes’ rule).

Voters hold a prior belief about γ that is given by a $Beta(\alpha, \beta)$ distribution. This distribution is a natural choice because it is the conjugate prior of the binomial distribution underlying the present model. Assuming a symmetric loss function (for example, the traditional quadratic loss function underlying OLS, which also has the convenient property that the point estimate for gamma summarizes all relevant information and therefore becomes the sole object of interest) then implies that for $\alpha = \alpha_t$ and $\beta = \beta_t$, the time t estimate of γ equals

$$\hat{\gamma}_t = \frac{\alpha_t}{\alpha_t + \beta_t}. \quad (6)$$

Expression (6) is intuitive: $(\alpha_t + \beta_t)$ represents the total sample of outcomes we have gathered so far, whereas α_t is the fraction of winners in this sample relative to the unit interval. The ratio of these two is the time t estimate of γ .

After revealing $\Delta\alpha_{t+1}$ additional winners and $\Delta\beta_{t+1}$ additional losers during period t , Bayes’ rule implies that the posterior estimate of γ , which is the prior at the beginning of period $t+1$, equals (see, e.g., Kvam and Vidakovic (2007, chapter 4))

$$\hat{\gamma}_t = \frac{\alpha_t + \Delta\alpha_{t+1}}{\alpha_t + \Delta\alpha_{t+1} + \beta_t + \Delta\beta_{t+1}}. \quad (7)$$

From equations (6) and (7), one can confirm the intuitive notion that the application of Bayes’ rule leads to an upward revision of the expected fraction of reform beneficiaries (γ) when $\Delta\alpha_{t+1}$ additional winners are revealed (and vice versa after the revelation of losers). Because this implies that beliefs about γ can change over time, the critical value for α (α^*) also becomes time varying. In particular, after plugging (6) into (5), we obtain

$$\alpha_t^* = 2\hat{\gamma}_t + \beta_t - 1. \quad (8)$$

Now, the key question is whether we can obtain $\alpha_t \geq 1/2$ before $\alpha_t \geq \alpha_t^*$. If this is the case, the government is able to reveal that the median voter is a reform winner (which

happens when α_t crosses $1/2$) before this pivotal voter begins opposing the reform package (which happens if and only if α_t exceeds α_t^* while $\alpha_t < 1/2$). Subsequently, the government can complete the reform with no risk of losing majority support.

Because empirical studies such as those of Carlin and Mayer (1992), Frydman, Rapaczynski, and Earle (1993), Marcincin and Van Wijnbergen (1997), and Gupta, Ham, and Svejnar (2008) all present evidence that reforms start by revealing ex post winners, it is interesting to see what our model predicts would happen if the reform follows such a selective path.¹³ To investigate this, we make the following assumption on the sequencing within the reform:

Assumption 1. *Sequencing is such that the reform starts by revealing ex post winners.*

The reason for the presence of this selection bias can be twofold. First, it can result from a situation of asymmetric information in which the government knows ex ante who will benefit and who will lose from reform (but, as in Perotti (1995), the government is unable to transmit this information credibly to the public).¹⁴

Especially in our privatization example, this assumption seems realistic because the government (as the incumbent owner of the firms that are to be privatized) has inside information on firm profitability and future policies that may benefit or harm each firm. If this government then follows the conventional wisdom and begins by reforming the ex post winners (which is often recommended to reformers in practice; cf. Roland 2000, 49), Assumption 1 materializes.

Second, in light of our application to privatization, Assumption 1 can also result from the fact that better firms tend to find buyers more rapidly (Roland, 2000, 248). This point has also been recognized by policy makers. According to Egyptian government officials in the *New York Times* of June 27, 2010, Egypt suspended its privatization program in 2009 because “most of the likely candidates had already been either privatized or dissolved, leaving hard-to-sell industries that were technologically outdated and overstaffed with ill-trained workers”.

In the Appendix, we explore the alternative case in which the government is not able to identify winners and losers in advance. That case is probably more relevant to trade reform because this type of reform does not come with a natural selection process, and it is not clear that the government knows the identity of the winners and losers up front in that setting. Then, reform outcomes are sampled randomly from the true underlying distribution. Crucially, the Appendix shows that the “sampling without replacement” effect continues to be

13. Roland (2008, 4) nicely summarizes this literature by noting, “The few studies on the determinants of privatizations suggest that the more profitable firms were privatized first, which is consistent with political economic theories of privatization where the sequencing of privatization is used to gather support for further privatization”.

14. In this sense, the government in our model is a bit like Monty Hall in the “Monty Hall problem”: he knows ex ante behind which doors the gains and losses are located. Note that the counterintuitive solution to the Monty Hall problem follows from the fact that sampling takes place without replacement.

present under random sampling. This leads to two regions in the (α, β) space where the dynamics are anomalous (i.e., favorable reform outcomes decreasing support for reform and vice versa). More generally, the importance of the “sampling without replacement” effect is increasing in the tightness of the prior belief on γ : the tighter the prior on γ is, the less responsive voters’ beliefs on γ are to news and the more dominant the “sampling without replacement” effect becomes (because the latter works independently of the tightness of γ ’s prior). In the limit, because the prior on γ converges on a point, the model collapses to the one discussed in the previous subsection (without aggregate uncertainty).

Turning to the setting in which Assumption 1 holds, it is instructive to first think through what would happen if voters do *not* account for the selection bias and hold a diffuse prior belief on γ at the start of the reform (call this “time 0”). In particular, let us assume that both α_0 and β_0 (the fractions of winners and losers whose identities are clear ex ante) are close to zero (which minimizes the tightness of the prior). Then, Bayes’ rule implies that voters’ beliefs about γ are revised upward when winners are revealed ($\partial \hat{\gamma}_t / \partial \alpha_t = \beta_t / (\alpha_t + \beta_t)^2 > 0$). In particular, voters’ beliefs about γ will quickly converge to 1 because voters only observe favorable reform outcomes and erroneously think that this is the result of random sampling from the underlying true distribution of winners and losers. This implies that $\alpha_t^* \rightarrow 1$ (cf. equation (8)), which allows the reforming government to reveal that the median voter is a reform winner before $\alpha_t \geq \alpha_t^*$ (provided that $\gamma > 1/2$, of course). Subsequently, the reform can be completed with no risk of losing majority support.

Therefore, when voters have a diffuse prior belief on γ at the start of the reform and when they do not take the selection bias into account, the government is able to complete efficient reforms gradually by revealing winners while running no risk of losing majority support.

This case, however, imposes an unrealistically high degree of naivety on voters; they think that the reform is sequenced in a truly random way and do not take into account that the government (or nature; recall our discussion following Assumption 1) starts by revealing reform winners.

Consider, therefore, the more realistic case in which the public does consider the selection bias. Then, the revelation of additional outcomes provides no valuable information; the public realizes that these draws do not come from the true underlying distribution, as a result of which, Bayes’ rule no longer leads to a revision of the prior belief.¹⁵

15. The only thing that is revealed if α_t is increased to α_{t+1} , is that $\gamma \geq \alpha_{t+1}$ (which was already known given that $\alpha_{t+1} < 1/2$, whereas voters know that the reform is efficient, i.e., $\gamma > 1/2$). However, under Assumption 1, this is by no means informative about how many winners are located beyond $1/2$. More formally, Bayes’ rule states that $P(A|B) = P(A)P(B|A)/P(B)$, where “A” is a particular hypothesis of interest (e.g., $\gamma > 0.6$) and “B” represents the new incoming data ($\Delta\alpha_{t+1}$ or, equivalently, $\Delta\beta_{t+1}$). When sampling is such that the winners are revealed first, $\Delta\beta_{t+1}$ will *always* be equal to 0 as long as we move to an $\alpha_{t+1} < \gamma$, both conditional on hypothesis A as well as unconditionally (remember that the latter case still conditions on voters knowing that $\gamma > 1/2$). Consequently, $P(\Delta\beta_{t+1} = 0|A) = P(\Delta\beta_{t+1} = 0) = 1$ and Bayes’ rule implies that the posterior belief equals the prior belief (i.e., $P(A|B) = P(A)$). In this case, beliefs are no longer revised.

Consequently, agents cannot update their estimate of γ , and $\hat{\gamma}_t$ remains constant at $\gamma_0 \forall t$ (where γ_0 is the exogenously given belief on γ at the start of the reform). This leads to the following results.

Proposition 1. *If the public believes that the reform starts by revealing the ex post winners and if it believes that the reform is “sufficiently efficient” (in the sense that $\hat{\gamma}_0 \geq 3/4 - 1/2\beta_0$), the reform can still be completed gradually by revealing only winners from time 0 onwards.*

Proof. From equation (8), it follows that $\hat{\gamma}_0 \geq 3/4 - 1/2\beta_0 \Leftrightarrow \hat{\alpha}_0^* \geq 1/2$. Revealing only winners (keeping β_t constant at β_0) implies that $\hat{\alpha}_t^*$ remains constant at $\hat{\alpha}_0^* \geq 1/2$ over time. This implies that the threshold $\hat{\alpha}_t^* \geq 1/2 \forall t$, as a result of which the reforming government can reveal that the median voter is a reform winner before this voter begins opposing the reform (i.e., the government can push $\hat{\alpha}_t \geq 1/2$ before $\alpha_t > \hat{\alpha}_t^*$).

Proposition 2. *If the public believes that the reform starts by revealing the ex post winners but if $\hat{\gamma}_0 < 3/4 - 1/2\beta_0$, even reforms that are believed to be efficient (i.e., reforms for which $\hat{\gamma}_0 > 1/2$) can never be completed gradually by revealing only winners from time 0 onwards.*

Proof. From (8), it now follows that $\hat{\gamma}_0 < 3/4 - 1/2\beta_0 \Leftrightarrow \hat{\alpha}_0^* < 1/2$. This implies that the reform is not believed to be “sufficiently efficient” (as defined in Proposition 1). Revealing only winners (i.e., keeping β_t constant at β_0) then decreases the expected value of the median voter via (3). Because the constancy of β_t again implies constancy of $\hat{\alpha}_t^*$ (at $\hat{\alpha}_0^* < 1/2$), $\alpha_t > \hat{\alpha}_0^*$ before $\alpha_t > 1/2$, and majority support is lost before the reform is completed.

Hence, those reforms on which initial prior beliefs are not sufficiently optimistic (such that $\hat{\gamma}_0 < 3/4 - 1/2\beta_0$) can no longer be completed gradually by revealing only winners. This result arises even if the reform enjoys majority support at its beginning and even if the reform is believed to be efficient (in the sense that it is believed to generate more winners than losers, i.e., $\hat{\gamma}_0 > 1/2$).

The intuition for what is going on is exactly as in the previous subsection: every additional winner revealed reduces the perceived probability of ending up as a winner for those who remain uncertain. As a result, the median voter will, at some point, begin opposing the reform that he used to support.

Similar dynamics arise when we drop Assumption 1 and instead assume that reform outcomes are sampled *randomly* from the underlying distribution (see the Appendix for a discussion of this case). Then, the revelation of winners also implies that voters become more enthusiastic about the reform as they revise their estimate of the aggregate fraction of winners ($\hat{\gamma}$) in the upward direction. However, when the prior belief on the aggregate state is sufficiently tight, the

updating process in the aggregate dimension will proceed at a rather slow pace, and the “sampling without replacement” effect will dominate.¹⁶

Returning to the setup in which Assumption 1 does hold, revealing losers immediately ends majority support.

Proposition 3. *If the public believes that the reform starts by revealing the ex post winners, any reform will lose majority support as soon as a loser is revealed before $\alpha_t > 1/2$.*

The proof is intuitive and simply follows from the fact that the public expects the government to start by revealing reform winners. If a loser appears, the public thinks that all winners have already been revealed and that those individuals who are still uncertain about their identity will all be losers. If this happens while $\alpha_t < 1/2$, majority support is immediately lost.

Summarizing

Fernandez and Rodrik (1991) pointed out that welfare-enhancing reforms that would enjoy majority support ex post may not enjoy majority support ex ante because the reform winners cannot always be identified up front.

In a way, the message of this paper is more discouraging: even welfare-enhancing reforms that enjoy *both ex ante and ex post* majority support may still not come to completion because of the learning dynamics that are triggered through the initiation of the reform process. Revealing winners launches the “sampling without replacement” effect (as a result of which majority support will be lost at some point if initial beliefs about the aggregate dimension of the reform are not sufficiently optimistic), whereas revealing losers immediately ends support. The reforming government thus finds itself seemingly trapped and destined to lose majority support irrespective of what action it takes.

HOW CAN THE LOSS-OF-SUPPORT PROBLEMS BE AVOIDED?

Is there anything reformers can do to overcome these loss-of-support problems? The Chinese reform experience in particular suggests that a route exists toward successful gradual reform. After all, China also followed a more gradual path, and with quite some success. In sharp contrast to the experiences of many Latin American and Central and Eastern European countries, the initial Chinese reform successes seem to have only increased support for further reforms.

16. To see how the speed of updating is inversely related to the tightness of the prior, consider two priors with an identical mean estimate, $\hat{\gamma}_t = \alpha_t/(\alpha_t + \beta_t) = \alpha'_t/(\alpha'_t + \beta'_t)$, where $\alpha'_t > \alpha_t$ and $\beta'_t > \beta_t$. Consequently, by the standard formula for the variance of the Beta-distribution, $Var_{\gamma}(\alpha_t, \beta_t) > Var_{\gamma}(\alpha'_t, \beta'_t)$, so the prior driven by (α'_t, β'_t) is tighter than that driven by (α_t, β_t) . Now consider a *given* $\Delta\alpha_{t+1}$ (of equal size in both cases). Because $(\alpha'_t + \beta'_t) > (\alpha_t + \beta_t)$, equation (7) implies that this will lead to a smaller upward revision for the tighter prior driven by (α'_t, β'_t) . This holds for other distributions as well. Intuitively, an agent with a tighter prior is more certain that the true $\gamma = \hat{\gamma}_t$, as a result of which he revises his beliefs by less after receiving new information. Consequently, learning will occur at a slower pace for such an agent.

Obviously, the voting mechanism is absent in China (also recall footnote 9), but reforms there could still generate dissatisfaction and opposition, which does not seem to be the norm in China (cf. Litwack and Qian 1998; Qian, Roland and Xu, 1999, 2006). This raises the question of why the experiences with gradualism have been so different across countries.

In this respect, it is crucial to note that the Chinese gradual reform strategy differs from the Latin American and Central and Eastern European approaches. Whereas most countries in the latter regions tried to reform gradually along the *sectoral* dimension (which implies that one firm or sector is reformed after the other; cf. our discussion of the “case-by-case” approach to privatization in the introduction), China reformed gradually along the *spatial* dimension. In particular, China first introduced market forces in 1978 in the inland province of Sichuan (which was the first province to abolish collective agriculture and begin state-owned enterprise reform) and in the coastal province of Guangdong in 1980.

By reforming gradually along the spatial dimension, Chinese policy makers enabled the Chinese public to learn about the effects of new policies by looking at outcomes in reformed regions. Of course, the citizenry will only find the information generated via the spatial dimension useful if those regions are believed to be informative to the rest of the country. Here, China had an advantage over many other countries. As noted by Qian, Roland, and Xu (2006, 394), the Chinese economy is organized along territorial lines. This implies that its regions (such as Sichuan) are rather self-contained and relatively representative of the Chinese economy as a whole. As Démurger et al. (2002) argue, this is the result of a conscious decision made by Mao Zedong. In addition to the two key principles of Soviet development (common ownership and central planning), Mao added a third principle: regional economic self-sufficiency. This principle required each region to be self-sufficient, not only in food production but also in industrial goods. The Soviet Union did not adhere to this principle at all; their ideology called for an organization of the country along industrial lines with high degrees of industrial concentration (Qian and Xu 1993). Consequently, each Soviet region was much more specialized, dependent on other regions, and less representative of the Union as a whole.

Because of this, China had (in contrast to, for example, Russia) the possibility to start reforms by taking informative samples of small mass (in the form of certain regions) and using them to show the public where the gains and losses of the proposed reform were likely to occur. In particular, the coastal Special Economic Zones were instructive to inhabitants of other coastal regions (with approximately one-third of the total Chinese population living near the coast), whereas Sichuan fulfilled a similar role for inland areas. This strategy reduces individual uncertainty about the distribution of gains and losses (the root of all problems), but - crucially - insulates the main part of the country from the “sampling without replacement” effect.¹⁷

17. Obviously, this effect will be present within the reform region itself, but if that is only a region of relatively small mass, then the reform runs no risk of losing country-wide majority support.

For the main part of the country, this sampling strategy does not affect the distribution from which future sampling will take place because it is a form of sampling from a different, smaller urn (where the distribution of balls in this smaller urn is taken randomly from the large urn, the latter representing the main part of the country, which remains untouched in this sampling strategy).

An alternative way to think about this is by conceiving a model that has two dimensions:¹⁸ a sectoral one (as in our model presented above) and a spatial one. For a country that is perfectly diversified, these two dimensions are orthogonal to each other, as a result of which reforming gradually along the spatial dimension does not trigger the “sampling without replacement” effect (which exists along the sectoral dimension¹⁹). When there is some correlation between the two dimensions (as is the case when certain sectors are concentrated in particular areas), orthogonality no longer holds, and sampling without replacement begins to play a role again. In the extreme situation where there is a perfect correlation between “sectors” and “space”, the two dimensions merge and we are back in our original one-dimensional setup where *any* gradual reform strategy launches the counterintuitive support dynamics.

For this reason, a spatial reform strategy would not have been a viable option for Russia: over there, reform outcomes in one region were not only less relevant to those in other regions (due to the higher degree of spatial heterogeneity) but the higher degree of industrial concentration would have also given a “sectoral flavor” to any spatial reform strategy. After all, if certain sectors are concentrated in certain areas, reforming one area is equivalent to reforming one sector. Then, “sampling without replacement” would re-enter the story.²⁰

For the spatial strategy to work, it is crucial that agents who know that they will be among the winners (i.e., those with $i < \alpha_t$) cannot self-select into the reformed regions because doing so would imply that the zone becomes less instructive to the relevant other parts of the country. To continue the urn analogy, the smaller urn needs to be isolated from the larger one. Interestingly, this is precisely what the Chinese “hukou” system (which restricts the mobility of citizens within

18. Thanks to our referees for putting us on track for this interpretation. Formally developing such a two-dimensional model goes beyond the scope of this paper but could be an interesting avenue for future research.

19. The “sampling without replacement” effect exists due to heterogeneity (the existence of winners and losers), but in a country that is perfectly diversified, all regions are alike in terms of economic structure (each district is a miniature version of the country as a whole), so there is no heterogeneity along the spatial dimension.

20. As Dani Rodrik pointed out to us, the spatial strategy could be applied to reforms that entail only one particular sector (“sector X”). The “sampling without replacement” effect would then operate *within* sector X (because cross-sectoral heterogeneity no longer plays a role). In the case of one-sector reforms, governments could get around this effect by first transforming a region of small mass that contains a representative sample of sector X firms into a Special Economic Zone. Subsequently, the sector X firms inside the Zone could be used to reduce individual uncertainty by giving other firms in the sector outside the Zone an idea of what the reform would do to them. The Special Economic Zones in Malaysia and Mauritius seem to have fulfilled such a role successfully for the electronics and apparel sectors, respectively (Auty 2011).

China) achieves. Therefore, although one could debate the fairness of this system (just like one could debate the fairness of mobility restrictions between different countries), it does seem to play an economic role in the Chinese reform process.

Finally, this view of Special Economic Zones shines a new light on their *raison d'être*. In a static setup, Hamilton and Svensson (1982) show that Special Economic Zones are actually welfare decreasing in a second-best world where the suboptimal regime continues to apply outside the zone. This raises the question of why governments bother installing them in the first place. In this respect, the present paper argues that Special Economic Zones could produce large *dynamic* gains because they can facilitate the implementation of reforms that bring the entire country closer to the first-best.

Our results thus indicate that countries can ease their reform process if they have the possibility to begin the reform by first implementing it in a region of small mass that is instructive to the rest of the country (because such an action decreases individual uncertainty). In this sense, there is an important difference between sectoral and spatial gradualism. This difference may be key as to why the gradual reform strategy has worked for China but has worked much less for many other countries.

CONCLUSION

In this paper, we have modeled the learning process surrounding economic reforms when there is both aggregate and individual uncertainty. The process of revealing reform outcomes entails sampling *without* replacement. We have shown that this implies that the revelation of winners early in the reform process makes those who remain uncertain about whether they will gain or lose from reform more pessimistic about their chances of ending up as reform winners. This channel can be so strong that it can even induce the median voter to start opposing a reform that he used to support - in which case the reforming government loses majority support. As a result, even reforms that enjoy *both ex ante and ex post* majority support may still not come to completion. The conditions under which such a destructive interaction between rational learning and political support will occur are relatively mild. As soon as one combines the presence of individual uncertainty with rational belief updating, the “sampling without replacement” effect kicks in.

This situation challenges the conventional wisdom that sequencing should be such that favorable reform outcomes are revealed first. Instead, our model illustrates that a reform strategy based on revealing winners first may backfire. The reason is that such an approach leads to a deterioration in the quality of the remaining pool, triggering reform fatigue in spite of the successes of those firms that have already been reformed. Strikingly, this is consistent with the puzzling experiences that many reformers have had in practice. There are numerous examples of reforms that began while enjoying majority support but subsequently lost this support even though they were progressing in a successful way.

We have also outlined a strategy that is able to overcome the problems related to the learning process. In particular, if a country happens to contain individual regions that are instructive to the rest of the country, the “sampling without replacement” effect can be avoided by reforming gradually along the *spatial* (rather than along the *sectoral*) dimension. This could explain the success of the “Special Economic Zone” approach to reforms taken by China.

On a more general level, this paper has developed a theory of agents who are learning from realizations that are sampled without replacement, which may have other applications as well. In many environments, the distribution from which future sampling will take place is not static and time invariant; it changes over time, often as a result of past sampling actions (the process of authorities revealing the identity of good and bad banks in a financial crisis is a clear example). From the so-called “Monty Hall problem”, we know that this can give rise to counterintuitive dynamics that are not easy to understand without formalism (whereas a formal model clarifies things significantly). Applying the concepts developed in this paper to problems that entail sampling without replacement might therefore be a fruitful avenue for further research.²¹

APPENDIX

Although there are certain reform types (such as privatization) for which the government is likely to have an *ex ante* idea about where the gains and losses of the reform will be located, there are also cases in which the reforming government does not have such information. This Appendix therefore explores the properties of our model when we drop the assumption that the reform is sequenced in a non-random, selective way.

Suppose that the government cannot identify the reform winners and losers up front. In addition, we assume that there is no natural selection process that could lead to a non-random sequencing of events. Consequently, reform outcomes are sampled randomly from the true distribution. Moreover, we assume that the public *believes* that these outcomes are sampled randomly, as a result of which they perceive new observations to be informative and apply Bayes’ rule to update their estimate of γ (the aggregate fraction of winners) in response to new information.

Because we feel that this case may deserve a closer study in its own right, we leave a full analysis for future work, but we present some main results in this

21. In this respect, it was pointed out to us that the mechanisms underlying this paper also relate to those in Konrad (2004, although he does not establish the link to sampling without replacement, nor does he allow for aggregate uncertainty). However, that paper addresses strategic rather than dynamic issues and shows that a party that is campaigning to implement a particular reform may obtain an incentive to point out *ex ante* who will be a reform *loser* (a phenomenon that Konrad refers to as “inverse campaigning”).

Appendix. In particular, we show that the “sampling without replacement” effect discussed in the main text continues to in this alternative setting.

The core of the model is unaffected, and the critical value for α (α^*) is still given by

$$\alpha_t^* = 2 \frac{\alpha_t}{\alpha_t + \beta_t} + \beta_t - 1. \quad (\text{A.9})$$

Using expression (A.1), one can analyze how the distance between α_t (the fraction of sure-winners) and α_t^* (the cut-off level for α_t above which the median voter starts opposing the reform) varies with the revelation of additional winners and losers. In particular, it holds that

$$\frac{\partial(\alpha_t^* - \alpha_t)}{\partial \alpha_t} = \frac{2\beta_t}{(\alpha_t + \beta_t)^2} - 1. \quad (\text{A.10})$$

Here, the first term shows that the revelation of winners pushes up α_t^* (because it leads to an upward revision of the expected fraction of reform winners $\hat{\gamma}_t$ through application of Bayes’ rule), whereas the second term (“−1”) indicates that the revelation of winners simultaneously makes those who remain uncertain more pessimistic about their individual chances of ending up winners. In particular, this term reflects the fact that revealing a reform outcome is an example of sampling without replacement. From (A.2), one can derive that as long as

$$\alpha_t > \sqrt{2\beta_t} - \beta_t, \quad (\text{A.11})$$

the “sampling without replacement” effect dominates. Under condition (A.3), the public’s estimate of γ increases less than one-for-one with α_t (mathematically, $\partial(\alpha_t^* - \alpha_t)/\partial \alpha_t < 0 \Leftrightarrow \partial \hat{\gamma}_t/\partial \alpha_t < 1$), and the median voter becomes more *pessimistic* as *favorable* reform outcomes are increasingly revealed. Hence, under this condition, the revelation of additional winners produces an increase in $\hat{\gamma}_t$ that is insufficient to compensate for the fact that sampling occurs without replacement.

Similarly,

$$\frac{\partial(\alpha_t^* - \alpha_t)}{\partial \beta_t} = \frac{-2\alpha_t}{(\alpha_t + \beta_t)^2} + 1 \quad (\text{A.12})$$

captures the same two effects for the revelation of losers. In this case, the median voter becomes more *optimistic* when additional *losers* are revealed (i.e., $\partial(\alpha_t^* - \alpha_t)/\partial \beta_t > 0$) as long as

$$\beta_t > \sqrt{2\alpha_t} - \alpha_t. \quad (\text{A.13})$$

FIGURE 2. Regions Where Favorable Reform Outcomes Decrease Support for the Reform (A3) and Vice Versa (A5)

Conditions (A.3) and (A.5) yield two regions of (α, β) combinations, displayed as the shaded areas in Figure 2, where one can characterize the learning dynamics as “anomalous”. That is, in region A3, good reform outcomes decrease support for the reform, whereas the revelation of bad reform outcomes increases support for reform in region A5.

Now, one can ask whether the government is able to complete the reform without a loss of majority support along the way. Because the sequencing of the reform is random in this case (because the government (or nature) is no longer able to select the winners up front), it is no longer possible to analyze this question analytically. Instead, one would have to simulate the reform process, and the answer to the question would depend upon the amount of time a typical simulation spends in the shaded areas of the state space. Because we feel that this issue deserves a full discussion in its own right, we leave this for future work.

The main point to take away from this Appendix is that the “sampling without replacement” effect continues to be present when reform outcomes are revealed in a truly random fashion. This leads to two regions in the (α, β) space where the support dynamics can be characterized as “anomalous”.

REFERENCES

- Auty, R. 2011. “Early Reform Zones: Catalysts for Dynamic Market Economies in Africa.” In: T. Farole, and G. Akinci, eds., *Special Economic Zones*. Washington, DC: World Bank.
- Berg, A., and O. J. Blanchard. 1994. “Stabilization and Transition: Poland 1990-91.” In: O. J. Blanchard, K. A. Froot, and J. D. Sachs, eds., *The Transition in Eastern Europe*, Vol. 1. Cambridge, MA: MIT Press.

- Boycko, M., A. Shleifer, and R. W. Vishny. 1993. "Privatizing Russia." *Brookings Papers on Economic Activity* 2 : 139–92.
- Carlin, W., and C. Mayer. 1992. "Restructuring Enterprises in Eastern Europe." *Economic Policy* 7 (15) : 311–52.
- Démurger, S., J. D. Sachs, W. T. Woo, S. Bao, G. Chang, and A. Mellinger. 2002. "Geography, Economic Policy, and Regional Development in China." *Asian Economic Papers* 1 (1) : 146–97.
- Dewatripont, M., and G. Roland. 1992. "The Virtues of Gradualism and Legitimacy in the Transition to a Market Economy." *Economic Journal* 102 (411) : 291–300.
- . 1995. "The Design of Reform Packages under Uncertainty." *American Economic Review* 85 (5) : 1207–23.
- Fernandez, R., and D. Rodrik. 1991. "Resistance to Reform: Status Quo Bias in the Presence of Individual-Specific Uncertainty." *American Economic Review* 81 (5) : 1146–55.
- Fidrmuc, J. 2000. "Political Support for Reforms: Economics of Voting in Transition Countries." *European Economic Review* 44 (8) : 1491–513.
- Frydman, R., A. Rapaczynski, and J. Earle. 1993. *The Privatization Process in Central Europe*. London: Central European University Press.
- Gupta, N., J. C. Ham, and J. Svejnar. 2008. "Priorities and Sequencing in Privatization: Evidence from Czech Firm Panel Data." *European Economic Review*. 52 (2) : 183–208.
- Hamilton, C., and L. E. O. Svensson. 1982. "On the Welfare Effects of a Duty Free Zone." *Journal of International Economics* 13 (1-2) : 45–64.
- Iglesias, E. 1994. "Economic Reform: A View from Latin America." In: J. Williamson, ed., *The Political Economy of Policy Reform*. Washington, DC: Institute for International Economics.
- Konrad, K. A. 2004. "Inverse Campaigning." *Economic Journal* 114 (492) : 69–82.
- Kornai, J. 1990. *The Road to a Free Economy*. New York, NY: W.W. Norton and Company.
- . 2000. "Ten Years after 'The Road to a Free Economy': The Author's Self-Evaluation." *Economic Systems* 24 (4) : 353–9.
- Kramer, G. 1971. "Short-Term Fluctuations in US Voting Behavior, 1896-1964." *American Political Science Review* 65 (4) : 131–43.
- Kvam, P. H., and B. Vidakovic. 2007. *Nonparametric Statistics with Applications to Science and Engineering*. Hoboken, NJ: John Wiley & Sons.
- Lipton, D., and J. D. Sachs. 1990. "Privatization in Eastern Europe: The Case of Poland." *Brookings Papers on Economic Activity* 2 : 293–341.
- Litwack, J. M., and Y. Qian. 1998. "Balanced or Unbalanced Development: Special Economic Zones as Catalysts for Transition." *Journal of Comparative Economics* 26 (1) : 117–41.
- MacKuen, M. B., R. S. Erikson, and J. A. Stimson. 1992. "Peasants or Bankers? The American Electorate and the US Economy." *American Political Science Review* 86 (3) : 597–611.
- Marcincin, A., and S. van Wijnbergen. 1997. "The Impact of Czech Privatization Methods on Enterprise Performance Incorporating Initial Selection-Bias Correction." *Economics of Transition* 5 (2) : 289–304.
- Messner, M., and M. K. Polborn. 2004. "Voting on Majority Rules." *Review of Economic Studies* 71 (1) : 115–32.
- Morrow, J., and M. Carter. 2013. "Left-Right-Left: Income, Learning and Political Dynamics." NBER Working Papers 19498, National Bureau of Economic Research, Cambridge, MA.
- Murphy, K. M., A. Shleifer, and R. W. Vishny. 1992. "The Transition to a Market Economy: Pitfalls of Partial Reform." *Quarterly Journal of Economics* 107 (3) : 889–906.
- Perotti, E. C. 1995. "Credible Privatization." *American Economic Review* 85 (4) : 847–59.
- Pleskovic, B., and J. D. Sachs. 1994. "Political Independence and Economic Reform in Slovenia." In: O. J. Blanchard, and K. A. Froot, and J. D. Sachs, eds., *The Transition in Eastern Europe*, Vol. 1. Cambridge, MA: MIT Press.
- Qian, Y., G. Roland, and C. Xu. 1999. "Why is China Different from Eastern Europe? Perspectives from Organization Theory." *European Economic Review* 43 (46) : 1085–94.

- . 2006. “Coordination and Experimentation in M-Form and U-Form Organizations.” *Journal of Political Economy* 114 (2) : 366–402.
- Qian, Y., and C. Xu. 1993. “Why China’s Economic Reforms Differ: The M-Form Hierarchy and Entry/Expansion of the Non-State Sector.” *Economics of Transition* 1 (2) : 135–70.
- Remmer, K. L. 1991. “The Political Impact of Economic Crisis in Latin America in the 1980s.” *American Political Science Review* 85 (3) : 777–800.
- Rodrik, D. 1995. “The Dynamics of Political Support for Reform in Economies in Transition.” *Journal of the Japanese and International Economies* 9 (4) : 403–25.
- Roland, G. 1994. “On the Speed and Sequencing of Privatization and Restructuring.” *Economic Journal* 104 (426) : 1158–68.
- . 2000. *Transition and Economics*. Cambridge, MA: MIT Press.
- . 2008. *Privatization: Successes and Failures*. New York, NY: Columbia University Press.
- Sachs, J. D., and A. Warner. 1995. “Economic Reform and the Process of Global Integration.” *Brookings Papers on Economic Activity* 1 : 1–118.
- Stokes, S. C. 2001. *Public Support for Market Reforms in New Democracies*. Cambridge: Cambridge University Press.
- Strulovici, B. 2010. “Learning While Voting: The Determinants of Collective Experimentation.” *Econometrica* 78 (3) : 933–71.
- Tommasi, M., and A. Velasco. 1996. “Where Are We in the Political Economy of Reform?” *Journal of Policy Reform* 1 (2) : 187–238.
- Tucker, J. A. 2000. “It’s the Economy, Comrade! Economic Conditions and Election Results in Russia, Poland, Hungary, Slovakia, and the Czech Republic.” PhD Thesis, Harvard University.
- van Wijnbergen, S. 1992. “Intertemporal Speculation, Shortages and the Political Economy of Price Reform.” *Economic Journal* 102 (415) : 1395–406.
- Veldkamp, L. 2009. “Learning about Reform: Time-Varying Support for Structural Adjustment.” *International Review of Economics and Finance* 18 (2) : 192–206.