



UvA-DARE (Digital Academic Repository)

Essays on optimal experimentation

Willems, T.

Publication date
2012

[Link to publication](#)

Citation for published version (APA):

Willems, T. (2012). *Essays on optimal experimentation*. [Thesis, fully internal, Universiteit van Amsterdam]. Thela Thesis.

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.

Chapter 1

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

Down here it's just winners and losers and don't get caught on the wrong side of that line.

Bruce Springsteen, *Atlantic City* (1982)

1.1 Introduction

Why have gradual economic reforms worked out well for China, while this is much less so for most Latin American and Central and Eastern European countries? And how is it possible that so many reforms that started successfully while enjoying majority support, subsequently lost this support, while there are simultaneously examples of reforms that did not start off that well, but nevertheless managed to maintain momentum among voters?

The most dramatic example of a reformist government that lost majority support despite impressive economic performance, is probably that of Slovakia in 2006. Back then, the *Wall Street Journal Europe* wrote:¹

"Imagine you're the leader of a country where economic growth

⁰This chapter is based upon joint work with Sweder van Wijnbergen.

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

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. [...]

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."

Indeed, Dzurinda lost the 2006 elections and his government was replaced by a much more conservative one led by Robert Fico of the SMER-party (a breakaway party from the SDL, the successor to the original Communist Party of Slovakia).

In an attempt to explain these puzzling support dynamics, this paper constructs a learning model that tries to provide more insight into the dynamics of public support for gradual economic reforms.² We believe that learning behavior plays a key role in determining the support dynamics for reforms, but so far the literature has remained relatively silent on this issue: although there are many informal discussions of learning from reform outcomes, formal treatments are much scarcer.³

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 be identified up

²Here, "gradual" can, but should not necessarily, be interpreted as opposing "big bang". As for example noted by Gupta, Ham and Svejnar (2008), even reforms that were supposed to be "big bang" (such as the Balcerowicz reform in Poland), were not completed instantaneously. Instead, it is almost inevitable that some parts of the reform are completed prior to others (introducing a sequential dimension). This is the notion of "gradualism" we have in mind in this paper. In this sense, all reforms are thus gradual, but some reforms are "more gradual" than others.

³The papers by Van Wijnbergen (1992, where voters learn about the effects of price reform), Dewatripont and Roland (1995, where the public uses the outcome of one reform to learn about the expected outcome of another one) and Veldkamp (2009, where laid-off workers learn about their re-employment chances) are three exceptions.

front. Consequently, voters update their beliefs on whether they will end up in the winner-group (or in the group of losers) as the reform progresses.

To see how such a learning process can lead to the aforementioned paradoxical situation in which successful reforms lose support along the way, consider the classic Fernandez-Rodrik case in which there is no aggregate uncertainty. In that setup the public knows that, say, 68 percent of the population will gain as a result of the reform, but *ex ante* no one knows the identity of these winners. Now let us start the reform: after reforming the first n firms (or sectors), it turns out that all individuals involved with these firms have gained from the reform. Hence, to the outside observer, the reform is a big success. However, to those who remain uncertain on the effects of the reform on their payoffs, this success is actually pretty bad news: after all, they know that 32 percent of the population is going to carry the losses of the reform, and none of these losers were among the first n reformed firms. Consequently, the probability that these uncertain individuals are among the losers has gone up - thereby making them less eager to continue the reform process. We show that these dynamics can even lead to a situation in which the median voter at some point starts opposing a reform of which he used to be supportive, in which case the reforming government suddenly loses majority support.

All this follows from the fact that the process of revealing reform outcomes is an example of sampling *without* replacement (an observation that seems to have been overlooked in the previous literature) and shows that good news can, in a sense, thus be bad. We emphasize that these effects result from rational economic thinking and that they come into play as soon as a reform is believed to generate some losers whose identity is *ex ante* unknown - an assumption that we see as being in accordance with many economic reforms in reality (also see the discussion in Fernandez and Rodrik (1991) on this).

The possibility of such a loss of support nuances the commonly expressed view that one should start with reforming the firms or sectors that are most likely to benefit from the reform (so as to increase public support for it). In fact, we show that such a policy may very well *decrease* support for the reform!

In our model we formalize the underlying logic of the above considerations. To the best of our knowledge, this is the first formal treatment of an economic model where agents are learning from realizations that are sampled without replacement,

which may be of independent interest to some readers. Compared to the example sketched earlier, our model is however more general as it encompasses the case with aggregate uncertainty, where voters are also uncertain on the exact share of the population that will benefit from the reform. In that setup we show that the "sampling without replacement"-effect is more important when voters believe that the sequencing of the reform is strategic (and hence non-random), or when they have a rather tight prior belief on the aggregate fraction of winners (as a result of which Bayesian updating in that dimension occurs rather slowly).

Although the learning mechanism potentially applies to many reform types (such as the phasing out of subsidies, the abolition of price controls, or the reduction in trade barriers), we will often make the link with privatization. Next to the fact that privatization is a nice example of a reform where learning dynamics may be important, it is currently also one of the most topical reforms. Both several Southern European countries as well as important countries in the Arabic world (such as Iraq), are privatizing significant parts of their economies at this very moment.

Our results question the political feasibility of the so-called "sectoral gradualist" approaches. Those strategies have for example been advocated by the World Bank through its plea for the "case-by-case" approach to privatization (Welch and Frémond, 1998). This case-by-case approach, also favored by János Kornai in his influential 1990 book, has been applied to many countries in both Latin America and Central and Eastern Europe, as well as to the UK during its liberalization phase. It implies that one sector (or firm) is reformed after the other. Berg and Blanchard (1994) nicely describe how this process evolved in Poland. At p. 53 they write:

"Progress on the microeconomic front was slow. Progress on privatization in particular was uneven, both in 1990 and in 1991. Once new local governments were in place in the spring of 1990, privatization of retail shops proceeded steadily, mostly through leasing. By the end of 1991, it was largely achieved. A comprehensive privatization law was passed in July 1990, after intense political debate. And, as a result, some progress was also made in the privatization of small- and medium-firms, usually through lease-to-buy arrangements. But,

in sharp contrast, there was in effect no progress in the privatization of large firms. The law envisaged privatization of these large firms mainly through case-by-case sales."

Further on, at p. 63, when discussing the privatization of large firms, they note:

"The results, at least for large firms, have been very limited. Not before November 1990 were the first five firms sold through public offerings. At the end of 1991, five more firms had been sold in the same manner, and another sixteen had been sold through public tenders or auctions."⁴

By following such a gradual, sequential strategy, one may however suffer from the "sampling without replacement"-effect sketched before, which may explain why practitioners have experienced political difficulties with the case-by-case approach to privatization: Lipton and Sachs (1990, p.298) for example note that "in almost all countries where privatizations have been attempted, there have been major political obstacles to the case-by-case approach", while Boycko, Shleifer and Vishny (1993, p. 148) state that reforms that proceed at a rather slow pace, are likely to reach a deadlock. As we will argue in Section 4 of this paper, "spatial gradualism" (reforming one *region* after another, as China for example did) can avoid the "sampling without replacement"-problem under some conditions - thereby contributing towards the political feasibility of the reform process.

The itinerary of this paper is as follows. First, Section 2 describes various examples of reforms for which support dynamics have been counterintuitive. In Section 3 we construct a learning model that provides an explanation for these puzzling dynamics. Section 4 will then deal with the question why gradualism worked quite well for China, while this is much less so for most Central and Eastern European and Latin American countries. Finally, Section 5 concludes.

⁴As noted in Berg (1994), only 11 percent of Polish state enterprises had been privatized or commercialized (the step prior to privatization) two years after the start of the reform.

1.2 The support dynamics for economic reforms

Next to the case of Slovakia (discussed in the Introduction), there are many more examples of economic reforms that lost support despite the fact that they progressed in a successful manner (and vice versa). In this respect, a volume edited by Stokes (2001) provides a thorough analysis of public support dynamics around various reforms. In the volume, several authors examine the public's reactions to economic reforms in Spain, East Germany, Poland, Mexico, Peru and Argentina. In her summary of the study, Stokes (2001, p.25) notes that "[the] 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 Tucker (2000), who analyzed election data from five post-communist countries and found that support for incumbents tends to *decrease* with economic performance.

Stokes (2001) gives 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.⁵ In East Germany, high inflation rates during the early phases of the transition fed optimism about the future and increased support for the incumbent Kohl government, while increased real wages in Poland did not generate support for the reforms, but created agnosticism instead.

With respect to the latter case, Rodrik (1995, p. 404) has also expressed his surprise. When discussing the return to power of the former Polish communist party in 1993, he writes that "why this should be so is not so easy to understand [...] By most standards, Poland must be judged a success case, even though much else remains to be done (particularly on the mass-privatization front)".⁶

Next to the cases studied in Stokes (2001), there are also several other examples of countries that have had similar experiences. Puzzled by Latin American experi-

⁵In a broader sample of 12 Latin American countries, Remmer (1991) finds similar paradoxical results.

⁶Poland did have a high unemployment rate at the time, but as Rodrik (1995, p. 405) notes it is not clear whether that is to blame for the deadlock: the unemployed are too small a group to be decisive in national elections, while it is also not straightforward that their interests are best served by policies that slow down the reforms.

ences, Tommasi and Velasco (1996) for example ask: "Why did Venezuelans riot, twice attempt to overthrow and eventually impeach a president (Carlos Andrés Pérez) who in 1990-92 brought them an average growth rate of 7.8% (the highest in Latin America), while Peruvians massively reelected Alberto Fujimori, under whose stewardship consumption dropped by 15.3% in 1990?" Similarly, Iglesias (1994, p. 497-498) 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".⁷ A similar story holds with respect to India: notwithstanding the successes of the Indian liberalization policies adopted in the 1990s, "India is a place that has fallen out of love with reform" (as stated in *The Economist* of March 24th 2012, p. 14) and is currently struggling to get new reforms implemented. More generally, Sachs and Warner (1995) have documented how many countries slowed down (or even reversed) their liberalization policies in the 1960s and 70s, even though economic performance under the more liberal regime was impressive.

In Europe, Slovenia faced similar difficulties in its post-communist reform process - despite the fact that it already had quite a few positive experiences with market forces from the past (Pleskovic and Sachs, 1994). And although the 1968 Hungarian reforms started off successfully, they subsequently ran into difficulties in the mid-1970s when the country went through 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 came back in power (very much like in Poland and later on also Slovakia) - a pattern that led Kornai (2000) to conclude that the gradual reform strategy may not be that feasible from a political point of view.

With respect to the general experiences of Central and Eastern Europe, Fidrmuc (2000, p. 1491) notes: "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

⁷In 1994, Luis Alberto Lacalle (of the Partido Nacional) was president of Uruguay. After taking office in 1990, he started with significant economic reforms (in the sphere of both taxation and liberalization), but later his initiatives lost support (despite the successes; cf. Iglesias' quote). Subsequently, he lost the elections and was replaced by Julio María Sanguinetti of the rivaling Partido Colorado in 1995.

voted out." In a similar spirit, Blanchard, Froot and Sachs (1994, p.6), speak of "reform fatigue" plaguing most reformist post-communist governments some years into the reform.

From all these examples, it thus seems that a successful reform start is by no means a sufficient condition for the reform to maintain majority support along the way - an observation that is at odds with the conventional belief that a favorable start facilitates continuation.

At the other side of the spectrum, however, there are the gradual economic reforms in China. There, the government installed certain Special Economic Zones in 1980, after which the economies of those regions started booming. Very much in contrast to the experiences of Central and Eastern Europe and Latin America, the successes of these initial Chinese reforms only seem to have increased support for further reforms over there (see Litwack and Qian (1998) and 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 it is simultaneously able to explain why the initial Chinese reform successes did not invoke such a paradoxical public response.

1.3 Model

In this section we spell out our model, which builds upon Fernandez and Rodrik (1991). To obtain a better insight into the underlying mechanisms and set the stage for our full model, we first consider the standard Fernandez-Rodrik setup in which there is only individual uncertainty, but no uncertainty in the aggregate. That is: in Section 3.1 individuals know *what fraction* of the population will gain from the reform, but *ex ante* they do not know yet *who* these winners will be.

In Section 3.2 we then augment the standard Fernandez-Rodrik setting with aggregate uncertainty: individuals are also not sure *what fraction* of the population will gain from the reform (a feature of reality that has been stressed before by Dewatripont and Roland (1995)).

1.3.1 Without aggregate uncertainty

To illustrate our point without aggregate uncertainty, we make a dynamic extension to the model developed in Fernandez and Rodrik (1991). In particular, time is discrete, the horizon infinite and there is a large number of risk neutral voters aligned uniformly between 0 and 1. We assume that these individuals are fully rational and forward looking.⁸ They are faced with a reform proposal R_γ , which is to replace the status quo. The latter is assumed to leave everyone with a net present value payoff of 0. Reform R_γ , on the other hand, is known to benefit a fraction $\gamma > \frac{1}{2}$ 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).⁹ This implies that there is no aggregate uncertainty and since $\gamma > \frac{1}{2}$ the reform is efficiency-enhancing and would always be welcomed by a majority ex post.

The public does face individual uncertainty however: although there is a fraction α_t that has already been reformed at time t and has found itself among the winners, and a fraction β_t that has found out to be among the losers (with $\alpha_t < \frac{1}{2}$ and $\beta_t < \frac{1}{2}$), the remaining fraction $(1 - \alpha_t - \beta_t)$ does not know at time t whether they will gain or lose from the reform. Consequently, 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 on the right, we get the configuration shown in Figure 1.

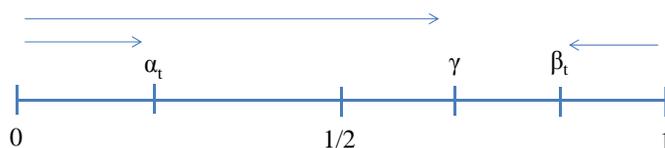


Figure 1.1: Graphical illustration of the setup

⁸Whether voters are for- or backward looking is somewhat debated. While especially 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) and Fidrmuc (2000)).

⁹Note that the distribution of winners thus follows a binomial distribution with success probability γ .

Voters with $i < \alpha$ know that they are among the reform winners, while voters with $i > \beta$ know that they are among the losers. The agents in 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. Since the identity of more and more individuals is revealed as the reform progresses over time, α and β become time-varying and thus have a time index; γ on the other hand is a structural parameter characterizing the reform and therefore time-invariant.

In this setup, the expected value of the uncertain individuals (i.e. those with $i \in (\alpha_t, 1 - \beta_t)$) equals:

$$\mathbb{E}\{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\} = (\gamma - \alpha_t) \cdot S + (1 - \gamma - \beta_t) \cdot (-S) \quad (1.1)$$

Individuals in that group follow the decision rule:

$$\delta_t = \begin{cases} 1 & \text{if } \mathbb{E}\{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\} > 0 \\ 0 & \text{if } \mathbb{E}\{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\} \leq 0 \end{cases} \quad (1.2)$$

Here, δ_t is a support indicator that takes the value 1 if the uncertain group votes in favor of the reform, and zero otherwise. Since α_t and β_t are both smaller than $\frac{1}{2}$, the median voter is located in this uncertain group.

The insight of Fernandez and Rodrik (1991) was that 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 $\alpha_t < \frac{1}{2}$ this implies that the reform does not enjoy majority support up front, even though it would be welcomed by a majority ex post (since $\gamma > \frac{1}{2}$).¹⁰ As Fernandez and Rodrik emphasize, this argument does not rely on risk aversion, irrationality, or hysteresis due to sunk costs.

Here, one sees how the presence of individual uncertainty can prevent an efficiency-enhancing reform from being accepted. 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! Since individual uncertainty is at the core of the problem, one may think that reducing individual uncertainty by

¹⁰Here, one should note that the Fernandez-Rodrik model assumes that it is not possible to compensate the losers ex post. Given the well-known difficulties that governments face in committing to future policies, this assumption may not be that unrealistic.

revealing some additional 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 a perspective on intriguing voter dynamics.

To see this, consider how the expected value for the uncertain individuals changes with α :

$$\frac{\partial \mathbb{E}_t \{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\}}{\partial \alpha_t} = -S < 0 \quad (1.3)$$

So, a decrease in individual uncertainty (brought about by the revelation of some additional winners), will make the individuals who remain uncertain only more *negative* about their chances of gaining from the reform. The reason is that in the absence of aggregate uncertainty, increasing α_t to α_{t+1} implies that there are $\Delta\alpha_{t+1}$ fewer gaining places left for those individuals who remain uncertain (since the revelation of reform outcomes is an example of sampling without replacement) - thereby making them more pessimistic about their chances of gaining from the reform. And since the median voter is located within this uncertain group, he also becomes more pessimistic.

Revealing some losers on the other hand (that is: increasing β_t to $\beta_{t+1} = \beta_t + \Delta\beta_{t+1}$), will increase the expected value of the reform for those who remain uncertain:

$$\frac{\partial \mathbb{E}_t \{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\}}{\partial \beta_t} = S > 0 \quad (1.4)$$

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. For example, if the vote is "No" to begin with, increases in α will only make $\mathbb{E}_t \{R_\gamma | i \in (\alpha_t, 1 - \beta_t)\}$ more negative and the median voter will continue to oppose the package.

But there is an intriguing possibility if the median voter initially supports the full reform package. To see this, hold β_t constant at $\bar{\beta}$ for a moment¹¹ and let us assume that the government tries to complete the reform gradually.¹² Then, if the

¹¹The 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.

¹²This choice for gradualism could either be the result of a conscious political decision (for example due to the fact that there are transitional adjustment costs as in Aghion and Blanchard (1994)), or simply by the fact that a true big bang reform (in which all reforms are completed simultaneously, before voters can update beliefs on their chances of benefiting from the reform) is practically infeasible (recall footnote 2).

increase in α is small enough, δ will remain 1 and the median voter continues to vote "Yes", pushing the overall vote in favor. But because of the effect captured by (3), one can define a critical value for α , call it α^* , such that if a subsequent series of increases brings α above α^* , the median voter swings around, causing a rejection of the package.¹³ The critical value α^* is thus the point at which the median voter starts opposing a reform that he used to support. Mathematically, α^* is defined as that value of α for which $\mathbb{E}\{R_\gamma|i \in (\alpha^*, 1 - \bar{\beta})\} = 0$, so from (1) we can derive that:

$$\alpha^* = 2\gamma + \bar{\beta} - 1 \quad (1.5)$$

More formally, one 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 if α increases to an $\alpha_{t+1} < \alpha^*$. Hence:

$$\Delta\delta_{t+1} = 0 \Leftrightarrow \Delta\Psi_{t+1} = \Delta\alpha_{t+1} (1 - \delta_t) = 0$$

if: (i) $\delta_t = 1$

$$(ii) \alpha_t + \Delta\alpha_{t+1} < \alpha^*$$

In this case, the revelation of $\Delta\alpha_{t+1}$ additional winners does not make α cross the critical value α^* . When we increase α , there are more individuals sure to vote "Yes" (since they have now found out that they are among the winners of the reform), 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 the reform remains unaffected.

But as α continues to increase, it will eventually exceed α^* at which point the median voter switches sides and will start to oppose the reform package that he used to support. A sudden loss of majority support for the reforming government results:

$$\Delta\delta_{t+1} = -1$$

if: (i) $\delta_t = 1$

$$(ii) \alpha_t < \alpha^* \leq \alpha_t + \Delta\alpha_{t+1}$$

¹³The last part of this statement of course assumes that $\alpha^* < 0.5$ - a crucial condition to which we will return to in the next section.

This opens up the possibility of a reform that starts off 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. Hereby, the model manages to reproduce support dynamics that are very much like the experiences that many reformist governments have had in practice (recall Section 2 of this paper).

Hence, when one takes into account that the process of revealing reform outcomes is an example of sampling without replacement, the conventional sequencing wisdom that one should start by reforming sectors that are most likely to benefit from the reform (in order to generate public support) is heavily challenged.

This conventional wisdom is for example expressed in Roland (1994, p. 1164). There he 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, when discussing the ongoing reforms in Cuba, *The Economist* of March 24th 2012 writes about the opposition the Cuban reform process is currently experiencing. *The Economist* then states that, in order to increase public support for the reform process, "Raúl Castro urgently needs to create some winners" (p. 20). This line of reasoning however overlooks the "sampling without replacement"-effect. By taking this effect into account, the present paper points out that Raúl Castro may very well decrease support for reforms even further by revealing winners.

1.3.2 With aggregate uncertainty

Now let us introduce aggregate uncertainty. In this case, voters do not know the true value of γ (the fraction of individuals that will benefit from the reform ex post) with certainty. In fact, they even do not know whether γ is smaller or greater than $\frac{1}{2}$ (i.e. whether the proposed reform is efficient or not). Instead, the public holds beliefs about γ . Let us use $\hat{\gamma}_t$ to indicate the beginning of period t estimate of γ . Additional information that becomes available during period t will lead to an updated estimate, $\hat{\gamma}_{t+1}$ (where updating occurs by application of Bayes' rule).

Voters hold a prior belief on γ that is given by a *Beta* (α, β)-distribution. As for example set out in Greenberg (2008), this distribution is a natural choice as

it is the conjugate prior of the binomial distribution which underlies the present model (cf. footnote 9). In this case, we have that for $\alpha = \alpha_t$ and $\beta = \beta_t$, the time t estimate of γ equals:

$$\hat{\gamma}_t = \frac{\alpha_t}{\alpha_t + \beta_t} \quad (1.6)$$

Expression (6) is intuitive: $(\alpha_t + \beta_t)$ represents the total sample of outcomes we have gathered so far, while α_t is the number of winners in this sample. 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:

$$\hat{\gamma}_{t+1} = \frac{\alpha_t + \Delta\alpha_{t+1}}{\alpha_t + \Delta\alpha_{t+1} + \beta_t + \Delta\beta_{t+1}} \quad (1.7)$$

Since 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 that:

$$\alpha_t^* = 2\hat{\gamma}_t + \beta_t - 1 = \frac{2\alpha_t}{\alpha_t + \beta_t} + \beta_t - 1 \quad (1.8)$$

Now, the key question is going to be: can we get $\alpha_t^* \geq \frac{1}{2}$ 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, before this pivotal voter starts opposing the reform package. Subsequently, the government can complete the reform with no risk of losing majority support along the way (which only happens when $\alpha_t^* \leq \alpha_t < \frac{1}{2}$). When analyzing this question, we make the following assumption on the sequencing within the reform:

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

This assumption is motivated on two grounds. First, it can result from a situation of asymmetric information where the government does know ex ante who will benefit and who will lose from the reform (but is, like in Perotti (1995), unable to transmit this information credibly to the public).¹⁴ Especially in our

¹⁴In this sense, the government in our model is thus a bit like Monty Hall in the "Monty Hall problem": he also knows ex ante behind which doors the gains and losses are located. Note that the counterintuitive solution to the Monty Hall problem also follows from the fact that sampling

privatization example this assumption seems realistic, as the government (being the incumbent owner of the firms that are to be privatized) has inside information on firm profitability, future policies that may benefit or harm each firm, etc. etc. If this government then follows the conventional wisdom and starts with reforming the winners, Assumption 1 materializes. Studies like Carlin and Mayer (1992), Frydman, Rapaczynski and Early (1993), Marcincin and Van Wijnbergen (1997) and Gupta, Ham and Svejnar (2008) all present evidence that such a selection bias indeed is present in reality.

Second, in the light of our application to privatization, Assumption 1 can also be the natural outcome resulting from the fact that better firms tend to find buyers more rapidly (Roland, 2000: p. 248).¹⁵

In the Appendix we explore the alternative case in which the government is not able to identify winners and losers up front either. That case is probably more relevant to trade reforms, as such a reform does not come with a selection process like the one described in footnote 15, while it is also not clear that the government has superior knowledge relative to the public on the identity of the winners and losers in that setting. Then, reform outcomes are just sampled randomly from the true underlying distribution. Crucially, the Appendix shows that the "sampling without replacement"-effect captured by equation (3) continues to be present under random sampling. In particular, that specification leads to two regions in the (α, β) -space where the dynamics are anomalous (i.e. favorable reform outcomes decreasing support for the 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 (since the latter works independently of the tightness of γ 's prior). In the limit, as the prior on γ converges on a point, the model essentially collapses to the one discussed in Section 3.1 (without aggregate uncertainty).

Turning to the setting in which Assumption 1 does hold, the following proposition shows that all efficiency-enhancing reforms can be completed if the public

takes place without replacement.

¹⁵This 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".

does not take the selection bias into account.

Proposition 1 *If an efficient reform (i.e. a reform for which $\gamma > \frac{1}{2}$) enjoys majority support at its start, it can always be completed gradually without loss of majority support if the winners are reformed first, without the public taking this selection bias into account.*

Proof. After the revelation of winners, voter beliefs about γ are revised upwards by Bayes' rule:

$$\frac{\partial \hat{\gamma}_t}{\partial \alpha_t} = \frac{\beta_t}{(\alpha_t + \beta_t)^2} > 0$$

The reform is assumed to enjoy majority support at its start (call this "time 0"), so $\hat{\gamma}_0 > \frac{1}{2}$. Simultaneously, the fact that $\partial \hat{\gamma}_t / \partial \alpha_t > 0$ implies that $\hat{\gamma}_t$ only increases as long as the government continues to reveal reform winners. Consequently, the location of the critical value α_t^* will increase as well (see equation (8)), thereby allowing the government to reveal that the median voters is a reform winner (which is possible because $\gamma > 1/2$) before $\alpha_t \geq \alpha_t^*$. Subsequently, the reform can be completed with no risk of losing majority support. ■

The reason that the government is able to complete the reform without a loss of support along the way, lies in the fact that it is assumed that the reform follows a strategic, selective path - without the public taking this into account. That is: favorable outcomes are revealed first, and as the public (wrongfully) believes that these outcomes are randomly selected from the true underlying distribution, their estimate on γ will only be revised upward.

This case however imposes an unrealistically high degree of naivety on voters: they think that the sequencing of reforms is truly random and do not take into account that the government (or nature, recall our discussion following Assumption 1 and footnote 15) starts with reforming the winners.

Consider therefore the more realistic case in which the public does take the selection bias into account. Then, the revelation of additional outcomes provides no valuable information since the public realizes that these draws do not come from the true underlying distribution, as a result of which Bayes' rule can no longer be applied.¹⁶ Consequently, agents cannot update their estimate of γ and

¹⁶After all, the only thing that is revealed if α_t is increased to α_{t+1} , is that $\gamma \geq \alpha_{t+1}$. However,

$\hat{\gamma}_t$ remains constant at $\hat{\gamma}_0$, $\forall t$ (where $\hat{\gamma}_0$ is the exogenously given belief on γ at the time of the start of the reform (which we refer to as "time 0")).¹⁷ This leads to the following results:

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

Proof. From equation (8) it follows that $\hat{\gamma}_0 \geq \frac{3}{4} - \frac{1}{2}\beta_0 \Leftrightarrow \alpha_0^* \geq \frac{1}{2}$. Revealing only winners (hence keeping β_t constant at β_0) implies that α_t^* remains constant at $\alpha_0^* \geq \frac{1}{2}$ over time. This implies the threshold $\alpha_t \geq \frac{1}{2} \forall t$, as a result of which the reforming government can reveal that the median voter is a reform winner before this voter starts opposing the reform (i.e. the government can push $\alpha_t \geq \frac{1}{2}$ before $\alpha_t \geq \alpha_t^*$). ■

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

Proof. From (8) it now follows that $\hat{\gamma}_0 < \frac{3}{4} - \frac{1}{2}\beta_0 \Leftrightarrow \alpha_0^* < \frac{1}{2}$. This implies that the reform is not believed to be "sufficiently efficient" (as defined in Proposition 2). Revealing only winners (hence keeping β_t constant at β_0) then decreases the expected value of the median voter via (3) and since constancy of β_t again implies constancy of α_t^* (at $\alpha_0^* < \frac{1}{2}$), $\alpha_t \geq \alpha_0^*$ before $\alpha_t \geq \frac{1}{2}$ and majority support is lost before the reform is completed. ■

Hence, those reforms on which beliefs are such that $\hat{\gamma}_0 \in (\frac{1}{2}, \frac{3}{4} - \frac{1}{2}\beta_0)$ can no longer be completed gradually by revealing only winners. Do note that this result arises even if the reform enjoys majority support at its start, and despite the fact that all of these reforms are believed to be efficient (in a sense that they are believed to generate more winners than losers, i.e. $\hat{\gamma}_0 > \frac{1}{2}$).

under Assumption 1 this is by no means informative on whether there are still any winners located beyond α_{t+1} .

¹⁷Similarly, α_0 (β_0) is the number of winners (losers) that has already been revealed at the start of the reform. This number can be equal to 0, but this does not necessarily have to be the case. It is perfectly possible that some firms already operate under the new regime before the reform starts (for example as a remnant of uncompleted past reform attempts).

The intuition for what is going on is exactly as in Section 3.1: because the public does not update its beliefs in response to the revelation of new reform outcomes (as it realizes that these outcomes do not carry any information since they are selected and do not reflect the true underlying distribution; cf. footnote 16), every additional winner revealed only reduces the perceived probability of ending up as a winner for those who remain uncertain. At some point, the median voter will therefore start opposing the reform that he used to support.

Similar dynamics will arise in the setup considered in the Appendix (where Assumption 1 is dropped and voters do apply Bayes' rule to update their estimate of the aggregate state) if the prior belief on the aggregate state is relatively tight. In that case, the updating process in the aggregate dimension will proceed at a rather slow pace, as a result of which the "sampling without replacement"-effect will get to dominate.

Moving back to the setup in which Assumption 1 does hold, revealing losers is no longer effective (in sharp contrast to the case without aggregate uncertainty).

Proposition 4 *If the public believes that the reform starts with reforming the ex post winners, reforms that are believed to be welfare-enhancing lose majority support as soon as a loser is revealed before $\alpha_t \geq \frac{1}{2}$.*

The proof simply follows from the fact that the public expects the government to start with revealing favorable reform outcomes. If an unfavorable reform outcome shows up, the public thinks that all winners have already been revealed and that those individuals who are still uncertain on their identity will all be losers. If this happens when $\alpha_t < \frac{1}{2}$, majority support is immediately lost.

1.3.3 Summarizing

The main lesson from Fernandez and Rodrik (1991) was already pretty dramatic: reforms that are welfare-enhancing may not enjoy majority support at their start because the reform winners cannot be identified upfront.

In a way, the message of this paper is even more depressing: even welfare-enhancing reforms that do enjoy majority support at their start (as a result of which they can get started along democratic lines), may not come to completion because of the learning dynamics that enter the story. In particular, by revealing

only winners the government suffers from the "sampling without replacement"-effect as a result of which it will at some point lose majority support if $\hat{\gamma}_0 < \frac{3}{4} - \frac{1}{2}\beta_0$, while revealing losers immediately ends support. The reforming government thus finds itself trapped and destined to lose majority support - irrespective of what action it takes.

1.4 How can the loss-of-support problems be avoided?

The depressing message from the previous section makes one wonder whether there is anything reformers can do to overcome the loss-of-support problems. In this respect, especially the Chinese reform experience suggests that there does exist a route towards 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 (described in Section 2 of this paper), the initial Chinese reform successes only seem to have increased support for further reforms (see Litwack and Qian (1998) and Qian, Roland and Xu (1999, 2006)). This raises the question why the experiences with gradualism have been so different across countries.

In this respect, however, it is crucial to note that the Chinese gradual reform strategy is quite different from the Latin American and Central and Eastern European approaches: while 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 introduced market forces in Shenzhen, Zhuhai, Shantou and Xiamen by designating those areas as Special Economic Zones in 1980 (soon to be accompanied by other regions, such as Hainan).

By reforming gradually along the spatial dimension, Chinese policy makers allowed the Chinese public to learn about the effects of the proposed reforms by looking at reform outcomes in these Special Economic Zones. Of course, voters will only find the information acquired via the spatial dimension useful if the Special Economic Zones are somewhat representative for the rest of the country. And it was in this respect that China had a major advantage over many other countries: as set out in Qian, Roland and Xu (2006, p. 394), the Chinese economy

is organized along territorial lines. This implies that all Chinese regions are relatively self-contained, in a sense that these regions can be seen as rough miniature versions of the Chinese economy as a whole. The former Soviet Union on the other hand was organized along industrial lines with high degrees of industrial concentration (mainly because of ideological reasons as argued in Qian and Xu (1993)), as a result of which each region was much more specialized, dependent of other regions, and less representative of the Union as a whole.

Consequently, China had (in contrast to the countries that emerged from the Soviet Union) the possibility to take more or less representative samples of the entire Chinese economy (in the form of certain regions), which they could then use to show voters where the gains and losses of the proposed reform are likely to occur - thereby reducing individual uncertainty (the root of all problems).¹⁸ Crucially, this approach does *not* suffer from the "sampling without replacement"-effect expressed by equation (3). The reason is that this is a form of sampling from a different, smaller urn (where the distribution of balls in the smaller urn was sampled randomly from the distribution of balls in the larger urn that represents the rest of the country). For this strategy to work, it is however crucial that agents who know that they will be among the winners (i.e. those with $i < \alpha_t$) cannot self-select into the Special Economic Zones, since that would imply that the zone becomes less representative of the country as a whole. Strikingly, this is exactly what the Chinese "hukou" system (which restricts the mobility of citizens within China) achieves. So although one could debate the fairness of this system (just like one could debate fairness of mobility restrictions between different countries), it does seem to play an economic role in the Chinese reform process.

Finally, this view on Special Economic Zones shines a new light upon 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 why governments bother installing them in the first place. In this respect, the present paper argues that Special Economic Zones could produce large *dynamic*

¹⁸In the former Soviet Union this strategy would not work: over there, reform outcomes in one region were less relevant to those in other regions due to the higher degree of spatial heterogeneity. Simultaneously, the higher degree of concentration would give any spatial reform strategy a sectoral flavor: after all, if certain sectors are concentrated in certain areas, reforming one area is equivalent to reforming one sector. Then, the "lack of replacement"-effect would enter the story.

gains, as they may facilitate the implementation of reforms that bring the entire country closer to the first-best.

Summarizing, we can thus conclude that when countries find themselves in the fortunate situation in which it is possible to experiment by first implementing the reform in a region that is somewhat representative for the rest of the country, gradualism along this spatial dimension may greatly facilitate the reform process. In this sense, there is an important difference between sectoral and spatial gradualism and this difference may be the key as to why the gradual reform strategy has worked for China, while this is much less so for many other countries.

1.5 Conclusion

This paper makes a previously overlooked observation, namely that the process of revealing reform outcomes is an example of sampling *without* replacement. In the presence of individual uncertainty, we show that this implies that the revelation of reform winners early on in the process, makes those who remain uncertain on their identity more pessimistic about their chances of ending up as a winner.

This "sampling without replacement"-effect challenges the standard view that sequencing should be such that favorable reform outcomes are revealed first and is consistent with the puzzling experiences that many reformers have had in practice: there are numerous examples of reforms that were started while enjoying majority support, but subsequently lost this support - even though they were progressing in a successful way.

Next to that, this paper has outlined a strategy that is able to overcome the problems related to the learning process. In particular, if a country happens to be organized along territorial lines (such that individual regions are rough miniature versions of the entire country), the "sampling without replacement"-effect can be avoided by reforming gradually along the *spatial* (rather than along the *sectoral*) dimension. This is the approach China has taken and implies that policy makers first introduce the reform proposal in certain regions of the country that are thought to be representative for the country as a whole. Subsequently, the public can look at reform outcomes in these regions to get an idea of where the gains and losses of the reform are going to be located. This enables the reforming government to reduce individual uncertainty (the root of all problems) without

running the risk of losing majority support due to the lack of replacement in the sampling procedure.

On a more general level, this paper has constructed a model in which agents are learning from realizations that are sampled without replacement. To the best of our knowledge, this has not been modeled before in the economics literature, despite the fact that many economic problems take this form (the process of revealing the identity of banks in a financial crisis (good or bad?) is a clear example). The model developed in this paper could therefore also be applied to different problems that entail sampling without replacement.

1.6 Appendix

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

Instead, here we will assume that the government cannot identify the reform winners and losers up front either. In addition, we will assume that there is no natural selection process present that could lead to a non-random sequencing of events. Consequently, reform outcomes are just sampled randomly from the true distribution. Moreover, we will also 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.

Since we feel that this case may deserve a closer study in its own right, we leave a full analysis for future work, but present some main results in this Appendix. In particular, we show here that the "sampling without replacement"-effect discussed in the main text continues to be at work in this alternative setting.

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

$$\alpha_t^* = \frac{2\alpha_t}{\alpha_t + \beta_t} + \beta_t - 1 \quad (\text{A1})$$

Using expression (A1), 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, one can verify that:

$$\frac{\partial (\alpha_t^* - \alpha_t)}{\partial \alpha_t} = \frac{2\beta_t}{(\alpha_t + \beta_t)^2} - 1 \quad (\text{A2})$$

Here, the first term shows that the revelation of winners pushes up α_t^* (since it leads to an upward revision of $\hat{\gamma}_t$ through application of Bayes' rule), while the second term ("−1") captures the fact that the revelation of winners simultaneously makes those who remain uncertain more pessimistic on their individual chances of

ending up as a winner. In particular, this term reflects the fact that revealing a reform outcome is an example of sampling without replacement. From (A2), one can derive that as long as

$$\sqrt{2\beta_t} - \beta_t < \alpha_t, \quad (\text{A3})$$

the "sampling without replacement"-effect dominates. Under condition (A3), the estimate of γ ($\hat{\gamma}_t$) 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 more and more *favorable* reform outcomes are 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 takes place without replacement.

Similarly:

$$\frac{\partial(\alpha_t^* - \alpha_t)}{\partial\beta_t} = \frac{-2\alpha_t}{(\alpha_t + \beta_t)^2} + 1, \quad (\text{A4})$$

captures the exact 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

$$\sqrt{2\alpha_t} - \alpha_t < \beta_t \quad (\text{A5})$$

Conditions (A3) and (A5) yield two non-overlapping 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, while the revelation of bad reform outcomes increases support for reform in region A5.

Now, one can ask: is the government able to complete the reform without a loss of support along the way? As the sequencing of the reform is random in this case (since 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. As we feel that this issue would deserve 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 then leads to two regions in the (α, β) -space where the support dynamics can be characterized as "anomalous".

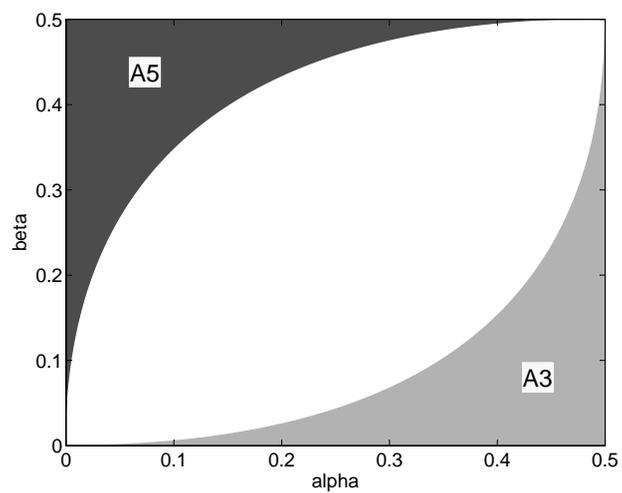


Figure 1.2: Region where favorable reform outcomes decrease support for the reform (A3) and vice versa (A5).