

Essays on Applied Microeconomic Theory

BY

David Andrés-Cerezo

European University Institute

APRIL 18, 2021

Contents

1	The Construction of National Identities	4
1.1	Introduction	4
1.2	Related literature	8
1.3	Model	10
1.4	Solution to the dynamic problem and main results	15
1.5	Non-interior steady state: A general result	19
1.6	Nation-building under electoral competition	23
1.7	Interpretation of the results: Case studies	28
1.8	Conclusion and ways forward	33
2	Storing Power: Market Structure Matters	36
2.1	Introduction	36
2.2	The Model	42
2.3	The Social Planner Solutions	44
2.3.1	The First-Best	44
2.3.2	The Second-Best	46
2.4	The Market Solutions	50
2.4.1	Competitive Storage	50
2.4.2	Independent Storage Monopolist	51
2.4.3	Vertically Integrated Storage Monopolist	55
2.5	Comparison across Market Structures	58
2.5.1	Non-Binding Storage Capacity	59
2.5.2	Binding Storage Capacity	61
2.6	Simulation of the Spanish Electricity Market	63
2.6.1	Scenarios	64
2.6.2	Results	65
2.7	Conclusions	70
3	Normative Analysis of Environmental Policy with Endogenous	

Preferences	73
3.1 Welfarist approach: preference satisfaction	78
3.2 Paternalism and substantive normative criteria	84
3.3 Non-consequentialist approaches	87
3.4 Normative analysis in a model of endogenous environmentalist preferences	89
3.4.1 Model set-up	89
3.4.2 Implications of different welfare criteria	93
3.5 Conclusion	100
 Appendix A	 102
 Appendix B	 156
 Bibliography	 178

Chapter 1

The Construction of National Identities

1.1 Introduction

During the past two centuries, nation-states have arisen in which the vast majority of the population shares a sense of collective belonging, represented by distinctive traditions, culture, language, and sometimes religion. By contrast, we have also observed cases of failed nation-building, in which governments were unable to create a unified national culture. Simultaneously observing such opposite outcomes naturally raises two questions: Which conditions lead to the promotion of a national identity? How do governments achieve this goal?¹

Providing answers to these questions is interesting for several reasons. First, many authors have identified a shared identity as a fundamental prerequisite to economic growth. The reason is that societies polarized along identity cleavages are prone to civil conflict and corruption, all of which are factors generally associated with poor economic performance (e.g., Besley and Reynal-Querol, 2014).² Second, the absence of a common national identity is a source

¹The analysis in this paper is only positive, and therefore we do not make any normative statement about the legitimacy of nation-building attempts or the desirability of having a unified national culture. These considerations depend on the particular context and on political philosophy issues beyond the scope of this paper. Our aim is solely to understand, albeit with many limitations, why some countries are divided along ethno-territorial identity cleavages while others are not.

²This problem is especially severe in Sub-Saharan Africa where, according to the Afrobarometer, around 50% of the population identifies more with their ethno-linguistic identity

of political instability, as is the case of countries characterized by a center-periphery tension such as Catalan and Basque regions in Spain, Scotland in the UK, Québec in Canada, Flanders in Belgium, Biafra in Nigeria, or Ogaden in Ethiopia. These regions are repeatedly involved in political processes claiming larger political autonomy, the recognition of their cultural distinctiveness, or even the formation of independent nation-states. Moreover, identity differences within a political union may prevent the provision of public goods, hinder redistribution, and lead to suboptimal economic policies (e.g., Guiso et al., 2016a), because they reduce social capital and increase the costs of collective action. Third, processes of nation-building are frequently related to political movements of state centralization and decentralization. In some cases, such as France at the end of the 19th century (Weber, 1976), the relatively fast creation of a unitary state was accompanied by the development of a common identity among citizens with different cultural backgrounds and a vague sense of national belonging. In other cases, such as the Soviet Union or Yugoslavia, the lack of state centralization within the political union coincided with the consolidation of opposing territorial identities (e.g., Bakke and Wibbels, 2006; Sekulic et al., 1994). In this regard, some authors suggest European leaders have relied on the expectation that a gradual political integration would promote the convergence of values and identities among the European population (Spolaore, 2013). However, these expectations have not been fulfilled yet (Guiso et al., 2016b), so whether we should expect the development of a European identity in the future remains an open question.

In the paper, we develop a theoretical model of state nation-building in a *peripheral* region where two identity groups co-exist: nationalists and regionalists. The cultural values of nationalists are aligned with those of the rest of the country, while regionalists have a culturally distinctive regional identity. Our framework incorporates two key ideas related to nation-building episodes. First, the ability of the central government to promote the national identity depends on its capacity to centralize the competences of the peripheral region. By having the power to tax and spend the revenues, the government can implement policies aimed at promoting certain cultural values.³ In our model, the main mechanism through which these policies operate is by changing the

than with the national identity.

³Nation-building policies include, for instance, the promotion of an official language, the introduction of a “collective memory” in school curricula, military conscription, or the provision of identity public goods such as patriotic celebrations commemorating important national events.

socialization incentives of the families, because parents want to educate their children to make them fit better in the region where they live. Second, centralization and nation-building policies usually come with episodes of political unrest in the form of protests or even civil war. The reason is that communities with a rooted identity feel aggrieved if their cultural and symbolic demands are not satisfied, or if an external government imposes a centralized state that disregards their cultural distinctiveness. Hence, by mobilizing in public demonstrations, protests, and riots, citizens can influence government policies through channels other than legal institutions.

Specifically, individuals in the peripheral region enjoy utility from consuming a public good associated with their identity (or an associated cultural trait, e.g., the regional language). The total amount of public goods is in fixed supply, giving rise to a zero-sum conflict, which allows us to capture cultural decentralization as the share of regionalist public goods. Identities are transmitted over generations through a process of cultural evolution with two channels of socialization: within the family (*vertical transmission*) and by society at large (*horizontal transmission*). As in standard models of cultural evolution (e.g., Bisin and Verdier, 2001), parents choose socialization efforts by weighting the benefits of transmitting their own identities against the potential costs for their children. This evolutionary process gives rise to a differential equation for the fraction of the population attached to the national identity, which is a function of the level of decentralization. A forward-looking government in the central region internalizes these dynamics and shapes them by modifying parents' socialization incentives. This government has a particular objective in mind, which for historical reasons is assumed to include the promotion of the national identity.⁴ However, the government is constrained by welfare losses created by political unrest, with protest participation levels increasing as nation-building attempts become stronger. Thus, the problem of the government becomes a non-linear optimal control problem over an infinite time horizon, in which a trade-off exists between its perceived benefits of having a more nationally-identified population and the costs created by political unrest.

In this setting, we begin by analyzing a problem that is linear in the control variable, which allows us to fully characterize the dynamics of nation-building episodes and to derive three key results. First, we find that optimal nation-

⁴The desire to promote a national identity could be motivated by pure ideological reasons, or by instrumental reasons such as processes of industrialization that require a population with a common language or the need of an army willing to defend the nation.

building policies involve quick and extreme interventions, so that homogenization is achieved in the fastest possible way.⁵ Concretely, optimal control trajectories follow a Most Rapid Approach Path (MRAP) (Spence and Starrett, 1975). Second, the optimal long-run steady states are characterized by cultural homogeneity within the peripheral region. Because identities are mutually exclusive, gains from one group always come at the expense of the other. Therefore, a forward-looking central government forced to solve this dispute will push homogenization to the maximum possible level, so that preferences in the population are aligned and the zero-sum conflict eventually disappears. Hence, in cases where promoting the full assimilation to the national identity is extremely costly, the government will grant full decentralization to the peripheral region and allow regional identities to thrive. Finally, the ability of governments to nation-build is highly dependent on the initial distribution of identities. In particular, we prove the existence of a population threshold above which the government promotes the national identity and vice versa. Small differences in economic, social, and political factors change the location of this threshold and may lead to extremely divergent trajectories, highlighting the contingent nature of nation-building processes. In addition, the model presents severe regime changes: an exogenous shock to the socioeconomic environment can trigger a sudden shift in the policies of the government. Therefore, our model can explain historical episodes in which states facing similar historical contexts followed disparate nation-building policies.

Then, we extend the benchmark model in several directions. First, we check how robust our results are to non-linear specifications of the objective function. The non-convex dynamics present in this problem prevent us from using standard optimal control techniques and hinder a full analytic characterization of the optimal path. However, we are able to identify general sufficient conditions under which the long-run distribution of identities remains homogeneous. Interestingly, these conditions are satisfied for certain cases in which parental socialization decisions are characterized by cultural substitution, in contrast to most results in the cultural-evolution literature (Bisin and Verdier, 2011). Second, we introduce democratic electoral competition for the central government. In particular, we characterize the Markov equilibrium of a differential game between two forward-looking political parties. Those two parties

⁵Throughout the text, we frequently refer to some policies or outcomes as optimal. This actually means they are optimal in a mere mathematical sense and from the point of view of a government that pursues its own goals, which may be different from a social welfare criterion.

have aligned nation-building motives but compete to win the benefits of being in office. Whenever the regionalist group is sufficiently pivotal, political parties face a trade-off between targeting this group to increase the probability of winning, and cooperating by proposing policy platforms that promote the national identity. Compared to the case of a secure government, we find that, everything else equal, the preservation of regional identities is more likely in democracies, because regional minorities have an additional instrument beyond protests to voice their demands.

The paper is organized as follows. Section 1.2 reviews the related literature. Section 1.3 describes the model. Section 1.4 includes the main theoretical results. Section 1.5 provides general sufficiency conditions for our homogenization results. Section 1.6 introduces electoral competition. In section 1.7 we discuss the results and we illustrate the predictions of the model with relevant case studies. Section 1.8 concludes. All proofs are relegated to Appendix A.1.⁶

1.2 Related literature

The main contribution of the paper is to formalize the conditions under which nation-building takes place and to characterize the dynamics of national identities. As argued by Smith (1992), “national identification has become the cultural and political norm, transcending other loyalties in scope and power.” Yet theoretical work on this topic is scarce.⁷ In the economics literature, Dessí (2008) studies the role of older generations in creating a shared identity by promoting a collective memory that praises the values and the history of the nation. Our work is complementary to hers. Whereas Dessí (2008) models the transmission of identities as a process of strategic communication of information between generations, in our model, identities correspond to different preferences. Moreover, we model its transmission as a cultural evolutionary process resulting from the interaction between vertical, oblique, and horizontal socialization.

Another set of theoretical contributions that explicitly deals with the forma-

⁶A supplementary online appendix, available at <https://bit.ly/2mtHtBa>, complements the paper.

⁷For an early discussion on nationalism from an economic point of view, see Breton et al. (1964).

tion of national identities emphasizes the role of interstate and modern mass warfare in triggering nation-building behavior by states. Alesina et al. (2017) explore how changes in warfare technology incentivized mass armies by conscription, which in turn pushed states to create a sense of national belonging that increased the willingness to fight for the country.⁸ Sambanis et al. (2015), building upon Shayo (2009) and Sambanis and Shayo (2013), argue that rulers who want to nation-build may find it optimal to promote war against an external power, as expectations of victory increase the international status of the country. We complement these works in two key respects. First, we depart from static models by introducing cultural dynamics that reflect more closely the behavior of nation-building policies over time. Dynamics allows us to make general predictions about the *process* by which national identities are formed. Second, these papers focus on the role of external wars and assume no internal opposition to nation-building policies. By contrast, our model allows for additional possible causes of nation-building and focuses on the domestic sources of political conflict and cultural resistance to these policies.

On the empirical side, some papers study the impact of different nation-building tools, such as language policy (Aspachs-Bracons et al., 2008; Caminal et al., 2018), school curricula (Cantoni et al., 2017; Fouka, 2019), compulsory schooling (Bandiera et al., 2015), mass media and propaganda (Voigtländer and Voth, 2015; Blouin and Mukand, 2019), public nationalist holidays (Madestam and Yanagizawa-Drott, 2012) and national football teams (Depetris-Chauvin et al., 2018). We contribute to this literature by developing a theoretical model that explores the long-run consequences of these policies and the political economy aspects that constrain their implementation.

Our paper also relates to a recent strand of the cultural-transmission literature that studies how different socialization agents shape preference formation, and how different modes of socialization compete with each other. Within this literature, our study is closely related to some papers analyzing how different “cultural leaders,” such as media or religious leaders, shape cultural evolution (Hauk and Immordino, 2014; Prummer and Siedlarek, 2017; Carvalho and Koyama, 2016). We contribute to this literature by studying the identity socialization role of a key cultural leader that has received little attention in

⁸In a closely related paper, Alesina and Reich (2013) analyze in a static setting the different incentives to promote a common identity that democracies and dictatorships face. Aghion et al. (2018) provides empirical support for the theory in Alesina et al. (2017), showing that external military threats encourage public investments in education that spur national identifications.

the literature: the nation-state.

We also connect to two strands of the literature that analyze the interplay between electoral competition and identity formation. On the one hand, within the cultural transmission literature, Bisin and Verdier (2000) and Tabellini (2008) study how majority voting maps the distribution of cultural traits into public policies, and how this mapping feeds back into socialization decisions. On the other hand, Shayo (2009) and Gennaioli and Tabellini (2018) analyze how individuals choose the strength of their national identification in the short-run in response to the relative salience of this dimension in the political realm as compared to the class-based dimension of conflict. We contribute to both strands of literature by introducing *perfectly forward-looking* political parties that purposefully shape individual identifications and fully internalize the impact of their actions on the future political power of different groups. As a result of not being purely opportunistic, they implement equilibrium policies that reflect the long-run goals of parties in addition to voters' policy preferences.

From a methodological point of view, our key contribution is to introduce a perfectly forward-looking socialization agent that solves a zero-sum conflict between cultural groups while internalizing cultural dynamics. Verdier and Zenou (2018) explore a similar idea in the context of the cultural assimilation of religious communities, although their paper includes no conflict between identities, and the leader chooses when to intervene. The presence of a zero-sum conflict makes our results fundamentally different from theirs: Whereas their optimal solution converges to a long-run outcome in which both cultural groups co-exist, our optimal steady states are homogeneous. Nevertheless, given the similar framework analyzed in both papers, some technical and conceptual similarities exist. For example, we both find that optimal trajectories with a linear per-period utility are characterized by an MRAP.

1.3 Model

Consider a continuous time model of a peripheral region of a country populated at time t by a stationary mass 1 of agents. Agents can be of two types: a fraction q_t of type N individuals, *nationalists*, and a fraction $1 - q_t$ of type R individuals, *regionalists*. In each period, all individuals receive identical income per capita normalized to 1. The government collects a fraction $r \in [0, 1]$ of

this income through taxes. The government uses total tax revenue r to provide identity (local) public goods $g_t^R = r\delta_t$ and $g_t^N = r(1 - \delta_t)$, where $\delta_t \in [0, 1]$ is the fraction used to provide public good g_t^R .

All individuals consume all of their after-tax income $c_t = (1 - r)$ deriving utility $f(1 - r)$, where f is continuous, increasing, and concave with $f(0) = 0$, and their associated identity public good, deriving utility g_t^i . Hence, total utility is given by $u_t^i(c_t, g_t^i, g_t^j) = f(c_t) + g_t^i$. Because $c_t = 1 - r$, $g_t^N = (1 - \delta_t)r$, and $g_t^R = \delta_t r$, indirect utilities are given by

$$U_t^N(r, \delta_t) = f(\underbrace{1 - r}_{c_t}) + \underbrace{(1 - \delta_t)r}_{g_t^N}; \quad U_t^R(r, \delta_t) = f(\underbrace{1 - r}_{c_t}) + \underbrace{\delta_t r}_{g_t^R}. \quad (1.1)$$

In this way, we capture national/regional identities as different preferences over mutually exclusive “identity” or cultural goods, that is, goods that can only be enjoyed by individuals with a particular religion, language or other cultural trait attached to the identity. Examples are national holidays commemorating a historical date, museums, monuments of past leaders of the nation, a national football team, and so on.⁹ In addition, the fact that the total amount of club public goods has constant total supply r in every period allows us to interpret δ_t , the fraction of the regionalist good, as the level of decentralization of this region.¹⁰

Cultural transmission

We endogenize preferences by including cultural dynamics following the cultural-transmission literature (Bisin and Verdier, 2001). Assume asexual reproduction where each parent has one child. Children are first exposed to parental (vertical) socialization, which, if unsuccessful, is followed by a random match

⁹We can also interpret them in a broader sense, as capturing the idea that some individuals within the peripheral region have different priorities and preferences with respect to government spending.

¹⁰In this view, decentralization corresponds to the *degree* of cultural recognition and accommodation of a differentiated nation within the boundaries of the state. This must be differentiated from fiscal or political decentralization, that requires the existence of a local government with fiscal and legislative powers. In any case, all these definitions of decentralization are highly correlated, as the provision of regionalist public goods is usually undertaken by regional governments with some political autonomy.

to an individual from the population, adopting her trait (horizontal socialization). Hence, transition probabilities are given by

$$\begin{aligned} P_t^{NN}(e_t^N) &= e_t^N + (1 - e_t^N)q_t & P_t^{NR}(e_t^N) &= (1 - e_t^N)(1 - q_t) \\ P_t^{RR}(e_t^R) &= e_t^R + (1 - e_t^R)(1 - q_t) & P_t^{RN}(e_t^R) &= (1 - e_t^R)q_t, \end{aligned} \quad (1.2)$$

where P^{ij} is the probability that a child of a parent with trait i is socialized to trait j , and $e_t^N \in [0, 1]$ and $e_t^R \in [0, 1]$ are parents' education/socialization efforts. Notice that the more present an identity trait is, the more likely agents in the young generation are to adopt it.

Parents take into account how rewarding is to have each identity in society, and based on that they choose how much effort to put in transmitting their own identities. More concretely, let V^{ij} be the utility that a type i parent derives from having a child with trait j , and let $C(e_t^i)$ be the socialization cost, assumed to be increasing and convex. Then parent i 's socialization problem at time t is

$$\max_{e_t^i \in [0,1]} P_t^{ii}(e_t^i)V_t^{ii} + (1 - P_t^{ii}(e_t^i))V_t^{ij} - C(e_t^i). \quad (1.3)$$

Under the assumption that parents' choices display *imperfect empathy*, socialization utilities are given by $V^{ii} = c_t + g_t^i$ and $V^{ij} = c_t$.¹¹ The following lemma characterizes optimal socialization decisions and the corresponding law of motion for identity dynamics:

Lemma 1 *Under imperfect empathy and quadratic costs*

$$C(e_t) = \frac{1}{2}e_t^2,$$

the optimal socialization efforts are given by

$$\begin{aligned} e_t^N &= (1 - q_t)g_t^N = (1 - q_t)(1 - \delta_t)r \\ e_t^R &= q_t g_t^R = q_t \delta_t r, \end{aligned} \quad (1.4)$$

and the law of motion for cultural transmission becomes

$$\dot{q} = q_t(1 - q_t)(e_t^N - e_t^R) = r q_t(1 - q_t)(1 - \delta_t - q_t). \quad (1.5)$$

¹¹As is standard in the literature, this assumption implies parents evaluate their children's utility using their own utility function.

The proof for this lemma is mechanical and can be found in Section A of the Supplementary Appendix. On the one hand, the parents' optimal choice of socialization effort takes into account how much welfare their children derive from holding their own identity, which depends on the provision of its associated public good. On the other hand, because horizontal transmission is a substitute for vertical transmission, parents' effort will decrease in the size of the group holding their identity in the population at large.

Observe that under cultural substitution and constant government intervention over time $\delta_t = \delta \in (0, 1)$, three steady states exist with a unique stable and interior steady state given by $q^{SS} = 1 - \delta$. Hence, our model preserves the standard prediction of the cultural-evolution literature of a heterogeneous steady state in which both identities co-exist. The main difference in our analysis is, precisely, that the government tailors dynamics at its own will by choosing a path for δ over time.¹²

Government dynamic problem

We assume the central government has “*de jure*” power to decide over δ_t and internalizes the cultural-transmission dynamics. The objective of the government is to choose a path $\{\delta_t\}_{t \geq 0}$ that maximizes:

$$\begin{aligned} \max_{\delta_t \in [0,1] \forall t \geq 0} \int_0^\infty e^{-\rho t} W(\delta_t, q_t; \omega) dt \\ \text{s.t. } \dot{q}_t = r q_t (1 - q_t) (1 - \delta_t - q_t) \\ q(0) = q_0, q_t \in [0, 1], \end{aligned} \tag{1.6}$$

where we define the parameter $\omega \in \Omega$ as the vector of all pertinent parameters in the model.

Assume the government has the following flow utility:

¹²We have two reasons to think that in the absence of government intervention, long-run heterogeneity will exist. First, taxes may not be collected or no identity public good may be provided, so that $e_t^N = e_t^R = 0$ and q_t remains constant. Second, both identity public goods g^N and g^R may be provided in a decentralized way by each of the groups. This reasoning could explain why both nation-building attempts and homogeneous nations have appeared recently in history, when nation-states have obtained coercive power and sufficient capacity to tax the people in all regions. The idea that fiscal capacity is a key determinant of the ability to nation-build can be found in Johnson (2015).

$$\begin{aligned}
 W(q_t, \delta_t) = & \underbrace{\psi^N q_t}_N + \underbrace{\psi^U [\alpha q_t U^N(\delta_t) + (1 - \alpha)(1 - q_t) U^R(\delta_t)]}_W \\
 & - \underbrace{\psi^S [\beta D_t^N(\delta_t, q_t) + (1 - \beta) D_t^R(\delta_t, q_t)]}_L, \tag{1.7}
 \end{aligned}$$

with three different goals:

N “Nation-building” motive: ψ^N captures factors that change the incentives to nation-build. Assume $\psi^N \geq 0$, so the government is biased toward the national identity.¹³

W “Welfare” motive: ψ^U captures how much the government cares about the utilities of individuals. The central government values utilities asymmetrically, so $\alpha \in (0, 1)$ represents the government’s weight on nationalists’ welfare.

L “Law and order” motive: ψ^S captures the loss in welfare created by protests $D_t^N(\delta_t, q_t)$ and $D_t^R(\delta_t, q_t)$. $\beta \in (0, 1)$ represents the government’s weight on nationalists’ protests. For now, assume the participation rates in protests, D^R and D^N , are given by¹⁴

$$D_t^N(\delta_t, q_t) = q_t \delta_t r \qquad D_t^R(\delta_t, q_t) = (1 - q_t)(1 - \delta_t)r.$$

Following Passarelli and Tabellini (2017), these participation rates in protests depend on the emotional reward for the individual of defending his group identity, and they are increasing in the distance between the policy implemented and the policy they deem fair.¹⁵

For simplicity, we normalize $\psi^U = \psi^S = 1$.

¹³All results carry through if $\psi^N \leq 0$.

¹⁴These participation rates could also be interpreted as the probability that political unrest reaches some threshold of, for example, secessionist attempts. Moreover, the parameter β captures how organized protesters are relative to the other group, or the relative capacity of group leaders to mobilize people.

¹⁵See Supplementary Appendix A.2.2 for the micro-foundations for the participation rate in protests. For the moment, note that throughout the paper we assume that individuals are entirely selfish with respect to the policy they feel entitled to. However, as discussed in A, the results of the paper are robust to situations in which the policy that individuals deem fair takes into account the size of each of the groups.

The final problem that the government solves is given by

$$\begin{aligned}
 \max_{\delta_t \in [0,1] \forall t \geq 0} \int_0^\infty e^{-\rho t} & \left\{ \psi^N q_t + \alpha q_t (f(1-r)(1-\delta_t)r) + (1-\alpha)(1-q_t)(f(1-r) + \delta_t r) \right. \\
 & \left. - r(\beta q_t \delta_t + (1-\beta)(1-q_t)(1-\delta_t)) \right\} dt \quad (1.8) \\
 \text{s.t. } \dot{q}_t &= r q_t (1-q_t)(1-\delta_t - q_t) \\
 q(0) &= q_0, q_t \in [0, 1].
 \end{aligned}$$

Problem 1.8 captures two key trade-offs of the government. First, a *static* constant-sum conflict on how to split the budget between the two types of public goods exists. Second, the government faces an *inter-temporal* trade-off when deciding whether to promote a common identity among a culturally diverse population. To see this dynamic trade-off, consider a situation in which the government only cares about the utility and size of the nationalist group, and only the protests of regionalist individuals create a loss of welfare. In this case, the government internalizes that increasing δ_t today in order to reduce the level of protests of the regionalist group also reduces q_t through the cultural-transmission mechanism, which in turn, will increase demand for more decentralization in the future, adding further pressure to set higher values of δ_t and making nation-building more difficult.

1.4 Solution to the dynamic problem and main results

The logistic equation shaping the evolution of the state variable in problem 1.8 prevents us from using standard optimal control techniques. Concretely, the first-order conditions of the Maximum Principle are not sufficient to fully characterize the dynamics, because the Hamiltonian is not jointly concave in δ and q . However, it can be shown that the optimal trajectory is characterized as an MRAP, which can be proved following Spence and Starrett (1975).

Proposition 1 *The optimal path for problem (1.8) is an MRAP. That is, the optimal solution approaches as fast as possible a steady state in which per-period welfare is maximized.*

The result that the government will approach a steady state where welfare is maximized at each point in time is intuitive: Staying where the highest benefits

are delivered is optimal. Moreover, Proposition 1 implies optimal paths do not involve a smooth approach to the steady state.

Theorem 1 *For any value of $\omega \in \Omega$ and any initial condition $q_0 \in (0, 1)$, optimal policies set either $\delta = 0$ or $\delta = 1$ forever with no switch in policies.*

Furthermore, no interior steady state exists so optimal paths approach one of the extreme stationary points of the state variable as fast as possible:

$$\lim_{t \rightarrow \infty} q_t^* = 0 \qquad \text{or} \qquad \lim_{t \rightarrow \infty} q_t^* = 1.$$

where q_t^* is the path of the state variable under the optimal policy δ^* .

Two important insights can be derived from the previous Theorem 1. First, the optimal stationary states in our model are culturally homogeneous. Second, the results above also imply optimal homogenization policies should be fast and intense, so any $\delta_t \in (0, 1)$ is sub-optimal, both along the optimal trajectory and in the steady state. The intuition for these results is simple: If at some point increasing q is optimal, further increasing it at the next instant must also be optimal, because fewer people engage in political unrest while more people enjoy the benefits of the public good. Eventually, full homogenization is optimal, so protests are minimized and the conflict about how to split the tax revenues fully vanishes. This result is driven by the constant-sum distributive conflict: A larger provision of one public good always comes at the expense of a reduction in the other good. Therefore, intermediate solutions for δ_t are sub-optimal because they imply investing in opposing goals: homogenizing toward N while homogenizing toward R .

Initial conditions

So far, we have established that the only steady-state candidates are $q = 0$ or $q = 1$. The next result characterizes, given initial conditions, which of those is the long-run optimum:

Theorem 2 *For any parameter values, a unique \bar{q}_0 exists such that the government is indifferent between setting $\delta^* = 0$ or $\delta^* = 1$ forever. That is, the*

optimal policy is characterized by threshold \bar{q}_0 as follows:

$$\delta^*(q) = \begin{cases} 1 & \text{if } q \leq \bar{q}_0 \\ 0 & \text{if } q \geq \bar{q}_0. \end{cases}$$

Theorem 2 implies that whether nation-building takes place depends on the initial distribution of preferences: When the national identity is held by a sufficient majority, the short-run costs of regionalist protests are relatively low. As a consequence, the government finds optimal to incur such costs for some time to obtain the long-run benefits of having a population fully homogenized to the national identity. Interestingly, when promoting the national identity is not worthwhile, the central government refrains from preserving a small nationalist group within the peripheral region and allows the regional identity to thrive, giving rise to a *multinational* state.

Our model sheds light on the initial question of why some distinctive regional identities persist within some countries. One of our predictions is that regional identities persist if the government finds that granting decentralization today is optimal if, for example, the welfare losses caused by *regionalist* political opposition are significantly large. By doing so, the demand for decentralization increases over time because *regionalist* parents socialize their kids to the cultural traits attached to the regional identity. By contrast, the group with an attachment to the identity of the central region observes that the policies of the central government do not represent their preferences, and they refrain from transmitting the national identity. Therefore, relatively strong political opposition to nation-building policies at early stages can prevent the development of a national identity.¹⁶

Proposition 2 *The threshold \bar{q}_0 is decreasing in ψ^N , α , and β .*

The proofs for Propositions 2, 3 and 4 are mostly based on algebra and can be found in Section A.2.3 of the Supplementary Appendix. We can see that a higher ψ^N implies that the benefits from increasing the size of the nationalist group are higher. Hence, everything else equal, an increase in ψ^N makes the government nation-build for a larger set of initial identity distributions. As expected, the incentives to nation-build change in the same direction when the weight attached to the utility of nationalist individuals (α) increases or when

¹⁶An alternative explanation for the survival of peripheral regional identities is a failure of the central government to fully internalize dynamics, either because of pure myopia or because it does not operate over an infinite time horizon.

the harm inflicted by political unrest of the nationalist group (β) is larger.

The comparative statics of \bar{q}_0 in ρ and r are significantly more complex. Furthermore, the sign of $\frac{\partial}{\partial \rho} \bar{q}_0$ depends on the other parameters of the model. The following proposition holds:

Proposition 3 *The comparative statics on ρ can go both ways:*

- *If α or ψ^N are large enough, then*

$$\frac{\partial}{\partial \rho} \bar{q}_0 > 0.$$

- *On the contrary, if α and ψ^N are small enough, then*

$$\frac{\partial}{\partial \rho} \bar{q}_0 < 0.$$

To understand the previous result, note that the long-run *differential* returns between setting $\delta = 1$ and $\delta = 0$ are decreasing in ψ^N and α , being negative for sufficiently large values of these parameters. Recall that \bar{q}_0 is the initial point for which the government is indifferent between setting $\delta = 1$ or $\delta = 0$ forever. Therefore, starting at \bar{q}_0 , for sufficiently high values of ψ^N and α , the differential returns between $\delta = 1$ and $\delta = 0$ are necessarily positive in the short-run. In other words, the government has a relatively strong desire to nation-build in the long-run but faces relatively high returns of setting $\delta = 1$ in the short-run. Therefore, when this government is more impatient, the short-run returns become more important and, as a consequence, fewer initial distributions of identities exist for which the government finds nation-building profitable (\bar{q}_0 increases). This last result highlights the fact that conditional on being sufficiently interested in nation-building, more *stable* governments (interpreted as smaller ρ) are more likely to develop a widespread national culture.

Proposition 4 *The comparative statics on r can go both ways and depend on the other parameters of the model:*

- *For small α , and sufficiently large ψ^N , it follows*

$$\frac{\partial}{\partial r} \bar{q}_0 < 0.$$

- On the other hand, for large α , and sufficiently small ψ^N , it follows

$$\frac{\partial}{\partial r} \bar{q}_0 > 0.$$

The previous proposition captures the fact that ρ and r play opposite roles in our model: An increase in r makes dynamics in any direction faster, so it is effectively equal to moving any future point closer to the present, or equivalently, putting more weight on the future. Hence, an increase in r can also be seen as a decrease in ρ . Finally, a second effect arises from the government per-period utility flow, but for sufficiently large ψ^N , the first effect dominates the second.

Therefore, everything else equal and conditional on having sufficiently large nation-building motives, countries in which the government has a greater ability to tax are more likely to have a shared national identity. The reason is that the government can implement stronger nation-building policies, which changes q faster and makes episodes of political unrest less prolonged.

1.5 Non-interior steady state: A general result

In our previous analysis, both the government's per-period utility flow and the law of motion were linear functions in the control δ . This linearity assumption allowed us to characterize our optimal control policy as an MRAP solution. In this section, we first provide a general result that holds for several specifications of the objective function, including non-linear functional forms. Then, we analyze the robustness of all the results of the baseline case when considering a per-period objective function that is quadratic in δ .

The following result gives sufficiency conditions under which no interior steady state exists:

Theorem 3 *Assume $\delta^*(q)$ is a solution to the following optimal control problem:*

$$\begin{aligned} & \max_{\delta_t \in \Delta} \int_0^{\infty} e^{-\rho t} W(q_t, \delta_t) dt \\ & \text{s.t.} \quad \dot{q}_t = g(\delta_t, q_t) \text{ and } q_0 = q. \end{aligned} \tag{1.9}$$

Denote by $\delta^S(q)$ the stationary policy function

$$\dot{q} = g(q, \delta^S(q)) = 0,$$

and define function

$$H(q) = W(q, \delta^S(q)).$$

If for some interior \tilde{q} , we have $\delta^*(\tilde{q}) = \delta^S(\tilde{q}) \in \Delta^\circ$, such that $\{g(\tilde{q}, \delta) | \delta \in [0, 1]\}$ is an open neighborhood of 0, then \tilde{q} is a local maximum of $H(q)$.

The intuition of the previous theorem is simple: If the per-period utility $H(q)$, derived from the policy $\delta^S(q)$ that keeps q unchanged ($\dot{q}(q, \delta^S(q)) = 0$), can be improved in some feasible direction at certain \tilde{q} , then staying at \tilde{q} cannot be optimal, because we can construct an alternative path delivering a higher discounted payoff. In other words, \tilde{q} cannot be an interior steady state.

Corollary 4 *If the function $H(q) = W(q, \delta^S(q))$ has no local maximum over all feasible values of q , the optimal path does not have an interior steady state.*

The previous results deliver sufficiency conditions for long-run cultural homogeneity. It is easy to check that for our equation describing cultural dynamics, we have that for any $q \in (0, 1)$, $\{g(q, \delta) | \delta \in [0, 1]\}$ is an open neighborhood of 0 and $\delta^S(q) = 1 - q$ is always feasible i.e. the government can always steer dynamics in its desired direction.¹⁷ Hence, it is enough to check whether the function

$$H(q) = W(q, 1 - q)$$

has a local maximum in $(0, 1)$.

The conditions for Theorem 3 apply to several objective functions of the government. In what follows, we illustrate its strength with two particular cases.

First, in the linear case of the baseline model, we have that

$$H(q) = W(q, 1 - q) = \psi^N q + (\alpha q + (1 - \alpha)(1 - q))f(1 - r) + r(\alpha q^2 + (1 - \alpha)(1 - q)^2 - q(1 - q)),$$

which is strictly convex in q because $H''(q) = 4r > 0$ and therefore does not have any local maxima in $q \in (0, 1)$. Hence, by Theorem 3 and its corollary,

¹⁷This condition does not hold in models where socialization efforts are always strictly positive. This would be the case, for instance, if we assume that there are private rewards of having an identity beyond the consumption of the associated public good. Under such conditions, it may not be possible to reach $q = 0$ or $q = 1$. However, the government still finds it optimal to provide only one type of public good and to achieve the maximum *feasible* level of homogenization.

no interior steady state exists under the optimal policy, a result already shown in Theorem 1.

Second, we can also show that there are no interior steady states when protests are given by a convex quadratic cost, that is, when the government solves the following problem:¹⁸

$$\begin{aligned} \max_{\delta_t \in [0,1]} \int_0^\infty e^{-\rho t} & \left(\psi^N q + \alpha q_t (f(1-r) + r(1-\delta_t)) + (1-\alpha)(1-q_t)((f(1-r) + r\delta_t) \right. \\ & \left. - r^2(\beta q_t \delta_t^2 + (1-\beta)(1-q_t)(1-\delta_t)^2) \right) dt \quad (1.10) \\ \text{s.t.} \quad \dot{q}_t & = r q_t (1-q_t)(1-\delta_t - q_t). \end{aligned}$$

Proposition 5 *If the government solves problem 1.10, no interior steady state exists. Therefore, under the optimal policy function,*

$$\lim_{t \rightarrow \infty} q_t = 0 \quad \text{or} \quad \lim_{t \rightarrow \infty} q_t = 1.$$

As in the linear case, the key result is that central governments will pursue homogenization to the maximum possible, because doing so minimizes long-run political unrest while maximizing the long-run benefits of having an homogeneous population. This result is robust across different specifications because its main driving force is the zero-sum nature of the conflict, in the sense that the gains for one group always come at the expense of the other. Intuitively, a heterogeneous steady state cannot be optimal, because in the long-run the central government still faces a conflict, which can be eliminated by further homogenizing the population. The previous argument goes through whenever the government faces a strong enough conflict on how to allocate finite resources between different groups.¹⁹

Finally, we can also prove the optimal policy and long-run dynamics also preserve the threshold property with intervals of fast and extreme interventions for largely homogeneous populations.

¹⁸See Supplementary Appendix A.2.2 for micro-foundations of this functional form of protests.

¹⁹Sufficient conditions for the latter are that citizens' valuations of any level of provision of the public good are large enough and participation rates in protests do not explode for low levels of public-good provision.

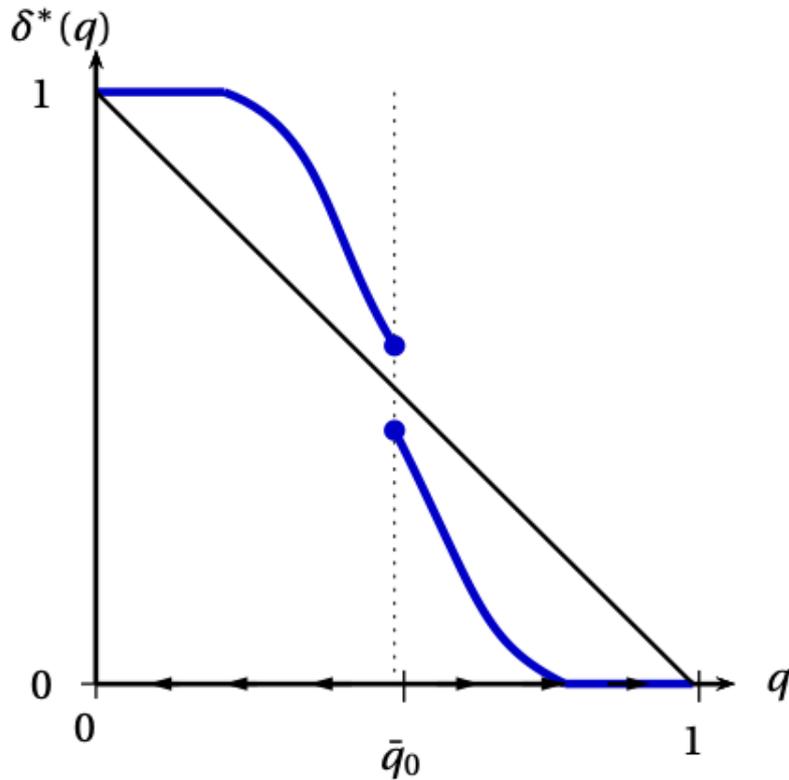
Theorem 5 A $\bar{q}_0 \in (0, 1)$ exists such that

$$\delta^*(q) > 1 - q \quad \text{if } q \leq \bar{q}_0, \quad \delta^*(q) < 1 - q \quad \text{if } q \geq \bar{q}_0.$$

Moreover, $\delta^*(q)$ is continuous on $[0, \bar{q}_0) \cup (\bar{q}_0, 1]$, and two open neighborhoods of $q = 0$ and $q = 1$ exist, say, $\mathcal{O}(0)$ and $\mathcal{O}(1)$ in $[0, 1]$, such that

$$\delta^*(q) = 1 \quad \forall q \in \mathcal{O}(0), \quad \delta^*(q) = 0 \quad \forall q \in \mathcal{O}(1).$$

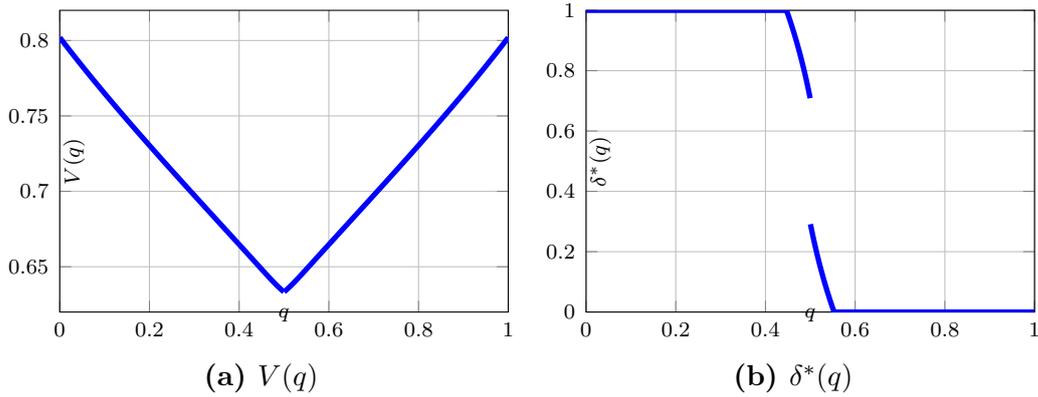
Figure 1.1 – Optimal policy $\delta^*(q)$ with quadratic protests



Unfortunately, a full analytic characterization of the optimal policy $\delta^*(q)$ becomes intractable. However, below we present a numerical example. We follow the approach presented in Achdou et al. (2017), which finds the value function using viscosity solutions. Observe that standard methods do not apply in our model, because our value function is not differentiable at the threshold $\bar{q}_0 \in (0, 1)$.

To summarize, the results of extreme homogenization and the existence of an indifference threshold go through beyond the linear specification. Regarding the optimal path for δ_t , the numerical results suggest (de)centralization becomes more gradual when protests are convex. To see why, consider the case

Parameter values: $\psi^N = 0, \alpha = 0.5, \beta = 0.5, r = 0.3, \theta = 0.3, \rho = 0.5$



in which the government starts with a balanced distribution of identities and eventually homogenizes the population toward the national identity. At intermediate values of q_t , the rate of change \dot{q} is higher and the size of protests bigger than for more homogeneous distributions of identities. Therefore, by setting an interior value for δ_t at early stages, the central government can avoid a large participation in protests of the regionalist group and still move in the desired direction. Interestingly, even with non-linear specifications, the optimal policy rapidly approaches corner solutions, suggesting the MRAP solution could be a good approximation of optimal nation-building policies. More importantly, although δ_t could take intermediate values along the transition path to the steady state, it eventually reaches a corner solution. Intuitively, an interior long-run value for δ_t cannot be optimal because, due to the identity dynamics, the government will be “investing” in two opposite goals at the same time.

1.6 Nation-building under electoral competition

In this section, we compare the results of our baseline model of a secure central government with the case in which the central government is democratically elected each period by citizens in the central and peripheral regions. For ease of exposition, we consider the limiting case in which the peripheral region is fully pivotal in national elections, so that only voters of groups N and R determine the result of elections. Nonetheless, in Supplementary Appendix A, we explicitly model voters in the central region and we show the results are robust.

We follow the probabilistic voting model with majority voting and aggregate uncertainty proposed by Persson and Tabellini (2000) based on Lindbeck and Weibull (1987). In this model, two parties A and B compete to win elections in every period by making simultaneous policy announcements δ^A and δ^B . Political parties commit to implement their announced policies if they happen to be elected. When announcements are δ^A and δ^B , party A 's probability of being elected is given by

$$p^A(\delta^A, \delta^B, q) = \frac{1}{2} + \frac{(1-q)\phi^R - q\phi^N}{q\phi^N + (1-q)\phi^R}(\delta^A - \delta^B) = \frac{1}{2} + \Phi(q)(\delta^A - \delta^B).$$

In Supplementary Appendix A, we provide an explicit microfoundation of $p^i(\delta^i, \delta^{-i}, q)$. For the moment, note that ϕ^i captures the intensity of preferences toward policy δ of group i .²⁰ By definition, the probability of winning the election for party B is $p^B = 1 - p^A$.

We assume political parties are forward-looking, maximize a discounted stream of utility payoffs, and internalize the dynamics of identities. They have an intrinsic nation-building motive, $\psi^N q$, with $\psi^N \geq 0$ equal for both parties, as well as office motivations, receiving per-period utility equal to p^i if they win the elections.²¹ Hence, when the proportion of nationalist is given by q , for given announcements δ^i and δ^{-i} , the per-period utility for party i is:²²

$$W^i(q, \delta^i, \delta^{-i}) = \psi^N q + p^i(\delta^i, \delta^{-i}, q).$$

We restrict our attention to Markov perfect equilibria, where strategies only depend on the current state q . The problem of player i is to choose a policy announcement δ^i taking the strategy of the other player, δ^{-i} , as given.

²⁰ ϕ^i is a measure of how much individuals within a group are concerned with the cultural policy/territorial cleavage of policy, relative to other policy dimensions or to intrinsic preferences toward political parties.

²¹The results hold when the strength of nation-building motives is different for both parties, as long as both have strictly positive nation-building motives.

²²Compared to the benchmark case, we have assumed $\psi^U = \psi^S = 0$. This choice of specification is made for tractability purposes. However, if political parties have “welfare” and “law and order motives,” as in the case of the secure government, similar results go through. In the benchmark model the benefits of holding office for the secure government enter as a constant (which we omitted) and therefore they do not alter the results of that section.

Equilibrium strategies are characterized by

$$\delta^{*i} = \arg \max_{\delta \in [0,1]^{[0,1]}} \left\{ \mathbb{E}_0 \int_0^\infty e^{-\rho t} W^i(q_t, \delta(q_t), \delta^{*-i}(q_t)) dt \right\} \quad (1.11)$$

$$\delta^{*-i} = \arg \max_{\delta \in [0,1]^{[0,1]}} \left\{ \mathbb{E}_0 \int_0^\infty e^{-\rho t} W^{-i}(q_t, \delta^{*i}(q_t), \delta(q_t)) dt \right\}, \quad (1.12)$$

where \mathbb{E}_0 is the expectation conditional on q_0 . Parental socialization decisions each period are made after elections have taken place and depend on the implemented policy. Hence, from the point of view of political parties, the value of \dot{q} is a random variable whose realization depends on the policy implemented by the winning party (as the result of elections is also a random variable). Hence, problems 1.11 and 1.12 are subject to the following dynamics:

$$\dot{q}_t = g(q_t, \delta^i, \delta^{-i}) = \begin{cases} q_t(1 - q_t)(1 - \delta^i(q_t) - q_t) & \text{with prob. } p^i(\delta^i(q_t), \delta^{-i}(q_t), q_t) \\ q_t(1 - q_t)(1 - \delta^{-i}(q_t) - q_t) & \text{with prob. } 1 - p_t^i(\delta^i(q_t), \delta^{-i}(q_t), q_t). \end{cases}$$

Note that for low values of q_t , electoral and nation-building motives are not aligned. Therefore, taking what the other party does as given, party i faces a trade-off between increasing the probability of winning elections by announcing a policy that favors the regionalist group R , or announcing a less popular policy today that increases q in the future. In this last case, party i faces the cost of reducing the expected benefits from office as well as the probability that this nation-building policy is implemented.

Given that two parties solve identical problems, we restrict our attention to symmetric equilibria.²³ The solution of the electoral-competition game with nation-building motives is characterized by the following theorem:

Theorem 6 *A unique equilibrium in symmetric strategies of the dynamic electoral-competition game with nation-building motives exists. The equilibrium strategies are described as a threshold policy given by \tilde{q}_D such that*

$$\delta^{A*}(q) = \delta^{B*}(q) = \begin{cases} 1 & \text{if } q \leq \tilde{q}_D \\ 0 & \text{if } q > \tilde{q}_D, \end{cases}$$

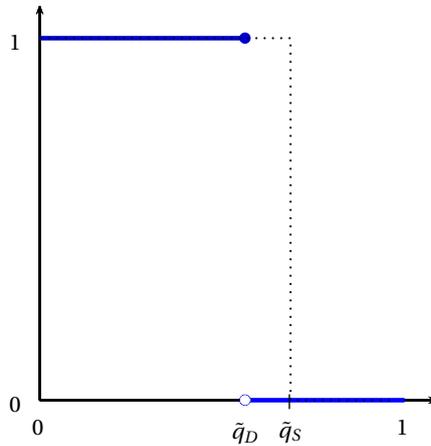
with $0 < \tilde{q}_D < \tilde{q}_S$, where \tilde{q}_S defines the threshold of the symmetric equilibrium for the static electoral competition game, and is given by

²³An interesting extension that merits a paper on its own would be to analyze the dynamic electoral game between forward-looking parties with opposite nation-building motives. In the Supplementary Appendix A we discuss this possibility as well as the case of shortsighted but ideologically motivated parties.

$$\tilde{q}_S = \frac{\phi^R}{\phi^N + \phi^R}.$$

As in previous cases, depending on the initial value of q , the system converges to one of the two homogeneous steady states. The intrinsic cultural-substitution properties of parental socialization are counteracted by the fact that it is better to be part of the most powerful group. Even if, *ceteris paribus*, parents exert more socialization effort when their group becomes smaller, the voting system favors individuals who belong to the majority, creating incentives for parents to socialize their kids to the predominant identity.

Figure 1.3 – Equilibrium policy of electoral-competition game with mixed motives



Proposition 6 *In the case in which parties only care about office ($\psi^N = 0$):*

$$\tilde{q}_D = \tilde{q}_S.$$

On the contrary, if parties only care about nation-building, ($O = 0$):

$$\tilde{q}_D = 0.$$

Note that when parties are purely office motivated, in every period they play the equilibrium strategy of the *static* game.²⁴ Starting below the threshold \tilde{q}_S , pure electoral reasons push parties to implement full decentralization $\delta = 1$, because this group is politically more powerful due to the combination of its

²⁴See Supplementary Appendix A for more details about our definition of a static political equilibrium.

demographic weight and the intensity of preferences of the group toward policy δ_t . However, when both parties also have a nation-building motive, a region between \tilde{q}_D and \tilde{q}_S exists in which parties are not constrained to target announcements to the powerful regionalist group and have some margin to implement policies that favor nationalist voters. The reason is that the nationalist group is big enough to guarantee a sufficiently large ex-ante probability of winning the elections. In equilibrium, both parties announce $\delta = 0$ and, due to the identity dynamics, the demographic weight q_t increases endogenously over time until \tilde{q}_s is reached and nation-building and electoral goals become aligned.

Finally, these results also suggest that, everything else equal, democracies tend to be more prone to accommodate and preserve regional identities than countries in which political power is not disputed. For example, consider the baseline model with $\psi^U = \psi^S = 0$, so per-period utility of the central government is given by $W(q_t) = \psi^n q_t + 1$, where the last term is the benefit of holding office obtained with certainty as no elections take place. Clearly, the optimal solution for the secure central government is to set $\delta(q) = 0$ for all q , so nation-building will take place for any initial $q_0 \in (0, 1]$. By contrast, when the central government is elected democratically, a region always exists in which equilibrium strategies are given by full decentralization, $\delta = 1$. Succinctly, the preservation of regional identities is more likely under electoral competition, because regional minorities have more voice in the political process. Moreover, the region where equilibrium strategies are $\delta^i = 1$ is increasing in ϕ^R , so democratic systems with powerful and ideologically motivated regional minorities are more likely to preserve their identities.

Finally, the following proposition characterizes how incentives to nation-build change with the parameters of the model:

Proposition 7 *The threshold \tilde{q}_D is decreasing in ψ^N :*

$$\frac{\partial}{\partial \psi^N} \tilde{q}_D \leq 0,$$

with limiting cases

$$\lim_{\psi^N \rightarrow 0} \tilde{q}_D = \tilde{q}_S, \quad \lim_{\psi^N \rightarrow \infty} \tilde{q}_D = 0.$$

On the contrary, \tilde{q}_D is increasing in ρ :

$$\frac{\partial}{\partial \rho} \tilde{q}_D \geq 0.$$

The proof can be found in Section A of the Supplementary Appendix. It is easy to see that when the incentives to nation-build are larger, nation-building will occur for more initial states. Similarly, when parties are less patient, they are more concerned about short-run electoral goals, and therefore nation-building will occur for fewer initial states. Moreover, as in the baseline model, small differences in the parameters or in the initial size of groups may have a large impact on the dynamics of identities.

1.7 Interpretation of the results: Case studies

In this section, we discuss the main predictions of the model and we illustrate them with empirical evidence from historical case studies. The non-linear and non-ergodic nature of nation-building processes that our model stresses poses a challenge for standard empirical work and emphasizes the importance of studying nation-building episodes on a case-by-case basis. In addition, as propositions 2, 3, 4, and 7 suggest, a myriad of different factors may tilt the balance for states about whether to implement nation-building policies. Moreover, how these sets of factors bundle together may differ across time and space. Hence, identifying the specific factors that played a role in different historical experiences may be more informative than looking for universal causes of nation-building.²⁵

Despite the absence of a specific set of causes triggering nation-building episodes, our results highlight some important characteristics that they share. First, our model has multiple steady states, accounting for the diversity we observe across nations with respect to the spread of national sentiments in peripheral regions. Second, the steady state that is finally reached is very sensitive to initial conditions, because small initial differences could lead to opposite choices with a high degree of path dependency. This prediction fits a rich set of historical cases in which regions that were initially similar in terms of economic, political, and social factors have experienced completely divergent trajectories.²⁶

²⁵In fact, the vast literature in political science on this topic also suggests that general theories on the *causes* of nation-building and national identities are unhelpful for particular cases. See McCrone (1998) for further elaboration of this argument, and Wimmer and Feinstein (2010) for supportive empirical evidence.

²⁶In addition, this dependence on previous choices explains why short-sighted governments that do not internalize identity dynamics sometimes pursue cultural decentralization policies that become difficult to reverse.

Third, the model displays drastic regime changes: Sudden and extreme shifts in (de)centralization policies can be caused by exogenous shocks to the relative costs and benefits of nation-building, leading to opposite long-run outcomes.²⁷ Fourth, nation-building episodes are relatively fast and extreme (MRAP behavior). Fifth, once sufficiently spread across the population, territorial identities are remarkably stable and difficult to reverse, because identities are reproduced over time inside the family (vertical socialization) and reinforced by the community (horizontal socialization). Sixth, regional identities are more likely to thrive in democracies than in dictatorships, because regionalists may be pivotal in reaching electoral majorities. Seventh, the model predicts a two-way causality between state centralization (δ) and the spread of the national identity ($1 - q$), because they tend to reinforce each other and to bundle together. Finally, although peripheral regions exist in which both identities have similar levels of attachment in the population, the model suggests that in the long run, one of the two identities will tend to become predominant.

Moreover, our theory highlights the fundamental role played of the state in purposefully shaping national identifications, which is in contrast to other explanations that emphasize a more bottom-up nature of this process and see national identifications as a byproduct of state modernization.²⁸ Although these explanations are able to explain the rise of nationalism as an historical phenomenon and identify some necessary conditions for nation-building, they cannot explain why regional identities survived within countries that went through the process of modernization.²⁹ Moreover, they cannot account for the divergence of outcomes for initially similar regions. In this sense, the two case studies discussed below provide support for the key role of top-down state socialization, because the crucial difference in both cases was the capacity and willingness of governments to promote national identities.

²⁷Nevertheless, changes in the parameters on the model or in the location of q_0 are not necessarily exogenous, because states may have additional instruments that affect them. For instance, Sambanis et al. (2015) note that some states, such as Prussia in the 19th century, engaged in international wars with the purpose of unifying the country by changing the value of q_0 . Others have relied on internal migration to start the process of nation-building in peripheral regions (McGarry, 1998).

²⁸The general idea underlying these theories is that the processes of urbanization and industrialization broke traditional parochial networks and spurred labor mobility. As a consequence, citizens from different backgrounds interacted with each other, which, in turn, promoted the emergence of a national identification. Moreover, this process of creating “imagined communities” was catalyzed by the advent of technologies permitting mass literacy in vernacular languages (Anderson, 1991).

²⁹For instance, (Robinson, 2014) argues the strength of ethnic identifications intensified parallel with the process of modernization of African countries after de-colonization.

Case study I: Spanish vs. French Catalonia

Catalonia constitutes a paradigmatic case in which small differences in the initial conditions could amplify over time through the evolutionary process. At the beginning of the 17th century, Catalonia was an homogeneous region that was later split between France and Spain by the treaty of the Pyrenees in 1659. Since the split, Catalan national identity has persisted in both countries. However, this identity is prominent in today's Spanish political and social life, whereas it is almost negligible in France.³⁰

According to McRoberts (2001), both parts of Catalonia remained relatively similar until the 18th century. In particular, Catalan regions on both sides of the border shared a common past, presented analogous degrees of linguistic and cultural homogeneity, and even kept some common legal institutions that were relatively independent from the central states. Despite these similarities, an obvious crucial difference was the ruling state and, in particular, the different set of incentives and constraints each state faced during the first decades of the 19th century. On the one hand, the Spanish state lacked the means to implement and enforce nation-building policies due to its inability to collect taxes (De Riquer, 2001). Moreover, its political power in relation to Catalan elites was relatively low, and the latter were able to mobilize the population around the idea of a Catalan identity (McRoberts, 2001). On the other hand, as Weber (1976) notes, during this time, the French state was characterized by a relatively strong state capacity. Therefore, it had the resources to implement mass schooling with a school curriculum designed to “make” French citizens.³¹ By contrast, as Enguita (2012) argues, universal schooling in Spain was not effective until the late 20th century, because the government did not have the financial means to provide schooling in rural areas. Moreover, the incentives for both states were relatively different. In particular, industrializing opportunities in France were greater at the beginning of the 19th century, increasing the benefits for the French central government of having an homogeneous population that could communicate in the same language (Gellner and

³⁰For instance, according to Ethnologue, Catalan is the main language of communication for around 50% of the population in Spanish Catalonia, whereas this figure is 1% in France.

³¹The case of French Catalonia also shows that successful nation-building episodes tend to be fast and characterized by homogenization policies that disregard regional particularities. As Weber (1976) notes, within a single generation, locally identified peasants became loyal “Frenchmen.” This identification was achieved by what can be considered extreme (MRAP) interventions, such as the implementation of mandatory schooling and military services, the suppression of the use of the Catalan language in public administration, and the abolition of local institutions.

Breuilly, 1983). In addition, the French state was involved in several external wars that needed an army of soldiers willing to fight for their nation (Aghion et al., 2018). In our model, these initial differences correspond to slightly different levels of fiscal capacity (r), damage created by regionalist political unrest ($1 - \beta$) and the nation-building motive parameter (ψ^N), but not by substantive initial differences in q_0 .

The Catalan case also highlights the high degree of path-dependency implied by the model. As Balcells (2013) shows, the failure of the Spanish state to spread the Spanish identity during the period of the “scholastic revolution” enabled the appearance of a regional revival movement (the “*Renaixença*”) that promoted Catalan cultural values. As a result, more people were socialized to the Catalan identity. Therefore, when industrialization and state capacity levels in Spain resembled those in France a few decades earlier, instilling the national identity was no longer desirable for the Spanish government. In fact, the size of the regionalist group continued to grow and its political importance pushed the Spanish state to progressively grant more decentralization during the early 20th century.

Finally, the Catalan case exemplifies the difficulties in reverting formerly instilled territorial identities, as well as the importance of exogenous shocks in provoking drastic changes in nation-building policies. After the civil war (1936-1939), Franco’s regime started a process of massive centralization, and implemented repressive and brutal measures aimed at eliminating the Catalan identity, in sharp contrast to the federal approach and policies of cultural recognition of the Second Spanish Republic (1931-1939). In terms of the model, the civil war and the posterior establishment of a fascist regime could be interpreted as a large shock that triggered a severe regime change in policies. In particular, it can be seen as a shock to the ability of Catalan regionalists to organize and protest, because during the civil war many of the Catalan leaders and citizens were killed, imprisoned, or forced into exile (a decrease in $1 - \beta$). Also, the ideological shift of the new fascist regime, based on extreme Spanish nationalism, can be interpreted as a shock to ψ^N . However, Catalan identity was already widespread in society due to the weakness of the Spanish state during the 19th century (small q_0), and the measures implemented by Franco’s regime did not significantly alter these previous identifications. In fact, with the advent of democracy 40 years later, the majority of Catalans

were still attached to the regional identity.³² This new “shock” altered again the balance between the costs and benefits of promoting the national identity (because the success of the new democracy depended in part on the acceptance by the Catalan population), which triggered a radical switch in policies in favor of more federalism and the recognition of the cultural distinctiveness of the Catalan region.³³

Case study II: Tanzania vs. Kenya

The case of Tanzania and Kenya is even more paradigmatic of the contingent nature of nation-building attempts and of the difficulties present in identifying a set of fundamental causes behind them. In a fascinating work, Miguel (2004) shows how the post-independence governments of Kenya and Tanzania, two countries that were similar in many respects, pursued radically different nation-building policies: Whereas the Tanzania government promoted Swahili as a national language, praised a national identity in schools, and dismantled tribal authorities, the Kenyan government allowed ethnic division. Over time, this difference in approaches created a national identity in Tanzania that is not present in Kenya.

The institutional and historical similarities between these two countries were remarkable, both in the colonial and early post-colonial period. Both countries have a similar geography and population density, were former British colonies, became independent in the 1960s, and started afterwards from similar economic conditions. In the political realm, they both formed a one-party system and inherited similar administrative structures from the colonial period (Weber, 2010). Despite these similarities, (Barkan, 1994) claimed that the fundamental difference between the two countries lay in the fact that Tanzania had a large number of small ethnic groups, whereas Kenya was populated by fewer but larger ethnic groups with sufficient power to oppose the government’s policies. Moreover, the concentration of capital was relatively higher in Kenya (Iliffe, 1979), which enabled some ethnic groups to arm themselves

³²Given the intergenerational nature of this process, possibly not enough time passed for nation-building policies to have a significant effect, especially given the low rate of change at early stages.

³³This case also highlights one of the predictions of the electoral model in section 1.6 regarding the possible different behavior of democracies and dictatorships. In particular, democracy gave Catalan regionalists the ability to influence the territorial and identity policies of the central government. In fact, Catalan parties have been pivotal in the national parliament several times during the democratic period.

and fight violently the against the state, preventing nation-building. In our model, this explanation corresponds to a different initial value of q_0 and different value of β . Nevertheless, as Miguel (2004) notes, it seems that the crucial but small initial difference that led to divergent nation-building experiences were just “the personalities and philosophies of the respective independence leaders, Jomo Kenyatta and Julius Nyerere.”

Moreover, the Tanzanian case nicely illustrates how nation-building policies are generally implemented in a fast and radical way, as implied by the MRAP solution of our model. As Miguel notes, “The Tanzanian regime quickly pushed for total Swahilization of government administration after independence and established the National Swahili Council to promote its use in all spheres of public life.”

1.8 Conclusion and ways forward

In this paper, we develop a theoretical framework to illustrate the main mechanisms in nation-building processes, highlighting the importance of contingent historical circumstances in shaping the ability of states to nation-build. To conclude, we summarize the key results of the paper and outline a few potential extensions.

Our key theoretical contribution is to analyze the problem of a forward-looking leader who internalizes cultural dynamics and solves a zero-sum conflict between identity groups. This exercise yields three main results. First, the model displays multiple steady states and dependence on initial conditions, a typical characteristic of models with a logistic differential equation. Second, although the optimal trajectory of the control may vary with the choice of the objective function, all optimal paths eventually take extreme values for several specifications. That is, the government eventually provides only one type of public good, because providing both implies “investing” in opposite goals at the same time. Third, the optimal long run steady states are culturally homogeneous, even under cultural substitution between vertical and horizontal transmission channels, in contrast to previous results in the literature (e.g., Bisin and Verdier, 2001; Verdier and Zenou, 2018). We have shown that two sufficient conditions must be satisfied for this long-run behavior. First, the cultural leader must be able to shape any non-degenerate identity distribution in any direction. Second, a strong enough conflict must exist over scarce

resources between the two groups. Interestingly, the qualitative results are similar when we introduce dynamic electoral competition, and solve for the Markov Nash equilibrium of the differential game.

One of the main limitations of the benchmark model is that total tax revenues are assumed to be exogeneously given and constant over time. To amend this shortcoming, in Supplementary Appendix A, we explore the robustness of our results to the case of a government that also controls the tax rate. Interestingly, the ability of the government to soften the budget constraint does not change the qualitative results of the analysis. Because the zero-sum conflict between the two groups is still preserved for any positive tax rate, implementing policies that preserve both identities in the long run cannot be optimal.

The framework developed here can incorporate a number of important questions that may be addressed in future research. First, the model can be easily adapted to analyze the dynamic positive feedback between the fiscal capacity of states and the formation of national attachments (Johnson, 2015). In our model, a greater power to tax equips governments with more resources to forge stronger attachments to the nation. However, causality may also work in the other direction, because a shared national identity makes citizens more loyal to the state and facilitates tax collection by relying less on tax enforcement policies and more on quasi-voluntary compliance (Konrad and Qari, 2012). Therefore, endogenizing tax compliance in our model would help understand the role of nation-building policies as a state capacity investment.³⁴

Second, in our model, nation-building policies operate by changing the (vertical) socialization incentives of parents. However, a vast literature in the social sciences and several recent papers in economics emphasize the essential role of the educational system in promoting national identities.³⁵ Our framework can easily accommodate educational tools such as school curricula or mandatory schooling, by allowing the government to affect horizontal socialization directly.

Third, our analysis is based on the assumption of polarized and mutually exclusive identities. Although this type of identity cleavage has been the norm historically (Marx, 2005), in some regions, a number of individuals have “dual” or “mixed” identities, in the sense that they identify simultaneously with both

³⁴See Besley and Persson (2011) for an extensive overview of questions related to state and fiscal capacity.

³⁵See the related literature (Section 1.2) for some examples.

regional and national cultural groups (Hierro and Gallego, 2018; Stepan et al., 2011). Introducing in our model a third group that derives utility from both types of public goods would allow the exploration of whether dual identities limit the scope of conflict that the government faces, which in turn may alter the full homogenization result. Solving the methodological difficulties present in this problem merits a separate paper and constitutes a very interesting extension.

Finally, in the paper, we analyzed the strategic interactions of two cultural leaders whose nation-building goals are aligned. However, we did not consider real-world examples in which central governments compete against local leaders that can encourage some identity resistance. Analyzing this possibility is an excellent topic for future research.

Chapter 2

Storing Power: Market Structure Matters

2.1 Introduction

The transition to a low carbon economy will require massive investments in renewable energy. Renewables provide substantial environmental and economic benefits (Borenstein (2012)), but their deployment is not free of obstacles. In particular, the intermittency of renewables poses a challenge for power systems, in which demand and supply have to be equal at all times. For this reason, the pathways to decarbonizing the power sector increasingly rely on energy storage as a means to counteract the volatility of renewable output.^{1,2} Whether this objective is actually achieved will crucially depend on firms' incentives to operate and invest in storage facilities. The goal of this paper is to characterize how such incentives shape market outcomes, and to understand how they depend on market structure.

By storing electricity when renewables' availability is high and releasing it when availability is low, storage facilitates the integration of renewables in

¹Demand response is also an important source of flexibility. Some of the economic issues it raises are similar to the ones raised by storage, with two important differences. First, consumers are usually considered as price-takers. And second, storage requires heavier investments as compared to demand response. However, behavioral, informational and political considerations often introduce obstacles to demand response (Fabra et al. (2020)).

²For instance, in the big five markets in Europe (Great Britain, France, Germany, Spain and Italy), energy storage could grow from 3 GW today, to 26 GW in 2030, and 89 GW by 2040, representing one fifth of the total capacity additions that are needed to decarbonize the power sector (the rest being wind, solar, interconnectors and gas peakers (McCarthy and Eager (2020))).

electricity markets. Furthermore, because storage improves security of supply, it reduces the need to invest in oil-fired or natural gas back-up generators (European Commission (2020)). And last, but not least, by smoothing production over time, storage reduces generation costs and flattens the price curve, which translates into improved production efficiency and lower prices for consumers. The downside is that the costs of investing in energy storage remain high, despite substantial cost reductions over the past decade (BloombergNEF (2020)).

We investigate if markets send adequate signals for firms to invest in storage, or if it is necessary to put in place regulatory arrangements to align social and private incentives. As it is well known, markets fail in internalizing positive externalities, such as the ones listed above, and this naturally leads to underinvestment. Such externalities, however, may not be the only relevant market failures in this context. Leaving aside externalities, in this paper we show that, on the one hand, perfectly competitive markets (both in storage as well as in generation) induce the socially optimal storage decisions (see also Ambec and Crampes (2019) and Schmalensee (2019)). On the other, we also show that market power distorts storage decisions (operation and investment) in ways that increase costs and consumer payments. For this reason, market structure is a key determinant of the ability of markets to induce efficient signals for storage operators.

We build a stylized theoretical model that captures the key drivers of storage investment and pricing incentives in wholesale electricity markets. In particular, we assume that the market is served by a fringe of non-strategic producers, one strategic producer, and a set of storage owners. In order to endogenize investment decisions, we assume that storage capacity is chosen once and for all, followed by competition in the wholesale market. Demand moves deterministically over time, from low to high levels over a compact interval,³ while production entails increasing marginal costs.

Under the welfare maximizing solutions,⁴ the planner uses storage to shift production from high to low demand periods in order to minimize generation

³In the context of electricity markets, demand can be understood as total load minus renewable generation. Our model focuses on seasonal demand variation while omitting demand uncertainty. Adding uncertainty would complicate the model without changing its main results.

⁴We characterize the first-best (the planner can decide both upon generation and storage) and the second-best (she can only decide upon storage, as generation decisions are market-based).

costs. Moreover, it invests in storage capacity so as to equate the additional marginal cost savings brought about by storage with its per unit investment cost.⁵ At the optimal capacity, production is not fully flattened across time, as the marginal cost savings of adding storage would fall down to zero, i.e., below the investment cost. Under the competitive market solution, storage owners make profits by arbitraging price differences across demand levels. Since competitive prices reflect marginal costs, the arbitrage gains capture the cost savings that storage brings about. Hence, the social and private incentives are aligned, absent other market imperfections.

Market power in generation and storage distorts this outcome in opposite directions. Consider first the case in which there is market power in generation, but not in storage. Since the strategic firm's incentives to withhold output are stronger in high demand periods, the price curve becomes steeper the higher the degree of market power. This makes arbitrage more profitable, inducing storage firms to over-invest.

Consider now the case in which a storage monopolist serves a perfectly competitive energy market. The storage firm is no longer a price taker, i.e., it internalizes the impact of its storage decisions on the prices at which it either buys or sells the stored amounts. This leads the storage owner to smooth its storage decisions over time in order to avoid a strong price reduction when it sells and a strong price increase when it buys (i.e., acting as a monopolist or as a monopsonist, respectively). In turn, this smoothing reduces the profitability of storage, and thus leads to under-investment.

These distortions are enhanced in the case in which a vertically integrated firm has market power in both storage and generation. The reason is that the vertically integrated firm not only internalizes the price impacts on its stored output but also on its own generation. This leads to a greater distortion in the allocation of output across firms. For this reason, under reasonable assumptions, this market structure yields the least efficient market outcome, the lowest level of investment in storage capacity, and the lowest level of consumer surplus.

In sum, we find that total welfare and consumer surplus decline as we introduce more layers of market power. Market power in production creates static

⁵Under the first-best, these marginal cost savings are computed along the industry marginal cost curve. Instead, under the second-best, they are computed along the market supply curve, which is steeper. This implies that the second-best capacity exceeds the first-best capacity.

productive inefficiencies as it distorts the optimal market shares across producers; while market power in storage creates dynamic productive inefficiencies as storage fails to flatten production across demand levels. In both cases, market power gives rise to additional inefficiencies as it distorts the incentives to invest in storage. These impacts ultimately translate into higher prices for consumers.

We illustrate the predictions of our model by simulating the Spanish electricity market under the 2030 energy and environmental targets (MITECO (2020)). Using detailed data on electricity demand, generation units and generation costs, we quantify the improvement in productive efficiency and the reduction in carbon emissions brought about by storage. We also compute the arbitrage profits made by competitive storage firms, and show that they decrease as installed storage capacity increases. Interestingly, our results point at the complementarity between investments in renewables and storage. On the one hand, storage boosts the profitability of renewables by reducing curtailment at times of excessive renewables availability. On the other, renewables enlarge arbitrage profits, as a result of increased price volatility and more frequent zero-price episodes.

However, even in scenarios with large renewables penetration, arbitrage profits are several orders of magnitude lower than the costs of investments. Accordingly, if regulators want to boost investments in storage (as mentioned in their decarbonization pathways), they will have to complement the market revenues of storage owners with public support. For this purpose, they could resort to storage capacity auctions to select those firms that are willing to carry out the investments at least cost. By bundling support to price caps or reliability options (Cramton and Stoft (2008)), they could at least partially correct the distortions created by market power on the optimal use of storage. The auctions' eligibility criteria could also serve to prevent dominant generators from increasing their market power by investing in storage, as that would also result in an inefficient use of the storage facilities.

Related Literature Our paper relates to a long-standing literature on the role of storage technologies in commodity markets. The canonical theory (Newbery and Stiglitz (1979); Wright and Williams (1984)) focuses on the role of storage in balancing stochastic production in a perfectly competitive environment. Subsequent papers in this literature consider alternative market struc-

tures and explore the impact of storage on price volatility and social welfare (McLaren (1999); Newbery (1990); Allaz (1991); Williams and Wright (2005); Thille (2006); Mitraile and Thille (2014)). Our contribution to this literature is two-fold. First, we abstract from issues related to stochastic demand to put the spotlight on the role of strategic interactions and ownership structure. Encompassing different market structures in a single tractable framework allows us to provide a welfare ranking across market structures. Second, in contrast to the previous literature, we characterize endogenous storage investment decisions and relate them to the degree of market power. Interestingly, our results imply that analyzing production and storage decisions in isolation underestimates the welfare distortions created by the exercise of market power.

Our paper also contributes to an emerging literature on the economics of energy storage.⁶ First, a strand of engineering-oriented studies quantify the value of electricity storage for small storage operators that take prices as given (e.g. Shardin and Szölgényi (2016) Steffen and Weber (2016)). In contrast to these papers, our analysis reveals that abstracting from strategic interaction and storage-induced price effects overestimates the profitability of storage investments. Related papers analyze the level of storage capacity needed to deal with the intermittency of renewables (e.g. Pommeret and Schubert (2019)) or examine the economic properties of different storage technologies (e.g. Crampes and Trochet (2019)). More closely related to us, Ambec and Crampes (2019) and Schmalensee (2019) analyze whether perfectly competitive energy markets provide optimal incentives for investing in storage facilities. We depart from these papers by introducing strategic behavior in both generation and storage under various ownership configurations, revealing a wedge between private and social incentives regarding storage decisions.⁷ This incentive misalignment is also present in an empirical paper by Karaduman (2020), who builds a quantitative model of the South Australian Electricity Market to estimate the expected market outcomes under various levels of storage capacity. Our stylized framework complements this analysis in two respects. First, we

⁶Early contributions to this literature analyzed the strategic behavior of conventional hydropower plants (e.g. Garcia et al. (2001)). More recently, analyzes how the ability of hydro power affects firms' ability to exercise market power in the context of the Colombian electricity market.

⁷Sioshansi (2014) and Schill and Kemfert (2011) also compare market outcomes under different market and ownership structures, but do not analyze investment decisions. Nasrolahpour et al. (2016) explores storage investment incentives, but only under the assumption of perfect competition. We also depart from these papers in that, instead of their two period configuration, we allow for a continuum of demand levels. With only two demand levels, storage smoothing would not be possible.

provide analytical closed-form solutions that single out the differences across different market structures. Second, we expand the set of cases considered by analyzing the effects of vertical integration between generation and storage, which is common in most electricity markets in practice.

More broadly, our paper is related to the trade literature that allows for strategic arbitrage across countries. The reason is that trade links markets across space, while storage links markets across time. One notable difference is that trade flows are rarely constrained by the infrastructure linking two markets, while storage is typically limited by binding capacity constraints. Hence, while (in the absence of market power) the law of one price (up to transportation costs) applies to the trade context, it does not apply to the storage case. Energy trade is an exception, as electricity and gas trade require cross-border interconnection capacity. It is thus not surprising to find some similarities between our analysis and papers on electricity trade (Joskow and Tirole (2000, 2005) and Yang (2020)), or gas trade (Ritz (2014); Massol and Banal-Estanol (2018)). The main difference however is that the storage capacity allows to ‘stock’ energy over time, in contrast to the transmission capacity which allows energy to ‘flow’ at an instant of time. Hence, while it is particularly relevant to understand how and when is a binding storage capacity operated, this question becomes simpler in the context of energy trade (i.e., always use the transmission line at full capacity).

Last, our paper connects with the literature on exhaustible natural resources. Indeed, oil, gas, and minerals, among other natural resources, have two common features with electricity: they are storable and often vulnerable to the exercise of market power. This literature has shown that the optimal extraction path of natural resources follows the “Hotelling rule” both for price-taking storage firms (Hotelling (1931)) and for strategic firms (Salant (1976)). Interestingly, our analysis departs from the Hotelling model in that, unlike the case of natural resources in which reserves are exogenously given, in our storage problem firms have to decide not only when to store, but also how much to invest in storage capacity.

The remainder of the paper is organized as follows. Section 2.2 describes the model. Section 2.3 characterizes the solutions to the social planner’s problem when she can take production and storage decisions (first-best) or when she can only decide on storage (second-best). These solutions serve as benchmarks to assess the equilibrium market outcomes characterized in Section 2.4. The

analysis considers three alternative market structures for storage ownership: a fringe of competitive storage owners, an independent storage monopolist, and a vertically integrated storage monopolist. Section 2.5 compares the resulting equilibrium outcomes in terms of consumer surplus and total welfare. Section 2.6 conducts simulations of the Spanish electricity market for different levels of storage capacity. Section 2.7 concludes. All the proofs are postponed to Appendix B.

2.2 The Model

Consider a market for a storable good (e.g., electricity). Its demand can be met through production or through storage. The costs of producing q units are captured by the function $c(q)$, which is increasing and convex, i.e., $c'(q) > 0$ and $c''(q) > 0$. The costs of storing and releasing the good are normalized to zero up to the storage capacity K , while storing above K is impossible. The costs of investing in storage capacity K are given by the function $C(K)$, which is also increasing and convex, i.e., $C'(K) > 0$ and $C''(K) > 0$. To obtain closed-form solutions, we will assume linear marginal costs of production, i.e. $c'(q) = q$. We also assume that the costs of investments are increasing and (weakly) convex, i.e., $C'(K) > 0$ and $C''(K) \geq 0$.

Demand, denoted θ , is assumed to be perfectly inelastic up to the consumers' maximum willingness to pay v .⁸ Demand takes values in the interval $[\underline{\theta}, \bar{\theta}]$ during a certain period of time (i.e., the storage cycle), which we will refer to as a 'day'. However, in the interest of tractability, rather than describing demand by the values it takes during the 'day', we describe it by a load duration curve (Green and Newbery (1992)), i.e., an increasing cumulative distribution function $G(\theta)$ that gives the fraction of the day when demand is below a certain level. This characterization implies that demand is monotonically increasing during the 'day'.⁹ We assume that $G(\theta)$ is everywhere differentiable in the support, with density $g(\theta)$. The density is assumed symmetric around its expected value, denoted $\mathbb{E}(\theta)$.

⁸In the context of electricity markets, θ can be interpreted as total demand net of renewables. We assume that v is large enough so that it never binds. Results do not depend on this assumption, or on demand being price-inelastic, but they make the analysis simpler.

⁹It would be straightforward to extend our analysis to more general demand characterizations, as long as demand during the storage cycle had at most one minimum and one maximum, e.g. if demand follows a sine or cosine function during the day. The notation would be more involved but results would remain unchanged.

Investment, production and storage decisions take place in two stages. In the first stage, storage capacity K is chosen once and for all. In the second stage, production and storage decisions are simultaneously chosen by firms (in those cases in which there is more than one) for each of the possible demand levels of the ‘day’(s). It is inconsequential whether there is a single day or several ‘days’ as long as they are all identical. At the beginning of each ‘day’, the storage capacity K is empty, but it can be filled up during the ‘day’. The stored amounts can be released any time and without constraints. However, at the end of the ‘day’, any stored amount becomes valueless.

Market structure in the product market There are two types of producers: a dominant firm (D) and a set of fringe firms (F). Inspired by Perry and Porter (1985), we assume that the existing production assets are split between them: for each cost level, the dominant firm owns a fraction $\alpha \in (0, 1)$, while the remaining fraction $(1 - \alpha)$ is owned by the fringe. This means that their marginal costs are $c'_D(q) = q/\alpha$ and $c'_F(q) = q/(1 - \alpha)$, respectively. The competitive industry supply curve remains fixed at $q = c'(q)$ irrespectively of the distribution of assets across firms. Firms’ market shares at an efficient output allocation are α for the dominant firm and $(1 - \alpha)$ for the fringe. Any departure from those efficient shares would lead to higher production costs. Last, note that α is a measure of the dominant firm’s size, i.e., at any given price, the higher α the more it can produce without incurring in losses. Equivalently, α is a measure of the dominant firm’s efficiency, i.e., the higher α , the lower the costs that the firm incurs when producing a given quantity.

Market structure in the storage market Regarding storage, we will consider various market structures. First, we will analyze the first-best and the second-best solutions. Under both of them, a social planner chooses how much to invest in storage capacity and when to use it. The difference between the two is that under the first-best, the social planner can also take production decisions, whereas under the second-best, production decisions are market-based. We will compare these benchmarks with three alternative cases in which there is either (i) a continuum of competitive storage firms; (ii) a single independent storage monopolist; or (iii) a vertically integrated firm that owns both production and storage facilities.

2.3 The Social Planner Solutions

2.3.1 The First-Best

Under the first-best, the social planner takes investment, storage and production decisions in order to maximize total welfare. Because total demand is inelastic, total welfare is simply the sum of gross consumers' surplus net of production costs, minus the costs of investing in storage capacity. Since the total amount that has to be produced in order to meet demand is $(\theta - q_S(\theta) + q_B(\theta))$, the first-best solves the following maximization problem:

$$\max_{K, q_B(\theta), q_S(\theta)} W = \int_{\underline{\theta}}^{\bar{\theta}} [v\theta - c(\theta - q_S(\theta) + q_B(\theta))] g(\theta) d\theta - C(K),$$

subject to two intertemporal constraints. First, the storage facilities cannot store beyond their capacity. And second, they cannot release more than what they have stored. Given our assumptions on demand, these two constraints can be written as

$$\int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta \leq K \quad (2.1)$$

$$\int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta \geq \int_{\underline{\theta}}^{\bar{\theta}} q_S(\theta) g(\theta) d\theta. \quad (2.2)$$

Our first lemma characterizes the optimal use of the storage capacity at the first-best solution, denoted as FB . Figure 1 provides an illustration.

Lemma 2 *At the first-best, for given K , the optimal storage decisions are given by:*

$$q_B^{FB}(\theta) = \max\{\theta_1^{FB}(\mu) - \theta, 0\} \quad \text{and} \quad q_S^{FB}(\theta) = \max\{\theta - \theta_2^{FB}(\mu), 0\}$$

where

$$\theta_1^{FB}(\mu) = \mathbb{E}(\theta) - \frac{\mu}{2} \leq \theta_2^{FB}(\mu) = \mathbb{E}(\theta) + \frac{\mu}{2}, \quad (2.3)$$

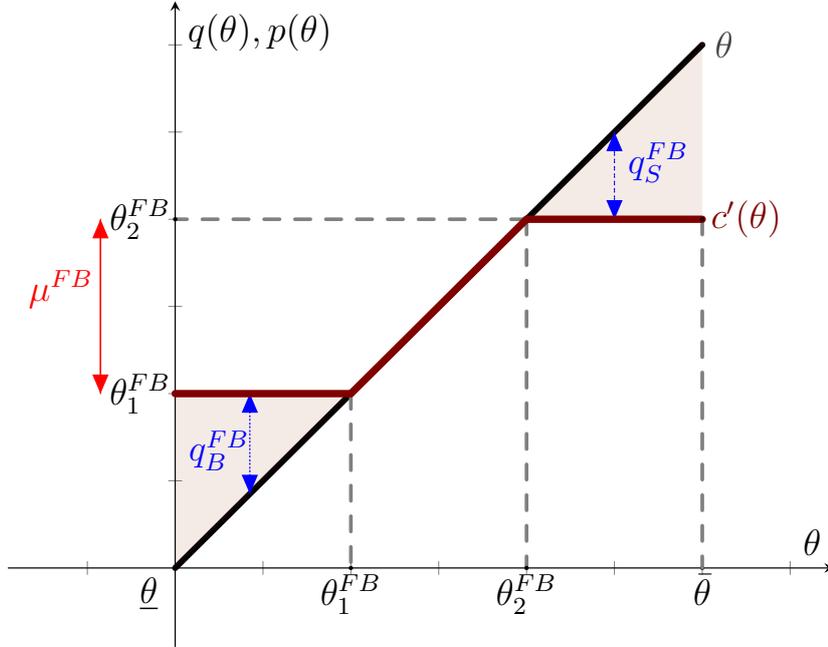
and where μ solves the capacity constraint (2.1) with equality.

Proof. See Appendix B. ■

For given capacity K , storage reduces production costs by smoothing production across time. It is optimal to store so as to flatten production at θ_1^{FB} for $\theta < \theta_1^{FB}$, and to release the stored amounts so as to flatten production at θ_2^{FB}

for $\theta > \theta_2^{FB}$. If the storage capacity does not bind ($\mu = 0$), production and marginal costs are equalized at $\mathbb{E}(\theta)$ across all periods. Instead, a binding capacity constraint ($\mu > 0$) partially prevents this as, for demand levels between θ_1^{FB} and θ_2^{FB} , the storage capacity remains inactive.

Figure 2.1 – Optimal storage decisions under the first-best solution



Notes: This figure illustrates the solution provided by Lemma 2. The brown line represents market demand plus/minus storage decisions. The shaded area represents the amount of stored goods. As can be seen, demand and hence marginal costs are fully flattened whenever the storage facilities are active. The marginal value of storage is found along the industry's marginal cost curve, as depicted by the red arrow.

The marginal value of storage capacity is given by $\theta_2^{FB} - \theta_1^{FB}$, i.e., the marginal cost savings from storing an extra unit of output that costs θ_1^{FB} in order to substitute production that would have cost θ_2^{FB} instead. The higher K , the lower the marginal value of storage as the cost savings from transferring output from θ_2^{FB} to θ_1^{FB} become smaller as θ_2^{FB} and θ_1^{FB} get closer to each other.

This leads to our first Proposition, which characterizes the optimal investment in storage.

Proposition 8 *At the first-best, the optimal investment in storage capacity, $K = K^{FB}$, is the unique solution to*

$$C'(K) = \theta_2^{FB}(K) - \theta_1^{FB}(K) > 0. \quad (2.4)$$

Proof. See Appendix B. ■

At the optimal investment, the marginal value of storage capacity is equal to its unit cost. This implies that the capacity constraint must be binding in equilibrium ($\mu^{FB} > 0$). Otherwise, the marginal value of storage capacity would fall below its unit cost. As a consequence, at the social optimum, storage allows to smooth production and marginal costs, but it does not lead to full price equalization across time.

2.3.2 The Second-Best

The first-best solution assumes that production is efficient, i.e., the market share allocation between the dominant and the fringe firms is efficient. However, in many instances, the social planner has no control over production decisions. Her role is limited to choosing how much to invest in storage capacity and how to operate it. We refer to the solution of the constrained planner's problem as the second-best.

The equilibrium in the product market is simultaneously determined by the storage decisions of the social planner, $\{q_S(\theta), q_B(\theta)\}$, and the output decisions of the dominant firm and the fringe, denoted $q_D(\theta)$ and $q_F(\theta)$ respectively (Cournot assumption). Since the fringe is willing to produce whenever prices are at or above its marginal costs, the fringe's supply is given by $q_F(\theta) = (1 - \alpha)p(\theta)$. Last, because of market-clearing, the inverse residual demand faced by the dominant firm is given by

$$p(\theta; q_S, q_B, q_D) = \frac{\theta - q_S(\theta) + q_B(\theta) - q_D(\theta)}{1 - \alpha}. \quad (2.5)$$

Taking $\{q_S(\theta), q_B(\theta)\}$ as given, the dominant producer chooses its output $q_D(\theta)$ in order to maximize profits over its residual demand for every demand level θ ,

$$\max_{q_D(\theta)} \pi_D = \int_{\theta}^{\bar{\theta}} [p(\theta; q_S, q_B, q_D) q_D(\theta) - c_D(q_D(\theta))] g(\theta) d\theta. \quad (2.6)$$

Our next Lemma gives the resulting output allocation between firms, as well as the market price as a function of $\{q_S(\theta), q_B(\theta)\}$.

Lemma 3 *For given $q_B(\theta)$ and $q_S(\theta)$, the quantities produced by the dominant and fringe producers as a function of the storage decisions are given by*

$$q_D(\theta) = \frac{\alpha}{1 + \alpha} (\theta - q_S(\theta) + q_B(\theta)) < q_F(\theta) = \frac{q_D(\theta)}{\alpha},$$

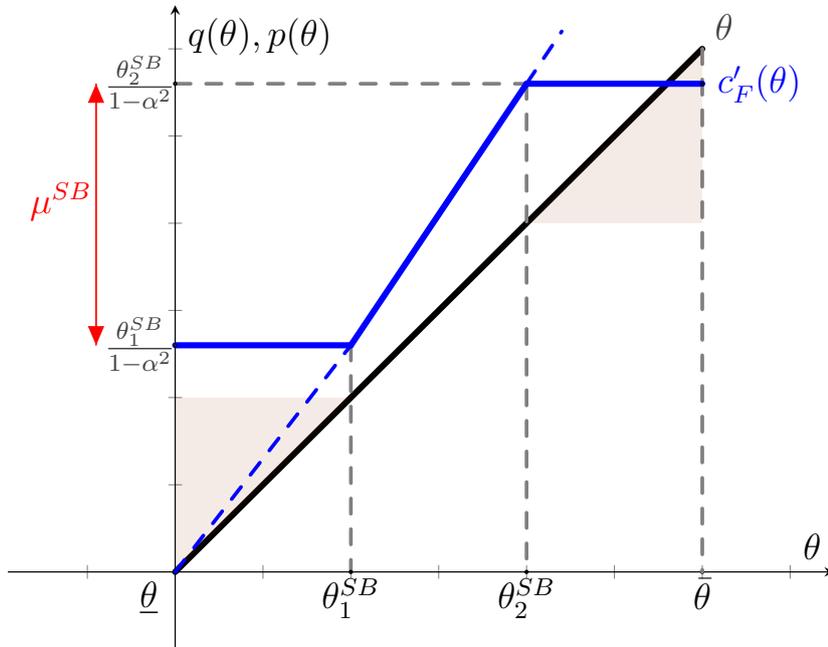
resulting in a market price given by

$$p(\theta; q_S, q_B) = \frac{\theta - q_S(\theta) + q_B(\theta)}{1 - \alpha^2}. \quad (2.7)$$

Proof. See Appendix B. ■

The dominant producer charges a constant price-cost markup equal to α , for all demand levels. Since the fringe operates at marginal costs, firms' market shares depart from the efficient allocation, giving rise to productive inefficiencies. The higher α , the stronger the dominant firm's market power, and the larger the degree of productive inefficiency.

Figure 2.2 – Optimal storage decisions under the second-best solution



Notes: This figure illustrates the solution provided by Lemma 4. The brown line represents market demand plus/minus storage decisions. The shaded area represents the amount of stored goods. The blue line gives prices at every demand level. As can be seen, demand is fully flattened, and the marginal value of storage is found along the price curve, as depicted by the red arrow.

In turn, taking $q_D(\theta)$ as given, the social planner takes storage decisions $\{q_S(\theta), q_B(\theta)\}$ to maximize total welfare,

$$\max_{q_B(\theta), q_S(\theta)} W = \int_{\underline{\theta}}^{\bar{\theta}} v\theta g(\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} [c_D(q_D(\theta)) + c_F(\theta - q_S(\theta) + q_B(\theta) - q_D(\theta))] g(\theta) d\theta,$$

subject to the intertemporal constraints (2.1) and (2.2).

Our next Lemma characterizes, for given K , the planner's storage decisions

under the second-best. The solution is illustrated in Figure 2.2.

Lemma 4 *At the second-best, for given K , the optimal storage decisions are given by:*

$$q_B^{SB}(\theta) = \max \{ \theta_1^{SB}(\mu) - \theta, 0 \} \quad \text{and} \quad q_S^{SB}(\theta) = \max \{ \theta - \theta_2^{SB}(\mu), 0 \}$$

where

$$\theta_1^{SB}(\mu) = \mathbb{E}(\theta) - (1 - \alpha^2) \frac{\mu}{2} \leq \theta_2^{SB}(\mu) = \mathbb{E}(\theta) + (1 - \alpha^2) \frac{\mu}{2}, \quad (2.8)$$

and where μ solves the capacity constraint (2.1) with equality.

Proof. See Appendix B. ■

The storage decisions under the second-best are the same as under the first-best. In particular, storage serves to flatten production at θ_1^{SB} for $\theta < \theta_1^{SB}$ and at θ_2^{SB} for $\theta > \theta_2^{SB}$. Since the storage capacity K is fully used, such demand thresholds are the same as under the first-best. There is however one key difference between Lemmas 2 and 4. Namely, μ^{SB} is now given by the fringe firms' marginal cost savings from moving production from θ_2^{SB} to θ_1^{SB} , and not by the marginal cost savings along the industry competitive supply. The reason is that the social planner takes the dominant firm's supply as given when deciding on the use of the storage facilities (Cournot assumption). Hence, the fringe's supply provides the production flexibility that accommodates the changes in the storage decisions. Since the fringe's supply is steeper than the industry competitive supply, $\mu^{SB} > \mu^{FB}$.

Turning into the optimal investment level, note that the impact of increasing storage capacity on total welfare can be decomposed into two terms:¹⁰

$$\frac{dW}{dK} = \frac{\partial W}{\partial K} + \int_{\underline{\theta}}^{\bar{\theta}} \frac{\partial W}{\partial q_D(\theta)} \frac{\partial q_D(\theta)}{\partial K} g(\theta) d\theta.$$

The first term is a direct effect, which results from relaxing the storage capacity constraint, i.e., it is given by $\mu^{SB} = (\theta_2^{SB} - \theta_1^{SB}) / (1 - \alpha^2)$. The second term is a strategic effect: an increase in storage capacity induces the dominant firm to withhold more output, which enlarges the productive inefficiencies and hence reduces total welfare. It follows that that marginal value of storage capacity is below μ^{SB} .

¹⁰Using the envelope theorem, the effect of the change in storage decisions vanishes out.

In particular, the marginal value of storage capacity is given by the marginal cost savings from storing an extra unit of output when demand is θ_1^{SB} in order to substitute production when demand is θ_2^{SB} . However, unlike the first-best, these cost savings are now evaluated at the equilibrium market shares, with the dominant firm (fringe) producing an inefficiently low (high) market share (Lemma 3). In particular, for given θ , the average marginal costs (weighted by firms' market shares) are given by

$$\frac{\alpha}{1+\alpha}c'_F\left(\frac{\alpha}{1+\alpha}\theta\right) + \frac{1}{1+\alpha}c'_D\left(\frac{1}{1+\alpha}\theta\right) = \frac{1+\alpha-\alpha^2}{(1+\alpha)(1-\alpha^2)}\theta.$$

Therefore, the marginal cost savings brought about by an additional unit of storage are given by the difference of the above expression evaluated at θ_2^{SB} and θ_1^{SB} .

Our next Proposition characterizes the investment decision at the second-best.

Proposition 9 *At the second-best:*

(i) *Equilibrium investment, $K = K^{SB}$, is the unique solution to*

$$C'(K) = \frac{1+\alpha-\alpha^2}{(1+\alpha)(1-\alpha^2)}[\theta_2^{SB}(K) - \theta_1^{SB}(K)].$$

(ii) *There is inefficient over-investment in storage,*

$$K^{SB} = \frac{1+\alpha-\alpha^2}{(1+\alpha)(1-\alpha^2)}K^{FB} > K^{FB}, \quad (2.9)$$

which is increasing in α .

Proof. See Appendix B. ■

How does market power in the product market, α , affect the optimal capacity decisions? The bigger the dominant firm, the more output it withholds. Hence, the marginal cost savings (weighted by firms' market shares) brought about by additional storage are higher the higher α . This implies that the optimal investment at the second-best is higher than at the first-best because it has the additional value of reducing the productive inefficiencies created by market power. Hence, this over-investment is inefficient: if the product market were perfectly competitive, the investment costs of the extra storage capacity would exceed the production cost savings.

The first-best and the second-best serve to assess the market solutions under various market structures, an issue to which we turn next.

2.4 The Market Solutions

In this section we analyze the optimal storage and investment decisions under three alternative ownership structures: (i) there is a fringe of storage owners; (ii) there is an independent storage monopolist; or (iii) there is a vertically integrated storage monopolist.

2.4.1 Competitive Storage

We start by considering the case in which storage facilities are in the hands of a large set of small owners, with free entry in storage. Since the storage and the production facilities are independently owned, for given storage decisions $\{q_B(\theta), q_S(\theta)\}$, the equilibrium in the product market remains as in Lemma 3.

Storage operators earn a return from buying the good when prices are low and selling the good when prices are high. Therefore, for given K , at every demand level θ , their problem is simply to choose how much to store, $q_B(\theta)$, and how much to release, $q_S(\theta)$, so as to maximize their arbitrage profits, taking market prices as given. Formally, their problem can be written as

$$\max_{q_B(\theta), q_S(\theta)} \pi_S = \int_{\underline{\theta}}^{\bar{\theta}} p(\theta) [q_S(\theta) - q_B(\theta)] g(\theta) d\theta, \quad (2.10)$$

subject to the intertemporal constraints (2.1) and (2.2). The free entry condition implies that there is investment in storage capacity until the returns from storage just cover the investment costs.

Not surprisingly, the operation of storage facilities by competitive firms results in the same pattern of storage use as under the social planner solutions. The planner flattens production, which is equivalent to flattening prices, just like the competitive owners do.

Lemma 5 *Under competitive storage, for given K , the equilibrium storage decisions are the same as under the second-best.*

Proof. See Appendix B. ■

For the competitive storage owners, the marginal value of capacity is given by the extra arbitrage profits, i.e., the price difference between storing an extra unit at a price $\theta_1^C/(1 - \alpha^2)$ in order to sell it at a price $\theta_2^C/(1 - \alpha^2)$. Note that

the market price is equal to the marginal cost of the fringe, which is steeper than both the industry marginal cost curve and the average marginal cost of the two firms at the market equilibrium. Hence, the marginal value of capacity for the storage owners is greater than under the first-best and the second-best.

This alone would imply that equilibrium investment is inefficiently high, a result that is further strengthened by the combination of free-entry and cost convexity. In particular, because of the free-entry condition, firms invest in storage capacity up to the level at which the marginal value of storage equals average investment costs. Due to cost convexity,¹¹ average costs are below marginal costs, giving rise to even greater over-investment. In turn, investment is increasing in the degree of market power in the product market, α , as it enhances the marginal value of capacity by making the price curve steeper.

Proposition 10 *When storage is owned by a competitive fringe:*

(i) *Equilibrium investment, $K = K^C$, is the unique solution to*

$$\frac{C(K)}{K} = \frac{\theta_2^C(K) - \theta_1^C(K)}{1 - \alpha^2}. \quad (2.11)$$

(ii) *There is inefficient over-investment in storage, $K^C > K^{SB} > K^{FB}$, which is increasing in α .*

2.4.2 Independent Storage Monopolist

Consider now the case in which the storage facilities are owned by an independent storage monopolist. The main difference with respect to the previous case is that the storage owner now internalizes the effects of its decisions on market prices, and thus on arbitrage profits. Hence, the problem of the storage monopolist can be re-written as in (2.10), now replacing $p(\theta)$ by the inverse demand (2.5),

$$\max_{q_B(\theta), q_S(\theta)} \pi_S = \int_{\underline{\theta}}^{\bar{\theta}} \frac{\theta - q_S(\theta) + q_B(\theta) - q_D(\theta)}{1 - \alpha} [q_S(\theta) - q_B(\theta)] g(\theta) d\theta, \quad (2.12)$$

subject to the intertemporal constraints (2.1) and (2.2). The problem of the dominant producer is still given by (2.6).

¹¹In the investment cost function $C(K)$ were concave, then the comparison with the first-best and second-best would depend on the relationship between α and the degree of cost concavity.

Our next Lemma characterizes, for given K , the use of the storage facilities by the storage monopolist. Figure 2.3 illustrates the solution.

Lemma 6 *When storage is owned by an independent monopolist, for given K , the equilibrium storage decisions are given by:*

$$q_B^M(\theta) = \max \left\{ \frac{\theta_1^M(\mu) - \theta}{2 + \alpha}, 0 \right\} \quad \text{and} \quad q_S^M(\theta) = \max \left\{ \frac{\theta - \theta_2^M(\mu)}{2 + \alpha}, 0 \right\},$$

where

$$\theta_1^M(\mu) = \mathbb{E}(\theta) - \frac{\mu}{2}(1 - \alpha^2) \leq \theta_2^M(\mu) = \mathbb{E}(\theta) + \frac{\mu}{2}(1 - \alpha^2), \quad (2.13)$$

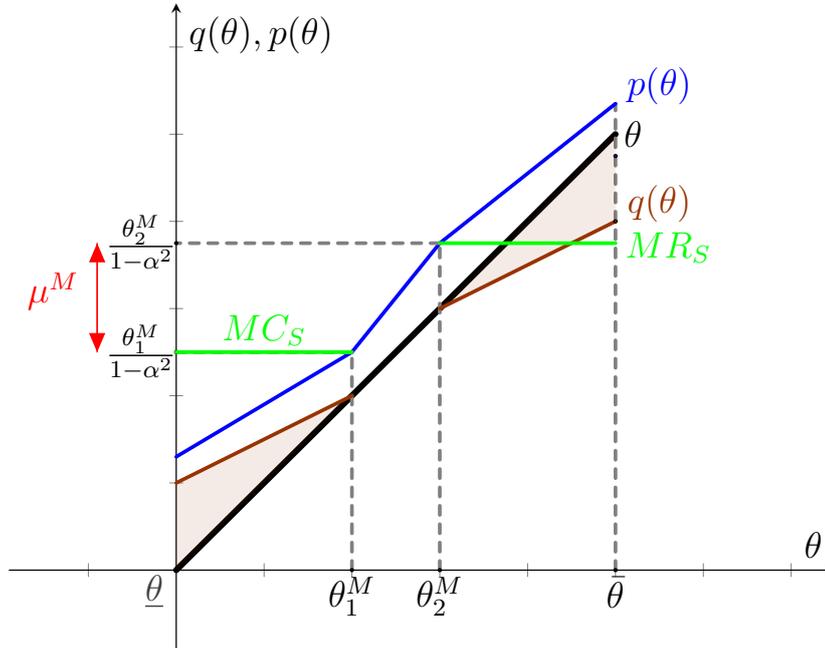
and where μ solves the capacity constraint (2.1) with equality.

Proof. See Appendix B. ■

As in the previous cases, storage allows to shift production across demand levels. Unlike the previous cases, however, it does not lead to a full flattening of production whenever the storage facilities are active. The reason is that the storage monopolist no longer equalizes prices, but rather marginal revenues when it sells (or marginal costs when it buys). As it is standard in a monopoly problem (or symmetrically, in a monopsonist problem), marginal revenue is below the market price because an increase in supply (i.e., an increase in q_S) reduces the price at which the inframarginal units are sold. Symmetrically, an increase in demand (i.e., an increase in q_B) makes it more costly to buy the inframarginal units. Thus, the storage owner smooths storage in order to avoid a strong price reduction when it sells and a strong price increase when it buys. In turn, this prevents production and prices from being fully flattened, and production costs from being minimized.

The comparison of Lemma 6 with Lemmas 2 and 4 shows that market power in storage creates an inefficient use of the storage capacity relative to both the first-best and the second-best. First, when the storage capacity is binding ($\mu > 0$), the region over which the storage facilities are not active is inefficiently short. Because of storage smoothing, the monopolist requires more demand levels to fill the same storage capacity.¹² Second, when the storage capacity constraint is not binding (implying $\mu = 0$), the monopolist under-utilizes the existing storage capacity. In particular, a fraction of the storage capacity

¹²For given K and the endogenous values of μ , we must have $\theta_1^M(\mu^M) \geq \theta_1^{FB}(\mu^{FB})$ and $\theta_2^M(\mu^M) \leq \theta_2^{FB}(\mu^{FB})$.

Figure 2.3 – Equilibrium storage decisions by the storage monopolist


Notes: This figure illustrates the solution provided by Lemma 6. The brown line represents market demand plus/minus storage decisions. The shaded area represents the amount of stored goods. The blue line gives prices at every demand level. As can be seen, the storage monopolist does not fully flatten production, but rather its marginal costs and revenues, as shown by the green lines.

remains idle despite the scope for marginal arbitrage, which would help to reduce production costs. Again, another source of productive inefficiency.

Note that the degree of storage smoothing is positively related to the degree of market power in the product market, α . The higher α , the steeper is the marginal cost of the fringe, and hence the steeper is the residual demand function faced by the storage owner (see equation (2.7)). This makes the storage monopolist willing to smooth storage more in order to avoid sharp price changes.¹³ In sum, market power in production amplifies the inefficient use of the storage capacity due to market power in storage.

For the storage monopolist, the marginal value of capacity is again made of two terms, a direct effect and a strategic effect:

$$\frac{d\pi_S}{dK} = \frac{\partial\pi_S}{\partial K} + \int_{\underline{\theta}}^{\bar{\theta}} \frac{\partial\pi_S}{\partial q_D(\theta)} \frac{\partial q_D(\theta)}{\partial K} g(\theta) d\theta.$$

¹³If the storage monopolist was a Stackelberg leader, there would be less storage smoothing than under a simultaneous quantity choice model. In the releasing region, the storage monopolist would be able to commit to sell more knowing that the dominant producer would respond by increasing withholding, which would mitigate the price reduction. Under simultaneous quantity choices, this strategic effect is not present.

First, as in the case of competitive storage, an extra unit of capacity allows the firm to increase its arbitrage profit by buying an extra unit at $p(\theta_1^M)$ and selling it at $p(\theta_2^M)$, thus making extra profits $(\theta_2^M - \theta_1^M)/(1 - \alpha^2)$. Due to storage smoothing, θ_1^M and θ_2^M are closer to each other than under competitive storage, thus implying that the marginal arbitrage profit is now lower.

However, there is now a second term that enhances the marginal value of capacity for the storage monopolist. In particular, when it adds new capacity and thus sells (buys) more output, the dominant producer restricts its own output (because of strategic substitutability, see Lemma 3). This strategic effect partially mitigates the price reduction (increase), thus making storage capacity more valuable. Since the effects when the storage operator buys or sells are of the same magnitude, this is formally captured by

$$\int_{\underline{\theta}}^{\bar{\theta}} \frac{\partial \pi_S}{\partial q_D(\theta)} \frac{\partial q_D(\theta)}{\partial K} g(\theta) d\theta = 2 \int_{\underline{\theta}}^{\theta_1} \left[\frac{\partial p(\theta)}{\partial q_D(\theta)} \frac{\partial q_D(\theta)}{\partial q_B(\theta)} \frac{\partial q_B(\theta)}{\partial K} q_B(\theta) \right] g(\theta) d\theta > 0.$$

This effect would not be present in the absence of market power in the product market (as the rivals' output decisions would not be affected by the storage decisions, $\partial q_D(\theta)/\partial q_B(\theta) = 0$). Similarly, it would not be present in the absence of market power in storage (as the storage operators would take prices as given, without internalizing the effects of their decisions on market prices, $\partial p(\theta)/\partial K = 0$). Hence, the combination of market power in both production and storage are necessary to uncover this effect.

Our next Proposition characterizes the equilibrium investment.

Proposition 11 *When storage is owned by an independent storage monopolist:*

(i) *Equilibrium investment $K = K^M$ is the unique solution to*

$$C'(K) = \frac{\theta_2^M(K) - \theta_1^M(K)}{1 - \alpha^2} + \frac{2\alpha K}{(1 - \alpha^2)G[\theta_1^M(K)]}. \quad (2.14)$$

(ii) *When $\alpha = 0$, $K^{SB} = K^{FB} > K^M$.*

(iii) *When $\alpha > 0$, if θ is uniformly distributed and $C'(K) = K$, then $K^{SB} > K^{FB} > K^M$.*

Proof. See Appendix B. ■

The comparison of the storage monopolist's solution versus the second-best de-

depends on countervailing forces. In the absence of market power in the wholesale market, storage smoothing reduces the marginal gain from arbitrage, thus leading to less investment than under the second-best. However, the presence of market power in generation pushes in the opposite direction. Whether one effect or the other dominates depends on the relative strength of the two sources of market power, which ultimately depends on the shape of $G(\theta)$ and $C(K)$, as well as on the value of α . We show that for uniformly distributed demand and a linear marginal cost function, the former effect dominates, thus leading to under-investment relative to both the first-best and the second-best.¹⁴

2.4.3 Vertically Integrated Storage Monopolist

We now consider the case in which the dominant producer owns all the storage facilities. Hence, the vertically integrated firm decides both on production as well as on storage. Its profit maximizing problem now becomes

$$\max_{q_D(\theta), q_S(\theta), q_B(\theta)} \pi^I = \int_{\underline{\theta}}^{\bar{\theta}} [p(\theta; q_S, q_B, q_D) [q_D(\theta) - q_B(\theta) + q_S(\theta)] - c_D(q_D(\theta))] g(\theta) d\theta,$$

subject to the intertemporal constraints (2.1) and (2.2), with the market price given by (2.5). As compared to (2.6), the firm now internalizes how its output decisions affect the arbitrage profits made through its storage facilities. Also, as compared to (2.12), the firm now internalizes how its storage decisions affect the revenues made through its own production.

By replacing $q(\theta) = q_D(\theta) - q_B(\theta) + q_S(\theta)$, the problem would be equivalent to

$$\max_{q(\theta), q_S(\theta), q_B(\theta)} \pi^I = \int_{\underline{\theta}}^{\bar{\theta}} [p(\theta; q) q(\theta) - c_D(q(\theta) - q_S(\theta) + q_B(\theta))] g(\theta) d\theta,$$

subject to the intertemporal constraints. As implicit in this formulation, the vertically integrated firm decides on how much output to offer to the market (regardless of whether it comes from its own production or from its storage facilities), and uses storage to minimize the costs of its in-house production. However, its own production is distorted by its incentives to push market prices up.¹⁵

¹⁴This result holds for more general investment cost functions as long as this curve is not very steep for low levels of investment.

¹⁵This shows why this solution differs from the first-best, even when α approaches one. Indeed, the vertically integrated firm withholds output to push prices up, which leads to a distorted output allocation between the fringe and the dominant firm. This source of

Our next lemma characterizes the production and storage decisions of the vertically integrated firm, for given K . Figure 4 illustrates the solution.

Lemma 7 *When storage is owned by the dominant producer, for given K , the equilibrium storage and production decisions are given by:*

$$q_B^I(\theta) = \max \left\{ \frac{\theta_1^I(\mu) - \theta}{2}, 0 \right\} \quad \text{and} \quad q_S^I(\theta) = \max \left\{ \frac{\theta - \theta_2^I(\mu)}{2}, 0 \right\}, \quad \text{and}$$

$$q_D^I(\theta) = \frac{\alpha}{1 + \alpha} \max \{ \theta_1^I(\mu), \theta \} \quad \text{for } \theta < \mathbb{E}(\theta) \quad (2.15)$$

$$q_D^I(\theta) = \frac{\alpha}{1 + \alpha} \min \{ \theta_2^I(\mu), \theta \} \quad \text{for } \theta > \mathbb{E}(\theta) \quad (2.16)$$

where

$$\theta_1^I(\mu) = \mathbb{E}(\theta) - \frac{\mu}{2}(1 + \alpha) \leq \theta_2^I(\mu) = \mathbb{E}(\theta) + \frac{\mu}{2}(1 + \alpha),$$

and where μ solves the capacity constraint (2.1) with equality.

Proof. See Appendix B. ■

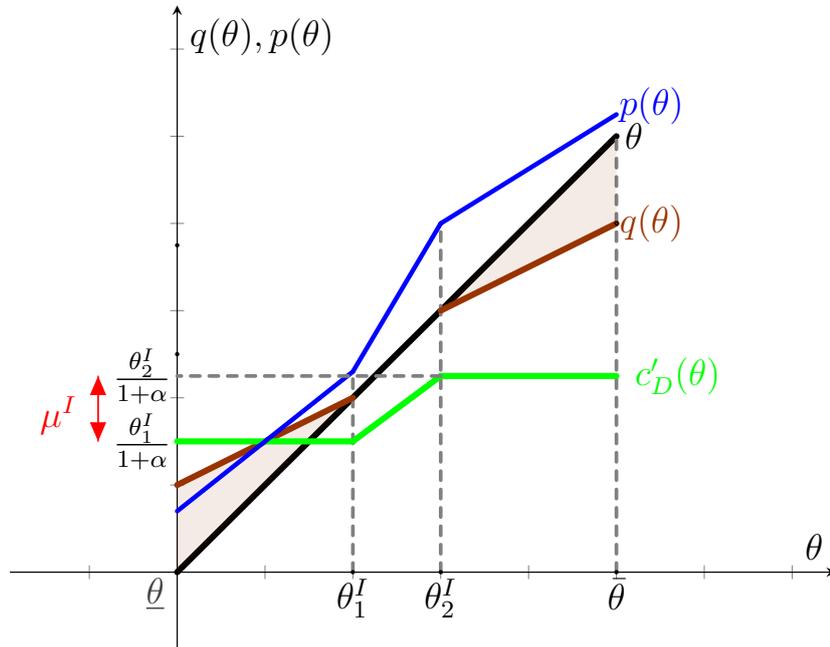
For given capacity K , the vertically integrated firm withholds output to push prices up and uses storage to smooth *its own* production across time. This minimizes its own costs, but gives rise to overall productive inefficiencies. First, because it produces inefficiently little. And second, because as it uses storage to flatten its own production, all the changes in demand are fully met by the fringe's production along its steeper marginal cost curve.

Interestingly, vertical integration changes the pattern of market power over time given that the firm now internalizes the price effects on its net position $q_D(\theta) - q_B(\theta) + q_S(\theta)$. In particular, the vertically integrated firm no longer charges a constant markup at α (as it was the case for the stand-alone producer). Instead, its mark-up is increasing in demand. For $\theta < \theta_1^I$, the firm charges a markup below α because its net position $q_D(\theta) - q_B(\theta)$ is smaller than in the case of the stand-alone producer. This mark-up even becomes negative when $q_D(\theta) < q_B(\theta)$, which is when the firm is a net-buyer and hence exercises monopsony power by reducing prices below marginal costs. Instead, for $\theta > \theta_2^I$, the firm exercises more market power than in the stand-alone case because its net position $q_D(\theta) + q_S(\theta)$ is now larger. This is summarized below.

Corollary 7 *The demand-weighted mark-up charged by the vertically integrated firm is higher than in the stand-alone case. In particular, the firms*

inefficiency is not present under the first-best.

Figure 2.4 – Equilibrium storage decisions by the vertically integrated storage monopolist



Notes: This figure illustrates the solution provided by Lemma 7. The brown line represents market demand plus/minus storage decisions. The shaded area represents the amount of stored goods. The blue line gives prices at every demand level. As can be seen, the vertically integrated firm operates the storage facilities to flatten its own production and thus its own marginal costs, as shown by the green line. Note that prices fall below marginal costs for low θ s for which the firm is a net buyer.

charges a mark-up below (above) α for low demand levels $\theta < \theta_1^I$ (for high demand levels $\theta > \theta_2^I$).

Proof. See Appendix B. ■

Using expressions (2.15) and (2.16), the dominant firm's marginal costs at θ_1^I and θ_2^I are $\theta_1^I/(1+\alpha)$ and $\theta_2^I/(1+\alpha)$, respectively. Hence, the marginal value of storage capacity is captured by the marginal cost savings from storing a unit of output that costs $\theta_1^I/(1+\alpha)$ in order to substitute production that would have cost $\theta_2^I/(1+\alpha)$. Accordingly, $\mu^I = (\theta_2^I - \theta_1^I)/(1+\alpha)$. Note that these are the marginal cost savings of the vertically integrated firm, which are below those at the industry level. As a result, there is inefficient under-investment in storage capacity as compared to the first-best. In turn, since the second-best capacity is above the first-best capacity, the equilibrium capacity is also inefficiently low with respect to the second-best.

Two key differences in investment incentives explain the departure from the first-best. The first difference comes from the ability of the vertically integrated

firm to exercise market power. To see this, note that in both cases the marginal value of storage capacity is equal to the marginal costs savings for the dominant firm. However, in the first-best these coincide with the marginal cost savings at the industry level, as the dominant firm is producing the socially efficient share of output. In contrast, a dominant firm that behaves strategically withholds output and depresses its marginal costs, which reduces the need to smooth its total production costs by investing in storage capacity. Second, even if the dominant firm priced at marginal cost in the production stage, its incentives to invest in storage would still remain lower, as it does not internalize the cost savings that storage facilities would provide to the competitive fringe.

Proposition 12 *When storage is owned by the dominant producer:*

(i) *Equilibrium investment $K = K^I$ is the unique solution to*

$$C'(K) = \frac{\theta_2^I(K) - \theta_1^I(K)}{1 + \alpha}. \quad (2.17)$$

(ii) *There is inefficient under-investment in storage, $K^{SB} > K^{FB} > K^I$. The distortion in increasing in α .*

Proof. See Appendix B. ■

Interestingly, Storage capacity K^I decreases in the degree of market power in the product market, α . A larger α implies that the dominant firm has lower and flatter marginal costs. Additionally, since a larger α implies that the dominant firm withholds more, the marginal cost savings are computed over a flatter region of the cost function. In sum, the marginal cost savings brought about by an extra unit of capacity are lower the more market power there is, thus making additional storage capacity less valuable the higher α .

2.5 Comparison across Market Structures

In this section, we compare equilibrium outcomes across market structures on equilibrium outcomes to assess the impacts on consumers surplus and overall efficiency. We start by performing the comparison for a given non-binding storage capacity, and then compare market outcomes under a binding capacity constraint. In all cases, we take K as given in order to understand how different

players would use a given storage capacity chosen by the regulator.¹⁶

2.5.1 Non-Binding Storage Capacity

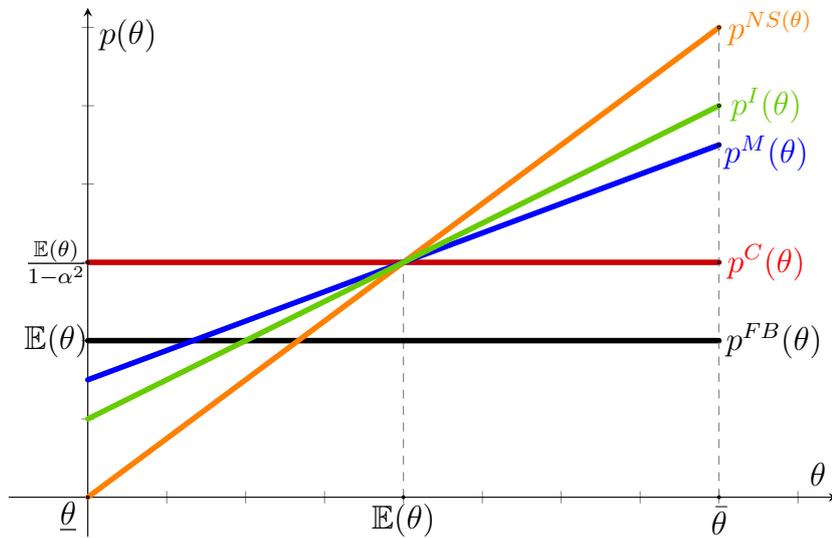
Consider first the case in which existing storage capacity is large enough so that it never binds, regardless of the market structure.

Consumer's surplus can be defined as

$$CS = \int_{\underline{\theta}}^{\bar{\theta}} [v - p^i(\theta)] \theta g(\theta) d\theta = v\mathbb{E}(\theta) - \mathbb{E}[p].$$

Hence, differences in consumer surplus across market structures are fully driven by differences in the market-weighted average price, denoted $\mathbb{E}[p]$. Market structures affect (i) the price levels for each demand realization, as well as (ii) the slope of the price pattern over time. Clearly, $\mathbb{E}[p]$ is higher under market structures that give rise to steeper price patterns, even if the unweighted average prices coincide.

Figure 2.5 – Equilibrium prices under different scenarios.



Notes: This figure depicts equilibrium prices for every demand level θ across all market structures: *FB* first-best (black), *SB* second-best and *C* competitive (red), *M* storage monopolist (blue), *I* vertically integrated firm (green) and *NS* no-storage (orange).

Using our previous results, Figure 2.5 plots equilibrium prices as a function of

¹⁶This approach facilitates the comparison across market structures, while providing a welfare ranking that extends to the case with endogenous storage capacity under some convexity conditions regarding the investment cost function.

demand θ under all market structures considered. We would like to highlight three main results. First, as compared to the case with no storage, storage smooths prices across time. However, only under competitive storage are prices fully equalized across demand levels (recall that we are assuming a non-binding capacity constraint). In contrast, market power in storage results in a steep price pattern, although not as steep as in the absence of storage. Second, regardless of who owns the storage facilities, market power in the product market increases the price level. This can be seen by comparing prices under competitive storage and the first best: both are flat, but the former are higher. Last, if storage facilities are monopoly owned, market power in the product market makes the price pattern both higher as well as steeper, more so under vertical integration than in the case of a stand-alone storage monopolist.

Averaging across all demand levels, the demand-weighted average prices under all market structures considered are given by:

$$\begin{aligned}\mathbb{E}[p]^{FB} &= \mathbb{E}[\theta]^2 \\ \mathbb{E}[p]^{SB} &= \mathbb{E}[p]^{FB} \frac{1}{1 - \alpha^2} \\ \mathbb{E}[p]^C &= \mathbb{E}[p]^{SB} \\ \mathbb{E}[p]^M &= \mathbb{E}[p]^{SB} + Var[\theta] \frac{1}{(1 - \alpha)(2 + \alpha)} \\ \mathbb{E}[p]^I &= \mathbb{E}[p]^{SB} + Var[\theta] \frac{1}{2(1 - \alpha)} \\ \mathbb{E}[p]^{NS} &= \mathbb{E}[p]^{SB} + Var[\theta] \frac{1}{1 - \alpha^2}\end{aligned}$$

Average prices under the first-best simply reflect the average across marginal costs. In all other cases, prices are increasing in α , reflecting two types of mark-ups (*i*) a mark-up due to market power in the energy market (which is a function of α), and (*ii*) a markup due to market power in storage (which depends on α and $Var[\theta]$ as both affect the slope of the price pattern faced by storage owners). Since the first mark-up is common across all market structures, the price comparison solely depends on the distortions due to the use of storage. Comparing these expressions, it immediately follows that

$$\mathbb{E}[p]^{FB} < \mathbb{E}[p]^{SB} = \mathbb{E}[p]^C < \mathbb{E}[p]^M < \mathbb{E}[p]^I < \mathbb{E}[p]^{NS}.$$

Regarding total welfare, since demand is assumed to be price-inelastic, it can

be expressed as simply the sum of gross consumer surplus minus total costs:

$$TW = v\mathbb{E}(\theta) - \int_{\theta}^{\bar{\theta}} \left(\frac{q_D^2(\theta)}{2\alpha} + \frac{q_F^2(\theta)}{2(1-\alpha)} \right) g(\theta) d\theta$$

Similarly as before, total costs can be decomposed into two terms:¹⁷ (i) one reflecting static productive inefficiencies, and (ii) another one reflecting dynamic production inefficiencies due to the distorted use of storage. On the one hand, total costs increase due to market power in the product market, which results in distorted market shares between the dominant firm and the fringe. Second, total costs increase due to the misuse of storage, which results in a lack of production equalization across time. As with prices, the second distortion is also amplified by market power in the product market.

The following Proposition summarizes the above results:

Proposition 13 *Suppose that K is large enough so that it never binds.*

(i) *The ranking of consumer surplus across market structures is:*

$$CS^{FB} > CS^{SB} = CS^C > CS^M > CS^I > CS^{NS}.$$

(ii) *Last, the ranking of total welfare is given by:*

$$W^{FB} > W^{SB} = W^C > W^M > W^I > W^{NS}.$$

In sum, as we add more layers of market power, the average price goes up, and consumer surplus and welfare go down. A scenario with market power in the product market is better (for consumers and for welfare) than one in which, in addition, there is market power in storage. In turn, the latter is better than the one in which the storage monopolist is vertically integrated with the dominant producer. In any event, the worst scenario is the one without storage.

2.5.2 Binding Storage Capacity

Consider now the case in which the storage capacity constraint is binding under all market structures considered.¹⁸ The following proposition summarizes the

¹⁷The expressions can be found in Appendix B.

¹⁸This implies that storage capacity is lower than the maximum level that a storage monopolist effectively utilizes when this storage capacity is not binding. That is: $K < \int_{\theta}^{\mathbb{E}[\theta]} \frac{\mathbb{E}[\theta] - \theta}{2 + \alpha} g(\theta) d\theta$. The results in this subsection extend to cases in which storage capacity

results.

Proposition 14 *Suppose that K is small enough so that it is binding under all market structures considered.*

(i) *For any cumulative distribution function $G(\theta)$, and for any binding storage capacity K , the ranking of consumer surplus and welfare across market structures is given by:*

$$CS^{FB}(K) > CS^{SB}(K) = CS^C(K) > CS^M(K) \& CS^I(K)$$

$$W^{FB}(K) > W^{SB}(K) = W^C(K) > W^M(K) \& W^I(K)$$

(ii) *When demand θ is uniformly distributed on $[\underline{\theta}, \bar{\theta}]$, for any binding K :*

$$CS^M(K) > CS^I(K)$$

$$W^M(K) > W^I(K)$$

Proof. See Appendix B. ■

Proposition 14 shows that storage capacity in the hands of smaller owners has a higher social value than when it is owned by either large storage firms or large generators. Indeed, similarly to the case with a non-binding constraint, the second-best is achieved under competitive storage. When storage owners are price-takers, they exploit all arbitrage until prices become flat (unless the capacity constraint is binding). In turn, this allows to minimize production costs, just as under the second-best. Therefore, one way for the regulator to implement the second-best is to allow small firms to invest up to K^{SB} and no more (recall that $K^{SB} < K^C$). An alternative would be to run storage auctions for a capacity equal to K^{SB} , but only allow small operators to participate. These alternatives would allow the regulator to correct the distortions arising from competitive over-investment while relying on the market to efficiently operate storage facilities.

Proposition 14 also reveals that it is not enough to promote investments in storage if market power in production remains. The reason is that storage facilities will be inefficiently operated if market prices are distorted due to market power.

is binding for some but not all market structures.

The comparison of consumer and total welfare in the stand-alone storage monopolist case versus the vertically integrated case depends on countervailing forces. On the one hand, the stand-alone monopolist spreads the use of storage more across time in order to avoid strong price effects. This results in higher production costs (recall that, for given K , $\theta_2^M(K) - \theta_1^M(K) < \theta_2^I(K) - \theta_1^M(I)$). On the other hand, when the dominant producer owns storage it has stronger incentives to withhold output in order to push market prices up. This creates larger static production inefficiencies as the dominant firm's output is replaced by the fringe's. Which of these two effects dominates depends on the degree of market power α and on the shape of the demand distribution $G(\theta)$. For reasonable assumptions on the demand distribution, e.g. uniformly distributed demand, the second effects dominates. This suggests that allocating storage capacity to vertically integrated firms may result in the lowest level of consumer surplus and overall efficiency.

In sum, regulators should avoid allocating storage capacity to dominant operators, particularly so if they are vertically integrated firms. This conclusion is further strengthened if regulators rely on pure market mechanisms to spur storage investments. As we have seen, the endogenous investment decisions of a vertically integrated firm departs the most from the second best solution, further compounding the inefficiencies arising from a distorted use of storage.

2.6 Simulation of the Spanish Electricity Market

In this section we illustrate some of our theoretical results using actual market data. In particular, we assess equilibrium market outcomes under different levels of storage capacity, and quantify the profitability of storage investment depending on whether electricity producers act competitively or strategically. For these purposes, we perform a series of simulations using the multi-unit auction model developed in De Frutos and Fabra (2012). The model characterizes equilibrium bidding by electricity generators who compete by submitting step-wise bid functions to the auctioneer. Production and prices are set according to a uniform-price auction.

The set of parameters used in the simulations closely replicate the Spanish wholesale electricity market. All simulations are conducted at the hourly level

over a one year period (8,760 hours). The technology mix (in terms of capacities) has been set according to the 2030 energy targets, following the 2021-2030 Spanish National Energy and Climate Plan. This includes the deployment of new renewable capacity (mainly solar, but also wind) and the phase out of coal plants and half of the nuclear capacity. The objective is that by 2030, 74% of electricity generation will come from renewable sources.¹⁹ For the plants' ownership structure, we have assumed that all new capacity additions are in the hands of competitive firms.²⁰ The hourly electricity demand patterns have been set as reported by the Spanish System Operator for 2017 (source: esios.ree.es). The hourly availability factors of the renewable resources have been set at the average of the previous five years, as also reported by the Spanish System Operator. By multiplying these factors times the installed capacity of each technology, we obtain the renewable production at an hourly basis. Because of its storability, hydro production has been allocated to shave the peaks of demand net of renewables. Last, the daily prices for CO2 (EUA) and gas (TTF Hub) have been set at the 2017 prices in international markets (source: Bloomberg). With detailed data about the gas plants' heat rates and CO2 emission factors, we compute the marginal cost used in the simulations.²¹

2.6.1 Scenarios

For different levels of storage capacity, we compute the optimal storing and releasing decisions when the storage capacity is operated by either an unconstrained social planner (first-best), by a constrained social planner (second-best), or by a set of competitive storage operators (competitive storage). Since the arbitrage gains of competitive storage owners are above the ones internalized by storage owners with market power, our analysis provides an upper bound to the investment incentives provided by market prices. We assume that the time frame for storage operation (i.e. the full storage/release cycle) is the natural day,²² and that the round-trip efficiency of storage is 0.85, i.e.,

¹⁹See "Plan Nacional Integrado de Energía y Clima 2021-2030", MITECO (2020).

²⁰Many capacity additions will certainly be in the hands of the large electricity producers. To the extent that this gives them more market power, our estimated mark-ups provide a lower lower bound of the degree of market power.

²¹De Frutos and Fabra (2012) provide a detailed description of how these costs are computed.

²²The time frame depends on several factors, such as the technical characteristics of the storage technology (batteries, pumped storage, etc...). Currently, the most common storage solutions use a 4-hour battery. In the Online Appendix we present simulations for alternative assumptions regarding the optimal cycle (two days, one week, one month) and

Figure 2.6 – Simulated outcomes

	K=0		K=10.000	
	Competitive	Strategic	Competitive	Strategic
Demand	29172	29172	29172	29172
Wind	46,9%	46,8%	47,1%	47,0%
Solar	27,0%	27,0%	27,3%	27,3%
Hydro	6,1%	6,1%	6,1%	6,1%
Nuclear	4,9%	4,5%	5,0%	4,7%
CCGT	3,6%	3,6%	3,0%	3,0%
Other renewables	11,5%	12,0%	11,7%	12,1%
Average price	17,25	21,65	17,67	21,88
Standard deviation	19,76	21,67	19,89	21,65
Price (max)	33,91	41,94	32,76	40,94
Price (min)	0,84	0,88	1,02	1,32

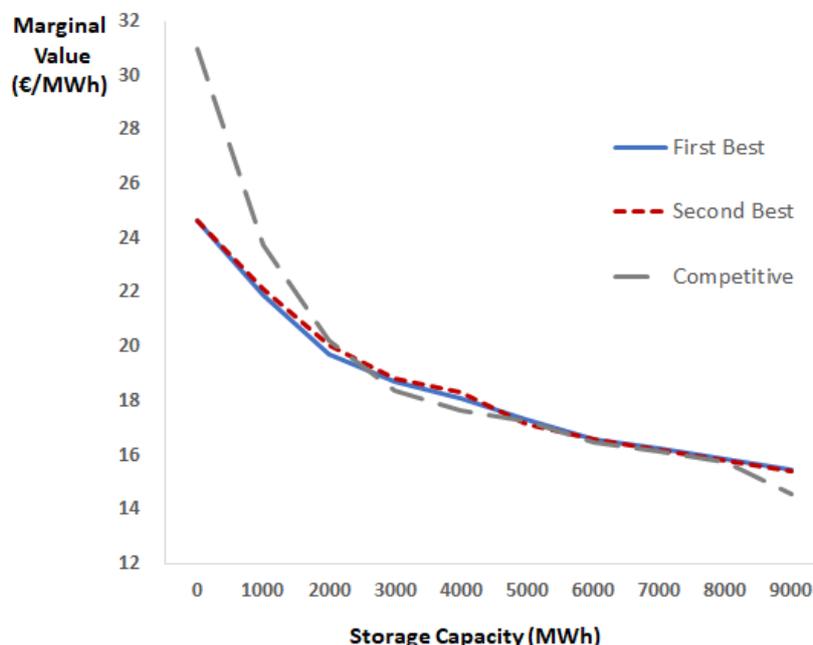
Notes: The table displays average hourly demand (MWh), the share of each technology in the generation mix (%), average demand-weighted hourly price (Euro/MWh), average maximum and minimum prices (Euro/MWh) and their standard deviation (Euro/MWh). Two scenarios are considered, one without storage (first two columns) and the other one with storage (last two columns). In each case, we report the results when all generators behave competitively, or when they bid strategically. Storage is operated by competitive non-integrated firms.

there is a 15% efficiency loss. For each scenario, we compute the equilibria that would arise if electricity producers behaved competitively (i.e. by bidding at marginal cost) or strategically (i.e. by playing the Nash equilibria in bid functions as in De Frutos and Fabra (2012)).

2.6.2 Results

For illustrative purposes, Table 2.6 first provides the summary statistics of two sets of simulations, with and without storage. As it can be seen, more than three quarters of total demand will be covered by renewables, the rest being nuclear, hydro and gas. Average prices will be between 17-22 Euro/MWh on average, being lower in the competitive than in the strategic scenario, and slightly larger in the scenario with storage. Storage has a clear impact on the maximum and minimum prices, as the price curve gets flatter.

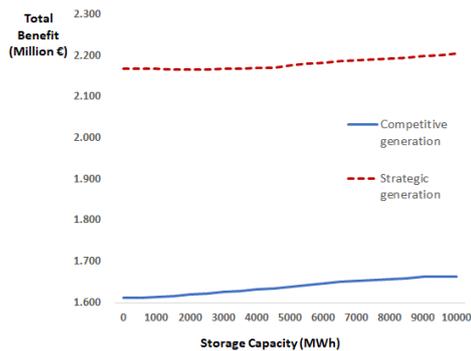
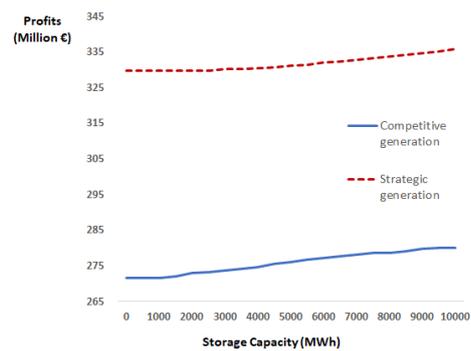
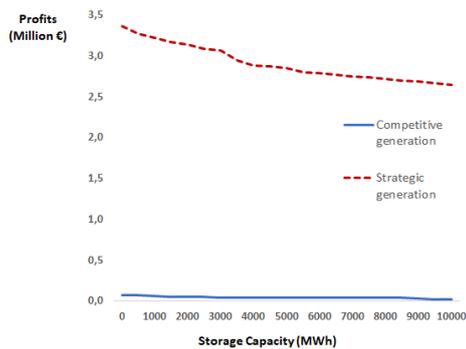
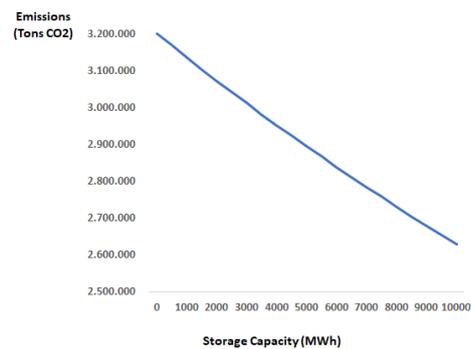
for daily cycles that do not correspond to the natural day. Obviously, a longer time frame comes with bigger marginal arbitrage benefits. However, when considering the lower usage over the life-cycle of the storage facility, daily cycles turn out to be the most profitable.

Figure 2.7 – Marginal value of storage capacity (high renewables penetration)

Notes: This figure shows the marginal value of storage capacity investments as a function of capacity, given the generation technology mix that is expected for 2030. It reports results for an unconstrained social planner (first-best), a constrained social planner (second-best), and competitive storage. The first best curve is computed as the system marginal cost saving when all generators supply at their marginal cost of production. The second best captures the marginal cost savings when generators act strategically, bidding above marginal production costs. The competitive curve is computed by calculating the marginal profit at different capacity levels when generators behave strategically. The value displayed is the marginal value per cycle/day averaged across all hours of the year.

Figure 2.7 reports the marginal value of storage capacity, for different levels of investment, given the generation technology mix planned for 2030. It reports the results under the three storage ownership scenarios: first-best, second-best, and competitive storage owners. Recall that, for the social planner, the marginal value of storage capacity is given by the marginal cost savings from storing one extra unit of electricity. Likewise, for the competitive storage owners, the marginal value of storage capacity is given by the marginal arbitrage profits.

The first thing to notice is that, in line with our theoretical predictions, these curves are negatively sloped, i.e., adding storage capacity becomes less valuable and less profitable the greater the amount of existing storage capacity. The reason is that the most costly production units are first replaced by the least costly production units, and so on, in decreasing and increasing cost order,

Figure 2.8 – Carbon emissions and profits by technology - Scenario 2030

(a) Profits of wind

(b) Profits of solar PV

(c) Profits of CCGTs

(d) Emissions

Notes: Panels (a), (b) and (c) display market profits by technology (not including investment costs), for different levels of storage capacity, when generators behave competitively (solid line) or strategically (dashed line). Panel (d) reports carbon emissions. Profits of renewable technologies are increasing in storage capacity as these technologies produce more (because there is less curtailment) and they do so at higher prices. Not surprisingly, CCGTs make almost zero profits under the competitive scenario, and strictly positive profits otherwise.

respectively. Hence, the cost savings (and thus the price differences) of storing become increasingly smaller as more storage capacity is introduced.

Also note that the differences across these three curves are small. A plausible explanation is that, with high renewables penetration, the degree of market power is low, leading to small production distortions and to prices that are close to marginal costs. This is particularly the case for high levels of storage capacity. Since this implies that marginal costs and prices are fairly similar in all scenarios, the marginal value of capacity is similar across of them. However, for lower levels of storage capacity, since the strategic producers can exercise more market power, the marginal arbitrage profits for the competitive storage owners exceed the marginal cost savings for the social planner. This might lead to inefficient over-investment under the market-based scenario, as we highlighted in our theoretical analysis.

If we consider an scenario with low renewables penetration, matters are strikingly different. We have also run simulations with 2017 data, when renewables' penetration was much lower: wind generation capacity was half and solar capacity was 8 times lower as compared to the 2030 targets. Results show that the marginal value of storage capacity was negative. This implies that investing in storage capacity was neither profitable for competitive storage owners nor desirable from a social point of view, even if investment costs were negligible.²³ This is so for a two-fold reason. With few renewables, market prices are almost always set by the conventional technologies, whose marginal costs, and the resulting market prices, are fairly constant. This implies that the marginal cost savings and arbitrage profits become so small that they are more than offset by the round-trip efficiency losses of storing and releasing electricity.

The contrast between the positive and the negative marginal value of storage capacity under high and low renewables penetration, respectively, indicate that renewables boost the value of storage capacity. The complementarity goes both ways, as storage also makes investments in renewables more socially valuable and more profitable. This can be seen in Figures 2.8a and 2.8b, which depict the market revenues of wind and solar as storage capacity goes up. Both curves are increasing for two reasons. First, storage prevents renewable curtailment in periods of relatively high renewables production relative to de-

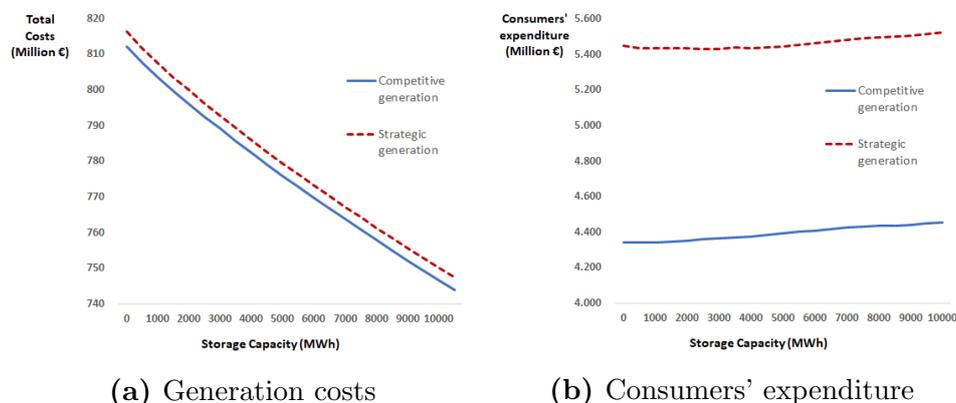
²³We should keep in mind that this analysis is limited to the arbitrage value of energy storage, as we only focus on using energy storage to shift load from peak to off-peak periods. However, energy storage creates other positive externalities which should also be taken into account when computing the optimal investment.

mand. This is particularly important for solar plants, as their production is strongly correlated during the sunny hours of the day, thus making curtailment more likely. And second, storage increases prices in low price hours (i.e., when renewables availability is high) and depresses prices in high price hours (i.e., when renewables availability is low).

Conventional technologies get the other side of the coin, as the reduced curtailment of renewables and the changes in price patterns imply that they sell less and they get paid lower prices on average. Hence, as can be seen in Figure 2.8c, their profits go down as storage capacity goes up. Because conventional production is increasingly replaced by renewables production, carbon emissions go down as well. This is shown in Figure 2.8d.

These various effects have important welfare implications. Figure 2.9a reports the effects of increasing storage capacity on generation costs, while Figure 2.9b reports the effects on consumers' expenditures. As expected, generation costs go down as storage capacity goes up. However, this does not necessarily translate into higher gains for end-consumers. The reason for this ambiguity is that storage decreases prices in some hours but increases them in others. The latter need not be the hours with lower demand, but rather those in which demand net of renewables is lower. Hence, the price increases might take place when consumers' consumption is high, thus implying that the average (demand-weighted) prices they face might well go up as storage capacity increases. In general, the relationship between consumers' prices and storage need not be monotonic.

Finally, to shed light on whether investments in storage are socially desirable, one would have to compare the savings in generation costs against the costs of investing in storage capacity. Current figures report costs of battery storage at around 150 Euro/MWh (IRENA (2017)). Hence, without considering other positive effects of storage (notably, security of supply and learning by doing externalities), it would not be socially optimal or profitable to currently invest in storage (recall from Figure 2.7 that the marginal value of storage is no greater than 24 Euro/MWh). Hence, the costs of storage have to fall and the installed renewables capacity needs to ramp up for there to be a clear case for investments in storage. Over the last ten years, we have witnessed sharp cost reductions in renewables and battery storage (65% to 85% since 2010). Only if this trend continues in the future, will the costs of investing in storage fall below their marginal value. This would strengthen the case for further

Figure 2.9 – Productive efficiency and consumers’ expenditure


Notes: For different levels of storage capacity, panel (a) reports generation costs (not including investment costs). Panel (b) shows firms’ market revenues, thus reflecting the total amount that consumers pay to buy their electricity consumption at market prices. Both panels report the results when generators behave competitively (solid line) or strategically (dashed line).

investing in storage, beyond the objectives of improving security of supply and learning by doing economies.

2.7 Conclusions

There is consensus among the relevant institutions and industry analysts on the strong growth potential of energy storage over the next decade (see for instance, McCarthy and Eager (2020) and European Commission (2020)). However, whether these expectations fully materialise will heavily depend on policy and regulatory decisions which will determine the incentives to operate and invest in storage facilities.

Our focus in this paper has been to analyze how such incentives depend on market structure. We have found that perfectly competitive markets replicate the first-best, absent other market imperfections. However, market power in storage and/or in generation reduces market efficiency through two channels: it induces an inefficient use of the storage facilities, and it distorts investment incentives. Whereas market power in the wholesale electricity market tends to induce over-investment in storage, market power in storage tends to induce firms to under-invest. Under reasonable assumptions, the combination of the two through vertical integration gives rise to the most distorted outcome, both

for consumers as well as for overall efficiency.

Our simulations of the Spanish electricity market show that the arbitrage profits made by storage owners are not enough to cover their investment costs. This implies that without public support, it is doubtful whether the socially optimal investments in energy storage (to the extent that they are positive) will actually take place. The mechanisms designed to grant public support should take into account that market structure matters, i.e., the same storage capacity in the hands of competitive storage owners is more socially valuable than if it is allocated to large storage firms or to generators.

Despite the scant attention given by academic research to these issues, there is an intense debate in the policy arena regarding the rules on who should own and operate storage facilities. In many jurisdictions, storage is considered a generation asset, which essentially bars system operators from owning and operating storage devices due to unbundling restrictions.²⁴ Yet, our analysis suggests that regulators should not put the spotlight on the integration between transmission and storage (which could potentially be positive),²⁵ but rather on the integration between generation and storage, as well as on the concentration in storage ownership. A vertically integrated firm or a large storage owner internalizes the price effects caused by storage on its own energy sale and purchase decisions. This causes them to distort the use of storage away from the cost-minimizing pattern, reducing its profitability, and thus weakening the firms' incentives to invest.

Throughout the analysis, we have assumed that storage owners are exposed to wholesale electricity prices. While this is generally true for large storage installations (e.g. pumped hydro), it need not be so for the distributed storage facilities (e.g. electric vehicles, or behind-the-meter batteries). In order to fully develop the potential of storage, it is paramount to foster dynamic electricity prices and time-of-use tariffs so that storage owners internalize its

²⁴For instance, in May 2019, the European Commission ruled that only under exceptional circumstances are transmission and distribution operators allowed to own and operate storage facilities (European Commission (2019)). Similarly, in 2019, China decided not to allow network operators to include storage costs in their fees, which led to a sharp decline in storage investment. Yet, other jurisdictions (such as Australia or Chile) allow network operators to own and operate storage assets under certain conditions. And in the US, the debate is still on-going as regulators are currently reviewing the storage ownership rules, which differ widely across states (European Commission (2019, 2020)).

²⁵For instance, that is the case if the transmission owner is required to operate storage so as to reduce system costs, just as a social planner would do. Also, storage in the hands of the System Operator could contribute to security of supply.

social benefits.

Chapter 3

Normative Analysis of Environmental Policy with Endogenous Preferences

What is the optimal level of taxes for polluting activities? Should we subsidize public transport in order to reduce private car usage, or should we instead promote electric vehicle adoption? Is it better to mandate recycling or to set a tax on waste disposal? More generally, which policies achieve the transition to a low carbon economy at the least cost for society?

These and other related questions are concerned with fixing some type of environmental externality. The traditional economic approach to address these externalities consists in devising policies that align private and social incentives by altering the costs of some targeted behavior. For instance, the classic instrument to reduce polluting emissions involves introducing carbon pricing that makes economic agents internalize the harm that their decisions are imposing on others. In this process of policy design, optimality is defined with respect to some exogenous criterion that generally takes into account people's wants. For example, if most people value private automobiles, policy design should put a higher weight on policies that favor this transport option as compared to non-motorized or public alternatives. In general, standard normative analysis in environmental economics is based on two premises. The first is that policies and regulations alter behavior by only changing individuals' choice set and not by altering their preferences, which are assumed to be stable and ex-

ogenous.¹ The second is that the notion of optimality should reflect the desires of the population, implying that a policy is superior to another if individual preferences are satisfied to a larger extent. This “preference satisfaction” approach in cost-benefit analysis has been championed on the grounds of being neutral among different visions of what is good for society and has proved successful in making policy recommendations across a wide range of issues.

However, the simplifying assumption of constant and exogenous preferences does not always hold. In many cases the preferences, values and norms of significant swaths of the population who are affected by a policy are also shaped by that policy. Examples of this possibility abound in the context of environmental policies and regulations, ranging from good public transport options crowding-out preferences for private transport alternatives (e.g., Weinberger and Goetzke (2010)) to health policy inducing preferences towards environment-friendly dietary choices (e.g., Birch (1999)). In general, many environmental policies directly or indirectly affect the values by which these same policies are judged, and this is more likely the longer the time scale, which makes it especially relevant for environmental processes that take place over the course of several generations. However, despite its pervasiveness, preference formation is rarely incorporated into cost-benefit analyses of climate policy.²

On the theory side, the importance of preference evolution for several socio-economic outcomes has stimulated a recent surge in the study of the endogenous formation of preferences from a positive perspective.³ However, these theoretical approaches to endogenous preferences rarely consider normative implications and they generally refrain from doing welfare analysis.⁴ Beyond the philosophical difficulties of judging whether some values and preferences are more socially desirable than others, the main reason behind this reluc-

¹In this paper we are concerned with doing welfare analysis in situations in which policies alter individual preferences, which is different from altering beliefs, with the latter defined as the subjective probabilities about the realizations of different states of the world. In this definition, policies that affect individual beliefs do not challenge standard normative analysis.

²A recent exception is Mattauch et al. (2018), who analyze the efficiency of different environmental policies when these policies affect some parameters of the consumers’ utility function.

³See for instance Akerlof and Kranton (2000), Becker and Murphy (1988), Tabellini (2008), Bisin and Verdier (2011), Alger and Weibull (2019), Bernheim et al. (2021), Bowles (1998), Güth (1995) and Besley and Persson (2019a).

⁴Some notable exceptions are Bowles and Hwang (2008) and Bezin (2015), but their analysis is limited to narrow welfare concepts.

tance lies in the well-known conceptual difficulties for the standard welfarist approach of dealing with situations in which preferences change. Welfarism is the normative criteria which states that social welfare is only a function of individual utilities (Sen (1979)). Therefore, when preferences are not stable, the measuring rod by which a policy is evaluated is shaped by the policy itself, which makes the task of establishing the best social policy quite ambiguous and raises several questions: should the impact of the policy be calculated with respect to pre-policy preferences, or instead should we use the preferences induced by the policy? Maybe we should rely on some combination of both?⁵ These shortcomings of the welfarist approach may explain the traditional assumption in economics about the exogeneity of preferences, but it would be misleading to proceed as if preferences are exogenous in cases when they are not. Then, should we walk away from welfarist approaches and rely on more paternalistic approaches, such as meeting certain desirable exogenous goals (e.g., emissions targets)? Or should we instead consider other methods that emphasize the process by which environmental problems came into being (e.g., who has the burden of decarbonization)?

The goal of this essay is to explore these questions and to assess how the main normative frameworks used in economics can be adapted to contexts in which policies shape preference formation. Any of the major positions in welfare economics is nothing but an attempt to answer the following question: what is the objective that policy should try to achieve? In general, one can distinguish three different ways of answering this question. First, as I have mentioned before, the most common approach in economics is to assume that policies should satisfy individual preferences. The second normative position also focuses on the consequences of different policies, but changes the source of welfare criteria from individual desires to substantive concepts of well-being specified by the analyst, such as overall happiness or scientific-based objectives (e.g., reducing global temperatures or preserving biodiversity). The third major line of thought puts the spotlight on processes rather than on outcomes, claiming that policy should be decided attending at how the status quo came about, the fairness of alternative rules and procedures for making social decisions, or the ability of the policy to maximize individual capabilities.⁶ In Sections 1,

⁵Moreover, time consistency problems arise, as the optimal policy for initially prevailing preferences may bring about new preferences for which the policy implemented becomes sub-optimal.

⁶One should keep in mind that these classification lines are sometimes blurred, although this taxonomy is useful to guide the discussion.

2 and 3 I examine whether these general welfare concepts can be applied to cases in which preferences endogenously respond to policy and I discuss some potential solutions that make them suitable for this purpose. While some of these approaches provide some useful guidance on how to conduct normative analysis, our review suggests that none of them completely solve the problems posed by endogenous preferences. However, a combination of them seem promising for dealing with environmental issues, in which actual policy practice combines more standard views that aim to satisfy individual preferences with more paternalistic and non-consequentialist approaches. Overall, while no general framework can be provided for dealing with situations in which policy changes preferences, the present essay highlights the importance of being explicit about the assumptions that must be made when making normative statements, an unavoidable task in the context of environmental policy.

Finally, in order to illustrate the implications of endogenous preferences for optimal environmental policy design, in the last section of the paper we sketch out a very simple model of environmental externalities and carbon pricing. In this model, some citizens (the environmentalist) have three distinctive characteristics: they attach a lower utility to their own consumption of polluting goods, they experience a higher negative environmental externality from aggregate pollution, and they derive additional utility from “green” goods due to social esteem effects. In this model, carbon taxes have two effects. First, they directly affect consumption behavior by changing the relative price of clean and polluting goods. Second, they indirectly crowd-in environmentalist preferences by increasing their relative utility payoff.⁷ In this setting, we derive the optimal carbon tax in the light of alternative normative criteria. The results show strikingly large differences in optimal tax rates, highlighting the importance of checking the robustness of environmental policy recommendations to different welfare criteria. Moreover, by being explicit about the dynamics of preferences, our model allows us to devise a method to micro-found the relative weights that welfare analyses could give to pre-policy, post-policy and intermediate preferences. This method partially addresses some of the shortcomings of the welfarist approach while keeping its essence, and its implementation is feasible as it relies on empirically observable parameters. Our model also illustrates the sizable differences in optimal policies that arise when comparing a myopic social planner against a forward-looking one that internalizes the endogeneity

⁷However, as more people adopt environmentalist preferences, the social esteem value from being environmentalist is reduced, putting a limit to the crowd-in effect.

of preferences.

This essay is mostly related to papers that address the challenges posed by preference endogeneity for normative work. Some of them will be discussed in the following sections, but I would like to mention the early contributions of Gintis (1974), Harsanyi (1953), Yaari (1977), Von Weizsäcker (1971), Peleg and Yaari (1973) and Goldman (1979). The modest contribution of this essay consists in updating some of these contributions and exploring how they relate to recent theory work on endogenous preferences, which up to this point has generally avoided normative analysis. Environmental policy serves as a guiding thread for the discussion, as it is an area in which normative analysis is unavoidable and preference endogeneity is likely to be important. In fact, several recent empirical papers point in this direction. First, carbon taxes and subsidies can increase or decrease the intrinsic motivation to take environment-friendly actions. For instance, Frey and Oberholzer-Gee (1997) find that offering money to accept a nuclear waste repository substantially reduce the willingness of residents in the area to accept it. In the context of dietary choices, Lanz et al. (2018) find that setting a carbon tax on food products with embodied carbon emissions reduces voluntary substitution towards these products. In contrast, other authors find crowding-in effects of carbon taxes when applied to gasoline prices (Rivers and Schaufele (2015)) or plastic bag usage (Convery et al. (2007)). Second, environment-related behaviors are influenced by transport policy and urban design (Schwanen et al. (2012); Weinberger and Goetzke (2010); Weinberger and Goetzke (2010)), public health measures related to dietary choices (Birch (1999); Hawkes et al. (2015)) and energy efficiency policies (e.g., Ito et al. (2018); Costa and Gerard (2015)). Third, there is evidence that climate change mitigation policies can also *directly* shape preferences over those same policies. For instance, Prieur and Bréchet (2013) suggest that investments in education that increase awareness of environmental issues increase the demand for those policies. Similarly, Eliasson (2008) shows that introducing a traffic congestion charging trial increased the public support for implementing the policy. Finally, several policies have been shown to influence preference traits directly related to environment-friendly behaviors, such as altruism (e.g., Byambadalai et al. (2019)) or time-discounting (e.f. Perez-Arce (2017)).

3.1 Welfarist approach: preference satisfaction

The most widely used normative approach in economics rests on two fundamental premises: that policies should only be evaluated by their effects on human welfare, and that individuals are the best judges of what constitutes a good outcome.⁸ Thus, when a policy-maker must decide between different policies, she analyzes how each policy will impact the welfare of every citizen concerned by the policy, as measured by their revealed *willingness to pay* to obtain the ensuing benefits or to avoid the resulting costs. Then these individual gains and losses are aggregated into a general assessment of the effects created by the policy.⁹

An additional fundamental premise which is sometimes overlooked is that preferences are stable and independent of the effects of the policy. While this is a reasonable assumption in several contexts, the examples in the previous section point out to cases in which policies inevitably forge the preferences of the individuals affected by the policy. In those cases, how do we tell which policies satisfy people's preferences to a larger extent, if these policies are shaping themselves what people want? In more practical terms, it is not clear whether the willingness to pay for different policy options should be measured with respect to ex-ante preferences, to ex-post preferences or perhaps to some combination of both (Harsanyi (1953)).¹⁰ Clearly, in several cases policy assessments will not necessarily yield the same ranking over policies depending on whether we calculate welfare with regard to pre-policy or post-policy preferences.¹¹ Moreover, choosing one policy over another presents additional ethical dilemmas, as it unavoidably implies privileging some preferences over others, a task which economists are usually reluctant to address due to its pa-

⁸See Kaplow and Shavell (2000) and Ng (2000) and for very good treatments of welfarism applied to economics.

⁹In most cases, the analyst relies on the Pareto principle (or on the less strict Kaldor-Hicks principle) as the method for aggregating individual measures of willingness to pay. Other widely used approach uses Social Welfare Functions (SWF). For interesting reviews of different methodologies see Adler (2016).

¹⁰A related challenge for normative analysis is posed by the existence of behavioral biases and other departures from the standard framework of rational choice (e.g., Kahneman (2011)) and is the subject of several works (e.g., Thaler and Sunstein (2003), Bernheim and Rangel (2009)). Although some issues are closely related, this literature does not directly address the problems posed by preference endogeneity.

¹¹Note that in dynamic settings changes in preferences give rise to the well-known time-inconsistency problem (e.g., Loewenstein and Prelec (1992)). The reason is that a policy that was optimal for some preferences may induce a new set of preferences under which the policy implemented is considered sub-optimal.

terrestrial connotations. Do these problems completely invalidate the welfarist approach, as some authors have suggested (e.g., Thaler and Sunstein (2003); Dold (2018)), or can we come up with modifications that allow us to preserve the advantages of the preference satisfaction approach?¹²

As I will try to argue, I believe that we can be mildly optimistic about the ability of the welfarist approach to deal with preference endogeneity. We should first note that, even in cases when preferences are exogenous and policies only affect the choice set of individuals, doing welfare analysis in its purest form (i.e., Pareto principle) presents significant problems, as any policy change requires unanimous support by the affected population. When individuals are heterogeneous along some dimension (for instance, because they have different preferences over environmental amenities) unanimity is very uncommon, as most policy reforms generate winners and losers.¹³ In more concrete terms: if policy A is preferred to policy B by person x and the opposite is true for person y , then which option should be considered socially optimal? The traditional way of solving this problem is by aggregating individual preferences to generate a *social* preference ordering over the two policies (usually through a Bergson-Samuelson social welfare function constructed as a weighted sum of utilities).¹⁴ As a result, the problem is pushed to another level, in which the analyst must decide which weight must be given to different individuals/preferences. Although there is an ongoing debate in welfare economics about how to justify and calculate these weights, in the last decades several interesting proposals have been put forward that provide a compelling rationale for inter-personal comparisons (see Fleurbaey and Abi-Rafeh (2016) for an overview of different approaches).

Note that the problem of assessing how *individual* welfare changes when preferences change is similar to the one of making inter-personal utility comparisons. If, for instance, the status-quo policy A is preferred to policy B when the individual has status-quo preferences x , and the opposite occurs when policy B is implemented and induces preferences y , then which option should

¹²Among others, the two main advantages proposed by advocates of this approach are its avoidance of paternalism and its respect for individual autonomy (e.g., Kaplow and Shavell (2000)).

¹³This is especially true in the context of environmental policy, as recently exemplified by the *yellow-vests* movement.

¹⁴Despite Arrow's famous impossibility result (Arrow (1950)), a lot of work has found ways to compare utilities, usually by giving a comparing the *intensity* of preferences. There are other approaches that do not rely on this cardinal interpretation of utilities, such as relying on equivalent and compensating variations (Fleurbaey and Tadenuma (2014)).

be considered better by the analyst?¹⁵ Following the reasoning in the previous paragraph, we can glimpse a potential solution consisting in considering an individual who changes preferences as two different individuals, and do comparisons of utilities among different selves in the same vein that we do inter-personal comparisons of utility.¹⁶

But then, which normative weights should be given to different preferences? Unfortunately, choosing weights for pre-policy and post-policy preferences is even more problematic than choosing weights for interpersonal utility comparisons. On the one hand, in some cases giving a relatively large weight to pre-policy preferences is tantamount to justifying an undesirable status-quo. This would occur when preferences are *adaptive* (Elster (1982); von Weizsäcker (2005)), which means that people tend to like what they can get in the socio-economic environment in which they live.¹⁷ This type of preferences may generate a circularity by which the analyst would reject a policy change because it is not preferred by the citizens, but that would occur because the status-quo policy has shaped preferences in its own favor. This could be especially relevant in the case of environmental policy. For example, large fractions of the population may not have experienced the value of living in non-polluted cities while they have experienced the benefits of having private cars. However, they would prefer cleaner cities if they had experienced them. In that case, would it be reasonable to justify an environmentally unconcerned society? On the other hand, putting instead a large weight on post-policy preferences may end up being notoriously paternalistic, in the sense of using for valuation purposes some ideal hypothetical preferences that the analyst regards as superior. For instance, should policy makers increase fuel taxes, imposing economic costs on large swaths of current populations, in the belief that they will come to value those taxes in the future? As implied by the question, this danger seems particularly relevant in real-world situations in which policy makers must estimate ex-post preferences for which no data is available.

¹⁵Of course, if a policy change is preferred under both ex-ante and ex-post policy preferences, the dilemma does not arise, as implementing the policy would yield a Pareto improvement. However, this “value-free” criterion would be indecisive in many instances.

¹⁶By analogy to the case of inter-personal comparisons, some technical details must be satisfied in order to do *intra-personal* comparisons of utility by means of a social welfare function. These include requiring the person’s preferences across different states of the world to be separable (Fleming (1952)), as well as requiring different selves to satisfy the standard preference axioms.

¹⁷This would be obtained in a large class of models of endogenous preferences, including evolutionary ones, which are the most widely used in the literature (e.g., Alger and Weibull (2019); Tabellini (2008)). One example is the model presented in Section 4 of this paper.

Some authors such as von Weizsäcker (2005) argue that under some circumstances policies should be evaluated with respect to post-policy preferences (i.e., only ex-post preferences have a positive weight). Those circumstances involve having *adaptive* preferences which, in a more concrete definition than the one provided above, are those preferences which have the property that if policy B is preferred to policy A under the preferences induced by policy A , then B is also preferred to policy A when the preferences are those induced by A . In more simple words, adaptive preferences imply that individuals generally resist change because preferences adapt to circumstances. In those cases, it would be possible to argue that a policy change is an improvement if and only if the new policy is preferred under the induced preferences (because it avoids “regret”).¹⁸ This logic has wide appeal for a large class of models that exhibit adaptive preferences. Unfortunately, it is of no use when preferences are not adaptive.

In contrast, other authors defend using only pre-policy preferences. Some of them implicitly, by just assuming that preferences are exogenous. A prominent example is Becker and Stigler (1977), who famously argued that “*all human behavior can be viewed as involving participants who maximize their utility from a stable set of preferences*”. Others, especially from the Postmodernist tradition in sociology (e.g., Harvey et al. (1990)), but also in economics (e.g., Galbraith (1998)), argue that evaluation should be done on the basis of pre-policy preferences due to the “manipulative” effect of policies on “natural” preferences.¹⁹

A more balanced method would entail given positive weight to pre-policy and post-policy preferences. To the best of my knowledge, we lack any consistent framework that guides the choice of these weights according to sound philosophical and ethical principles.²⁰ Although this type of framework is very much needed, in the meantime we can devise some approach that allows us to establish weights in an explicit and empirically appealing manner. A possi-

¹⁸Importantly, according to Von Weizsacker, with adaptive preferences choices would satisfy the strong axiom of revealed preferences. Note, however, that Von Weizsacker mostly has in mind models of habit consumption and evaluated different consumption bundles instead of different policies, but the logic can be adapted.

¹⁹See Benhabib and Bisin (2011) for a nice overview of this literature from the point of view of economics.

²⁰This sharply contrasts with the case of interpersonal comparisons of utility, for which several welfare criteria have been proposed in order to choose weights (See Fleurbaey and Abi-Rafah (2016) for a survey). In the context of redistribution policies, the most famous approach is the one proposed by Harsanyi (1955), based on expected utility theory.

bility that we illustrate in Section 3.4.2 of this paper is to derive the weights from an explicit dynamic model of preference formation. In general, the problems that endogenous preferences bring to the welfarist approach come from thinking about it in static frameworks, as in this case any choice of weights is necessarily arbitrary. In contrast, by putting a dynamic structure on the way preferences change (as captured by a law of motion) one can consider the effect of policy on initial and final preferences, as well as any intermediate preferences that arise during the adjustment process. Then, these varying preferences can be aggregated and time-discounted in an objective function. With this approach, the weights are justified through the model and can be estimated from empirical parameters such as the speed of preference change or the discount factor. Moreover, the method takes into account intermediate preferences in the process of adjustment and, once preference adjustment is over, it values the policy in accordance to the preferences that will be experiencing it in the long-run. Overall, the main goal of this approach is to put experienced utility as the key criteria by which policies should be evaluated, which can be seen as the ultimate goal of the preference satisfaction approach. Thus, if preference change is sluggish, the approach implies that a large weight is attached to initial preferences, whereas the opposite occurs when preferences change fast.

Nevertheless, this approach leaves several crucial issues unresolved and presents some problems. The first relates to the fact that it may not capture how individuals see themselves “choosing” preferences over time (i.e., do individuals evaluate outcomes by discounting future preferences?). Second, the choice of weights may be very sensitive to the specific model used to describe the dynamics of preferences (and, specially, to the speed of convergence of preferences towards the steady state).²¹ In that case, a question arises as to what the right model is to be used depending on the particular case at hand. Third, the results are likely to be very sensitive to the choice of discount factor, which itself is an endogenous preference parameter (Arrondel (2009); Robson and Samuelson (2009)). However, in the absence of other approaches to choose weights, this has the key advantages of being explicit about the assumptions and of relying on measures that in principle can be empirically identified. Moreover, it avoids choosing weights according to exogenous paternalistic criteria, inhibiting the analyst from making value judgements.

²¹Moreover, in some models preferences may not converge to a steady state.

Another alternative approach to keep the satisfaction approach when policies shape preferences was proposed by Becker (1998), who argues that we can redefine any situation in which policies affect preferences as a situation in which preferences remain unchanged, but the policy has an additional impact on some capital stocks. In other words, Becker proposes to rely on stable *meta-preferences* over the varying preferences that an individual may have, which could be formalized by means of an *extended* utility function. This utility function contains two additional terms, which are the stocks of personal capital and social capital, that are accumulated by past experiences and by the effect of the actions of others in the past. In turn, these capital stocks shift the utility from goods or from other circumstances an individual may face. In this view, any policy effect on preferences can be redefined as an effect on the level of these capital stocks. These leaves meta-preferences unchanged, which allows normative analysis to be conducted with respect to those meta-preferences. This is a creative way of preserving the preference satisfaction approach, although it does not fully solve the problem on how to conduct normative analysis, as it only reframes the important question: should policies be evaluated with respect to the initial or to the final stocks of personal and social capital?²²

A closely related approach is to make a distinction between fundamental preferences that remain unchanged and “external” preferences that are a function of the policies implemented (Dietrich and List (2013)). For instance, in the context of climate change mitigation policy, external preferences over dietary choices may depend on the context in which they are made and will usually depend on the policies in place (availability of vegetarian dishes, taxes on meat consumption, etc.), as in many cases food tastes are endogenous to past consumption (Becker et al. (1991)). However, one could argue that the fundamental preferences related to dietary choices may include a desire for long-run health, environmental sustainability and respect for animal rights. Clearly, the existence of fundamental preferences that are stable and exogenous solves the problems posed by endogenous preferences: the analyst has to identify what are those fundamental preferences and implement the policies that satisfy them. However, this approach presents difficulties. First, those fundamental preferences may be hard to identify, as they are unlikely to be

²²Moreover, this idea presents some practical difficulties. Among them, it stands out the problem of how the analyst measure these capital stocks. While measuring preferences is already a daunting task, these hypothetical capital stocks seem more difficult to elicit from surveys or individual choices.

revealed by choices that also depend on the superficial preferences. Thus, if the analyst cannot ascertain the content of those fundamental preferences, she might be tempted to specify them herself in a paternalistic way, which is precisely what the preference satisfaction criterion wants to avoid. Second, fundamental preferences may not exist over long-time scales. For instance, candidates for fundamental preferences in the context of environmental policy are the degrees of patience, risk-aversion, and altruism, as they determine optimal carbon pricing and decarbonization paths. However, these preferences are culturally determined and they change over time (Alger and Weibull (2019); Dohmen et al. (2012)). In this case, given that environmental policies operate at large time scales, it may be misleading to assume that policies would not affect these preferences, leaving the problem posed by endogenous preferences unresolved.

To sum up, this section has shown that relying on the preference satisfaction criterion to conduct normative analysis comes with several conceptual challenges. Some partial solutions have been proposed, which taken together may call for preserving this approach. However, we can take an alternative stand and reject the whole project of making welfare analysis in terms of satisfying individual preferences. The next two sections examine how other normative criteria may help to evaluate policies in the presence of endogenous preferences.

3.2 Paternalism and substantive normative criteria

The second stance on welfare criteria accepts that individual welfare matters but considers it to be of secondary importance as compared to other objectives, such as liberty, happiness, non-exploitation or other concepts of the good life (e.g., Layard (2011); Gul and Pesendorfer (2007); Frey et al. (2010)). In the context of environmental policies, alternative goals to individual welfare include air quality, preservation of biodiversity, sustainability of ecosystems or, in more prosaic and practical terms, policy goals such as targets for carbon emissions or aggregate meat consumption. Although these goals clearly contribute to individual welfare, the key difference of this approach is to regard these principles as important in their own right, beyond their effects on

individuals utilities.²³

Advocates of this approach generally question the assumption that individuals know what is best for them, as several biases permeate individual behavior (e.g., Thaler and Sunstein (2003); Loewenstein and Haisley (2007)). As a result, they argue that eliciting individual preferences is doomed, which increases the appeal to rely on objective measures of what constitutes a good outcome. The case in favor of substantive approaches is further strengthened when preferences can be shaped by policies. First, in those cases founding normative analysis on satisfying individual preferences presents all the shortcomings that we have discussed up to this point. In addition, some preferences may be considered better in terms of their ethical virtues, while others may be considered undesirable or morally repugnant. For instance, in the context of climate policy, preferences that place higher weight on the future or those that imply some degree of altruism towards future generations may be considered as superior to those that exhibit impatience or selfishness.

There are other considerations in favor of substantive approaches. An important one is that evaluating policies with respect to some exogenous target considerably simplifies welfare evaluation when preferences are endogenous. Consider the example of designing carbon taxes. If the goal of those taxes is to limit the rise of global temperatures to 2C (or, equivalently, to limit carbon emissions to a certain amount), then the social planner just needs to consider the direct effect on individual incentives, plus the potential crowding-in or crowding-out effects of carbon taxes on preferences (Mattauch et al. (2018)). However, whether policy shapes preferences or not does not affect the criterion by which the policy is evaluated. In other words, failing to consider the endogeneity of preferences may conduct to inefficient carbon taxes but does not complicate normative analysis once the effects of carbon taxes on preferences are considered. Moreover, from a more applied perspective, the policy maker avoids several issues encountered in eliciting preferences from observable choices.²⁴ With substantive welfare criteria the policy maker merely has to find a trustworthy way of measuring of the chosen goal and calculate how differ-

²³In formal models, it implies that these motives enter directly in the objective function of the social planner, in contrast to the welfarist approach in which they only enter indirectly through individual utilities.

²⁴According to several authors, behavioral biases permeate individual behavior, which complicates eliciting preferences from observed behavior (e.g., Kahneman (2011)). However, some solutions have been proposed (Adams et al. (2015)).

ent policies affect that metric.²⁵ These two advantages could provide a strong rationale for explaining actual policy practice in environmental issues, which is generally conducted to meet certain scientifically informed environmental targets.

The other side of the coin is that relying on substantive objectives undermines individual autonomy and can be excessively paternalistic. The conception of the good life specified by the policy maker may differ substantially from what individuals judge to be better for them. In fact, the idea of what constitutes a good life varies from person to person and, more importantly, every citizen has privileged knowledge about what her own desires and ambitions (e.g., Ng (2000); Kaplow and Shavell (2000)). The danger is that once we start choosing what's best for others, it becomes difficult to avoid setting policy to make people want whatever we think they should want instead of setting policy to assist people in achieving their own goals.

Nevertheless, further support for substantive approaches to welfare comes from the inability of preference satisfaction criteria to avoid some form of paternalism when preferences are endogenous to policy choice. When policies must be implemented and these policies affect individual preferences, the social planner inevitably privileges the development of some preferences over others, irrespective of whether it relies on ex-ante or ex-post preferences for policy evaluation. Similarly, relying on meta-preferences or fundamental preferences in actual policy practice when these are not observable may imply choosing some particular goal that the analyst considers to be desirable, even if this goal is not grounded on individual utilities. Then, if some degree of paternalism is inescapable, why not accepting it as a fact and discuss which substantive criteria should be maximized? Why not promoting the preferences that we consider socially desirable? In fact, the primary goal of some widely used policies is to promote certain preferences and values (think of the schooling system and of awareness-raising campaigns). Then, why not doing the same with policies that are primarily meant to change the choice set but also indirectly change preferences?

The main problem with this approach was pointed out in a famous article by Kaplow and Shavell (2001), who demonstrated that any method of pol-

²⁵In fact, a large literature in economics and social psychology has developed reliable measures of subjective well-being (such as happiness and life satisfaction), who are advocated as better reflecting individual welfare than preferences implied by observable choices. See for instance Frey et al. (2010) or Layard (2011).

icy assessment that does not depend *solely* on individual utilities violates the Pareto principle, meaning that the criterion will sometimes require the implementation of a policy measure that makes every individual worse-off. To see the astonishing implications of this result, consider a situation in which some individuals want a cleaner environment, while others are more concerned about other things such as greater economic growth or lower gas prices. In this case, their result implies that we will have circumstances in which a carbon price that is implemented with regard to some exogenous goal (such as an emissions target) will make environmentally unconcerned citizens worse-off but, surprisingly, also those that want a cleaner environment.

As with the preference satisfaction approach, substantive criteria present both advantages and disadvantages to deal with preference endogeneity. Moreover, the discussion points out that the boundary between the two approaches is blurred, as preference endogeneity inevitably implies some form of paternalism. This has led some authors to propose hybrid frameworks, in which the goal of satisfying individual utilities is combined with other objectives and moral principles. This is a promising approach which in fact seems to be guiding policy makers in designing actual environmental policies.

3.3 Non-consequentialist approaches

The previous welfare criteria have one important common aspect, namely, that the goodness of a policy is solely evaluated with respect to its consequences. However, critics of these approaches object that both the motives and the processes that lead to a consequence should not be completely ignored. Within this category, there are a bunch of approaches that in many cases bear little resemblance among themselves. Here I will briefly discuss those that are more relevant to environmental policy design (without aiming to be exhaustive).

Among these, the one that has had a greater reception within economics is the capabilities approach (Sen et al. (1999); Nussbaum (2009)). The main assertion of this line of thought is that the key metric by which policies should be assessed are “capabilities”, which are the combinations of various things that a person may value and can reach if given the freedom to do so,²⁶ such as having a clean environment or participating in the deliberative process for de-

²⁶ “Functionings” in the language of Sen et al. (1999).

ciding environmental policies. Then, welfare is the aggregation of the different valuations that people have over capabilities. As a result, instead of satisfying individual preferences or maximizing goals such as environmental sustainability, policy should be concerned with enlarging the ability of citizens to select among different goals. The main objective of this approach is to avoid taking directly substantive criteria, as these violates individual freedom of choice and the neutrality with respect to different goals. Unfortunately, this approach presents two important problems: the first is that the concept is difficult to translate into real policy practice, as it is difficult to translate capabilities into a unidimensional metric (Fleurbaey and Blanchet (2013)). Moreover, this theory also fails to satisfy the Pareto principle, because capabilities are not measured taking into account individual preferences (Kaplow and Shavell (2001)).

Another normative principle that had an important role in classical political economy, and that has recently been revived in optimal tax design (Weinzierl (2014)) is the principle of “equal sacrifice”. This view maintains that every member of society must make the same sacrifice for the common good. While it is surprisingly unexplored in environmental economics (at least explicitly), this principle could be easily adapted to climate change mitigation policy. In fact, this criterion strongly resonates with several proposals that defend that the burden of reducing polluting emissions falls upon developed countries, as reducing polluting emissions for under-developed countries would come together with their failure in achieving the same levels of development that their rich counterparts. A closely related view focuses on responsibilities and advocates that the burden of decarbonization falls upon those countries and individuals that have contributed most towards global warming (e.g., Neumayer (2000)).

Finally, following a tradition initiated by Buchanan (1979), other authors such as Sunstein (1993) entirely call into question standard welfare economics, as, according to them, it is based on the false premise that we can ascertain individual preferences at any point in time. In other words, well-defined preferences are a purely theoretical construct that has no close counterpart in reality. While in the case of Sunstein the way around this problem is to embrace a particular form of (libertarian) paternalism (Thaler and Sunstein (2003)), other authors have questioned this view, arguing instead for adopting a “contractarian” viewpoint which, instead of focusing on policy consequences, defends evaluating the fairness and quality of alternative rules through which policies are deliberated and implemented (Dold (2018)).

3.4 Normative analysis in a model of endogenous environmentalist preferences

In this section I sketch out a very simple model of “green” preferences, environmental externalities and carbon pricing in order to: i) clarify the previous discussion about the different possibilities on how to design optimal policy in the presence of endogenous preferences; ii) propose a reasonable methodology for preserving as far as possible the preference satisfaction approach and iii) illustrate the effect of different welfare criteria on optimal carbon pricing policies.

3.4.1 Model set-up

We consider an economy populated by a mass 1 of agents.²⁷ Agents can be of two types: a fraction q of environmentalists, labelled by e , and a fraction $1 - q$ of non-environmentalists, denoted by n . The utilities of both types are given by:

$$U^e(c, x, C, X) = (1 - \alpha) \ln c + x + \sigma(x - X) - (\psi + \phi)C$$

$$U^n(c, x, C, X) = \ln c + x - \psi C$$

where c denotes the consumption of polluting goods, x in the consumption of clean goods and C and X are the total (per-capita average) consumption of polluting and clean goods, respectively. Both preference types value consumption of each good, although environmentalists experience a lower utility from consuming polluting goods, as captured by parameter $\alpha \in (0, 1)$. However, clean goods constitute a status good that provides utility to environmentalists, and the more so the higher their consumption of these goods with respect to the average consumption in the population.²⁸ $\sigma > 0$ is a parameter that cap-

²⁷This model is a simplified version of the one in Andrés-Cerezo (2021), who explores the role of environmental activists in promoting decarbonization. A similar model can be found Besley and Persson (2019b), who focus in the political economy aspects of environmental policies.

²⁸I assume that environmental preferences have a pro-social preference component, which implies an intrinsic motivation to have a good behavior. This can come from having a higher status in society or from some form of “warm-glow” associated to behavior. One could also assume that the social esteem motive is increasing in the aggregate consumption of the

tures the strength of the status motive. The last two terms of each expression capture the disutility from aggregate pollution (with $\psi, \phi > 0$), implying that environmentalist experience higher disutility from environmental externalities. At time t , the budget constraint for both types is given by:

$$y + s = (1 + \tau)c + px$$

where $p \geq 1$ denotes the market price of the polluting good (the price of the other good is normalized to 1). $\tau \in [0, \bar{\tau})$ is a consumption tax on the polluting good, y denotes per-capita income (equal for all) and s is a lump-sum government transfer. Thus, the budget constraint of the government is given by $s = \tau C$.

Individuals of each type choose a consumption bundle, taking their effect on aggregate consumption as given (i.e., each consumer is small enough to ignore the effect of her consumption of environmental externalities). Optimal consumption choices at time t are given by:²⁹

$$\begin{aligned} c^e &= \frac{p}{1 + \tau} \frac{1 - \alpha}{1 + \sigma} \\ c^n &= \frac{p}{1 + \tau} \\ x^e &= \frac{y + s}{p} - \frac{1 - \alpha}{1 + \sigma} \\ x^n &= \frac{y + s}{p} - 1 \end{aligned}$$

and:

$$\begin{aligned} C &= qc^e + (1 - q)c^n = \frac{p}{1 + \tau} \left[1 - q + q \frac{1 - \alpha}{1 + \sigma} \right] \\ X &= qx^e + (1 - q)x^n = \frac{y + s}{p} - \left[1 - q + q \frac{1 - \alpha}{1 + \sigma} \right]. \end{aligned}$$

From the government's budget constraint, we have that:

polluting good. For the purposes of this paper, this distinction does not play an important role.

²⁹For y sufficiently large interior solutions are guaranteed.

$$s = \tau C = \frac{\tau p}{1 + \tau} \left[1 - q + q \frac{1 - \alpha}{1 + \sigma} \right].$$

In these expressions we can note the standard way in which carbon taxes affect individual behavior, consisting in favoring the consumption of clean goods by lowering their relative price.

In the literature on environmental externalities and carbon pricing, optimal policies are generally defined as those that maximize individual utilities, assuming that individual welfare is given by an exogenously specified utility function and that preferences of every individual in the population are identical. In this case, since preferences differ across individuals, the standard approach would usually adhere to the standard Pareto principle, defining optimal policies as those that cannot make any individual better off without making at least one individual worse-off.³⁰ If utility is given a cardinal interpretation and inter-personal utility comparisons are allowed, optimal policy can be defined as the maximum of a social welfare function in which individual utilities are aggregated, possibly with different weights attached to each individual. As it has become clear at this point, things are further complicated when preferences change in response to policies. We now turn to model how one could allow preferences to change in the current set-up.

Preference dynamics

The model has been defined in purely static terms. In these static environments, one could model preference change by exogenously specifying how policy instruments affect some parameters of the utility function (such as σ or α) or by changing the proportion q of each type of agent in the population. However, as I have argued before, being explicit about how preferences change over time in response to policy will help us to devise a method for doing welfare analysis in the presence of endogenous preferences. In this example we model preference change by using a standard model of cultural evolution, in which the fraction of environmentalists q changes over time as function of relative utility payoffs.³¹ Assuming a stationary population size, equal model structure in

³⁰Other authors use the similar but slightly less demanding Kaldor-Hicks efficiency criterion, by which agents that are made better off could compensate those that become worse off.

³¹These models are widely used in the literature on endogenous preferences (see for instance Tabellini (2008), Bisin and Verdier (2011) or Alger and Weibull (2019).) and

every period, and that the tax rate is chosen once and for all (mimicking what happens in static models), the typical equation describing preference dynamics takes the following form:³²

$$q_{t+1} - q_t = \theta q_t (1 - q_t) [U_t^e(q_t, \tau) - U_t^n(q_t, \tau)]$$

where θ is a parameter that affects the rate of change of preferences. We have that:

$$\begin{aligned} U_t^e(q_t, \tau) - U_t^n(q_t, \tau) &= (1 - \alpha) \ln \left(\frac{1 - \alpha}{1 + \sigma} \right) - \alpha \ln \left(\frac{p}{1 + \tau} \right) \\ &+ \frac{\alpha + \sigma}{1 + \sigma} [1 + \sigma(1 - q_t)] - \phi \frac{p}{1 + \tau} \left[1 - q_t + q_t \frac{1 - \alpha}{1 + \sigma} \right] \end{aligned}$$

On top of affecting individual consumption choices by changing relative prices, carbon taxes affect preference adoption by changing the relative payoffs of having one set of preferences over the other. In the expression above we can see that an increase in carbon taxes crowds-in environmental preferences, as the tax makes more attractive to consume clean goods and, as a result, to have preferences that value more those goods. This effect is additionally strengthened by the reduction in the additional harm that environmentalists experience from the aggregate consumption of polluting goods (as captured by ϕ).

Given the merely illustrative purpose of this section, from now on we assume that $\alpha = 0$ and $p = 1$ for ease of exposition.³³ The difference equation has three candidates for the long run steady state: $q = 0$, $q = 1$ and the value of q that such that $U_t^e(q_t, \tau) = U_t^n(q_t, \tau)$, which is given by:

$$q^I(\tau) = \frac{(1 + \sigma)(1 + \tau)}{\sigma[\phi - \sigma(1 + \tau)]} \left[\ln(1 + \sigma) + \frac{\phi}{1 + \tau} - \sigma \right]$$

their defining characteristic is that preferences change in the direction suggested by relative utility payoffs (cultural fitness). That is, those preferences that fit well in the socio-economic environment and provide higher utility tend to spread.

³²This “replicator dynamics” equation can be derived from several micro-foundations. The one that suits better the interpretation of this model is one in which every individual receives opportunities once in a while to switch to the other preference type if it provides a higher utility payoff (Güth (1995)).

³³Of course, this implies that the absolute price of the polluting good is higher than the price of the clean good, which is at odds with empirical facts. However, for the purposes of this paper we only need the tax to affect relative prices. A more formal and general analysis can be found in Andrés-Cerezo (2021)

which can be shown to be interior as long as σ is large enough. Moreover, it can be shown that there is a $\bar{\sigma}(\phi, \tau)$ such that the only stable long-run steady state is given by $q^I(\tau)$ when $\sigma > \bar{\sigma}$ whereas the only stable long-run steady state is $q = 0$ if $\sigma < \bar{\sigma}$. When σ is relatively small, choosing to being an environmentalist makes an individual worse-off due to the additional suffering from pollution externalities. As a result, environmentalism tends to fade away in the population. In contrast, when σ is large enough the social esteem effect compensates the additional externality cost. However, in the long-run both groups co-exist in the population, as the social status from being an environmentalist is reduced as more people adopt this type of preferences.

If we disregard the dynamics of preferences, the expression for $q^I(\tau)$ is the analogue of the expressions that we usually encounter in static models of preference change (e.g., Mattauch et al. (2018)), in which preferences automatically change as we change the tax rate. However, as we will see in the following section, being explicit about the *adjustment process* allows to devise a way to specify weights for pre-policy and post-policy preferences.

3.4.2 Implications of different welfare criteria

Welfarism: Pareto principle

We start with the most standard approach of doing welfare analysis, which is applying the Pareto principle. According to this criterion, policy τ^* is Pareto efficient if there is no policy τ' such that the utility of each agent (*preference type*) under policy τ' is at least as good as with policy τ^* and at least one of the agents (*preference types*) has strictly higher utility.

In its purest form, when preferences are not endogenous, the main advantage of this criterion is that it avoids dubious inter-personal utility comparisons. However, when some individuals change preferences over time, this approach presents the problem that the analyst must decide which preferences to be used for those individuals that switch to the other type. However, this problem may be solved under some circumstances by defining Pareto optimality with respect to *preference types* instead of individual agents. This may circumvent some of the problems posed by endogenous preferences as long as both types end up co-existing in the population, since this guarantees that for individuals who switch type the policy is optimal with respect to both pre-policy and

post-policy preferences. Thus, in our example, defining Pareto optimality with respect to types can be robust only for those parameter constellations in which the *population* post-policy preferences adjust to the steady state $q^I(\tau)$, while it is not robust when they adjust towards $q = 0$.

To see this, first consider the interior bliss points $\tilde{\tau}^e$ and $\tilde{\tau}^n$ for each agent when $\sigma = 0$, which are given by:

$$\begin{aligned}\tilde{\tau}^e &= \psi + \phi \\ \tilde{\tau}^n &= \psi\end{aligned}$$

Both when evaluated with respect to the ex-ante preference distribution ($q = q_0$) or the ex-post one when it is interior ($q = q^I(\tau)$), we have that the set of Pareto efficient carbon taxes τ^{PE} is given by every τ such that $\tilde{\tau}^n < \tau < \tilde{\tau}^e$, which guarantees that the preferences of individuals who switch type are respected according to both reference points. However, when the post-policy preference distribution is $q = 0$, the only Pareto efficient tax rate according to ex-post preferences is $\tilde{\tau}^n$, differing from the set of optimal policies with ex-ante preferences, which is given by $\tilde{\tau}^n < \tau < \tilde{\tau}^e$.

However, even when we define the Pareto criterion with respect to types instead of agents, we may encounter problems in cases in which the bliss points of one or both types depend on the size of each preference group. When $\sigma > 0$, the bliss points of both types are given by:

$$\begin{aligned}\tilde{\tau}^e(q) &= \frac{\psi(1 + \sigma) - q\sigma(1 + \psi) + \phi[1 + \sigma(1 - q)]}{1 + \sigma} \\ \tilde{\tau}^n(q) &= \frac{\psi(1 + \sigma) - q\sigma(1 + \psi)}{1 + \sigma}\end{aligned}$$

Clearly, in this case we have that the set of Pareto efficient carbon taxes depends on whether we use pre-policy preferences q_0 or post-policy preferences $q^I(\tau)$ or $q = 0$.

Apart from its inability to provide a set of optimal policies that do not change depending on the reference point, this example highlights another well-known problem of using the Pareto principle, which is its indecisiveness to pick one optimal policy within the set of Pareto efficient taxes. That is, even if we

agree on whether we should use ex-ante or ex-post preferences, we cannot discriminate what is the best tax rate in the region between $\tilde{\tau}^n(q)$ and $\tilde{\tau}^e(q)$.

Welfarism: interpersonal utility comparisons

A) Ex-ante preferences

Now we turn to methods that allow for inter-personal utility comparisons. We start with the standard approach in environmental policy, in which a social planner proceeds as if preferences were not affected by the carbon tax or, equivalently, as if policies were evaluated according to pre-policy preferences. With respect to the Pareto principle, the key difference of this approach is the assumption about the cardinality and comparability of different preferences, which allows the aggregation of the utilities of all citizens.

In this case, the optimal policy is defined as the argument of the maximum of a standard Bergson-Samuelson social welfare function and, in particular, the Utilitarian one.³⁴ Also, we preserve all utility motives in the objective function of the central government, although one could argue that social planners should not take into account purely subjective motives (such as the status one) when evaluating social outcomes, and just stick to “material” sources of utility (see Bergstrom (2006) for a discussion of this point). The optimal carbon tax τ^* is defined as:

$$\tau^B = \arg \max_{\tau \in [0, \bar{\tau})} \{q_0 U^e(q_0, \tau) + (1 - q_0) U^n(q_0, \tau)\}$$

which is given by:

$$\tau^B = \frac{[(1 + \sigma) - \sigma q_0](\phi q_0 + \psi) - \sigma q_0}{(1 + \sigma)}$$

Note that, when $\sigma = 0$, we have that $\tau^B(q = 0) = \psi$, which is the standard Pigouvian tax to correct externalities. This is the standard recommendation of static models of environmental taxation in which preferences are stable, identical and only defined over consumption and externalities. When we include

³⁴We could also work with other additive social welfare functions with different weights and concave transformations of utilities, but the qualitative results would remain.

environmentally concerned citizens, the recommendation changes to take into account the additional disutility from pollution experienced by these groups, as well as the status motive.

Importantly, social planners that fail to internalize the endogeneity of preferences (either because of ignorance or to avoid the normative issues implied by their existence) will choose an inefficient policy according to the preferences induced by τ^B . This highlights the time-inconsistency problem created by endogenous preferences, as the policy that is optimal for initial preferences is no longer optimal under the new preferences.

B) Ex-post preferences

As we have seen in our discussion of the previous section, there are reasons to defend the use of post-policy preferences. Therefore, we now turn to explore the problem that faces a forward-looking social planner who understands that carbon taxes do not only affect the costs and benefits of different consumption choices, but also the distribution of preferences in the population. The following expression defines optimal policy in this case:

$$\begin{aligned} \tau^F &= \arg \max_{\tau \in [0, \bar{\tau})} \{qU^e(q, \tau) + (1 - q)U^n(q, \tau)\} \\ \text{s.t. } & q = q^{SS}(\tau) \end{aligned}$$

Recall that when σ is large enough, the long run steady state $q^{SS}(\tau)$ is given by $q^{SS} = q^I(\tau)$, while for low levels of σ it is given by $q^{SS} = 0$.³⁵

C) Combination of ex-ante and ex-post preferences

Now we turn to what I consider to be the more satisfactory way of doing normative analysis in static models of endogenous preferences. Given the arguments in favor and against the use of ex-ante or ex-post preferences, a reasonable middle ground would give positive weight to both. The problem is deciding

³⁵Note that in both cases, the long-run steady state is independent of q_0 . In other models of this type, or in this one when $\sigma = 0$ and $\alpha > 0$, multiple steady states could arise with different basins of attraction. In those cases, the chosen policy and the initial condition q_0 would determine which steady state is reached. Nevertheless, the normative problems encountered in this case are conceptually identical to the one presented here.

which normative principles should guide the choice of those weights. However, by being explicit about the dynamic process by which preferences change we can attach weights on the basis on parameters that can be estimated empirically. That is, depending on how quickly preferences change in response to policy, and depending on how we discount the future, different weights would be obtained for pre- and post-policy preferences. This can be done by specifying the problem of a social planner that internalizes the dynamic path of preferences (instead of the steady state) and discounts each preference profile accordingly. More formally, optimal policy τ^* is defined as:

$$\begin{aligned} \tau^* = \max_{\tau \in [0, \bar{\tau}] \forall t \geq 0} & \sum_0^{\infty} \rho^t \{q_t U^e(q_t, \tau) + (1 - q_t) U^n(q_t, \tau)\} \\ \text{s.t.} & \quad q_{t+1} - q_t = \theta q_t (1 - q_t) (U_t^e(q_t, \tau) - U_t^n(q_t, \tau)) \\ & \quad q(0) = q_0. \end{aligned}$$

where ρ is the discount factor. With this approach, the weights on initial and final preferences depend on how fast preference adjust, on how we discount the future and on how the policy is valued by intermediate preference distributions. Also, given the *adaptive* nature of preferences in these models, this approach respects the idea by von Weizsäcker (2005) of weighting final preferences relatively more, as once the steady state is reached preferences remain constant indefinitely (although these are appropriately discounted to present value). In addition, this method respects the most important property of preference satisfaction criteria, which consists on relying solely on individual utilities, as it chooses the weights without introducing any exogenous substantive criteria.

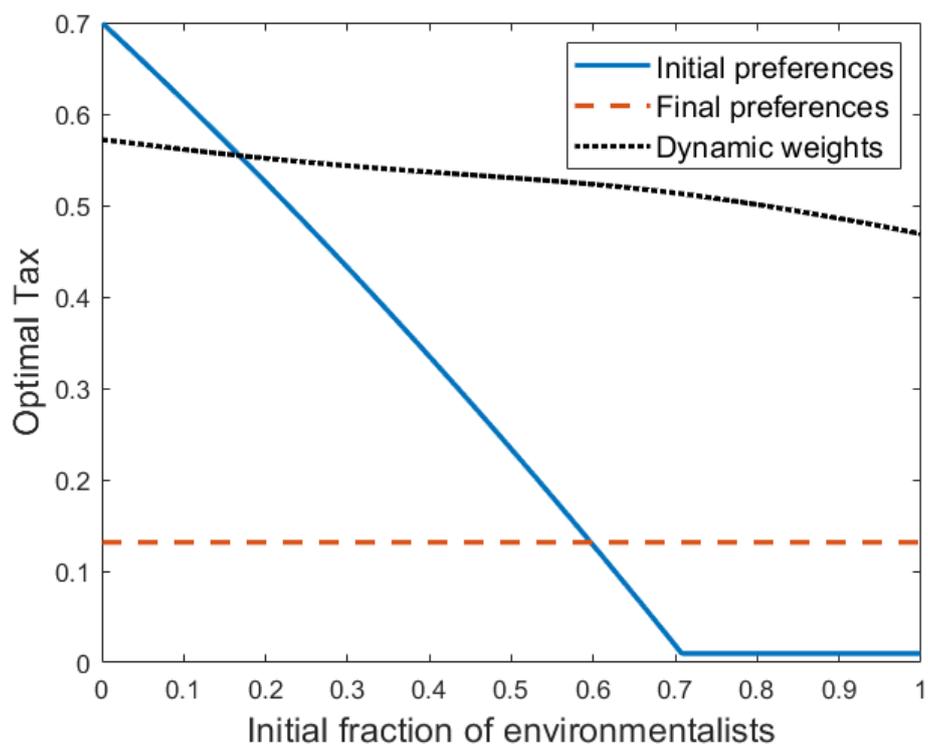
Unfortunately, the previous two problems do not admit an analytical solution. In Figure 1 I display the numerical solution for τ^B , τ^F and τ^* as a function of the initial preference distribution for a given choice of parameters that guarantees that the long-run steady state is given by $q^I(\tau)$ ($\sigma = 2$, $\psi = 0.7$, $\phi = 0.3$, $\rho = 0.9$, $\theta = 1$ and $y = 1$).³⁶

Apart from the specific mechanisms at work,³⁷ the main take-away from this

³⁶Other parameter constellations yield different results but most of them display a strong sensitivity of optimal policy to the choice of weights.

³⁷For an exhaustive analysis of the mechanisms and the full dynamics of preferences under different parameters constellations see Andrés-Cerezo (2021), who also explore long-run policy *paths* as well as the role of political leaders and activists in shaping the long-run evolution of environmental attitudes.

Figure 3.1 – Optimal tax rate as a function of initial preferences under different welfare criteria.



illustrative exercise is that, for any given initial preference distribution q_0 , the tax rate is extremely sensitive to whether the welfare effects of the policy are evaluated with respect to ex-ante, ex-post preferences or a combination of both. This is bad news for the important task of designing environmental policy and highlights the importance of developing sound normative principles for doing welfare analysis when preferences are endogenous. Nevertheless, we believe that relying on “dynamically derived weights” provides a compelling rationale for weighting pre-policy and post-policy preferences, as it allows to construct them in relation to parameters that could in principle be empirically identified. However, this approach faces the important problem of finding the right model of endogenous preferences for the situation at hand, as well as the one of determining values for the parameters that may be endogenous themselves. Although this method is far from settling the problem on which weights to use, which will always involve important value judgements on the side of the analyst, being explicit about the model and the parameters facilitates the task and makes it (partially) falsifiable and empirically grounded.

Non-welfarist criteria: emissions targets and desirable preferences

Our model can also nest some of the non-welfarist paternalist approaches discussed before.³⁸ The most common one in actual policy practice is to specify some target based on scientific criteria that take into account the sustainability of ecosystems and the impact of rising temperatures on socio-economic outcomes. In our model, this can be captured as a limit on the consumption of the polluting good i.e., $C \leq \bar{C}$. If the policy-maker totally disregards individual welfare in its calculations, the most direct approach would be to set the minimum tax τ^C such that the target is met (i.e., $C(\tau) = \bar{C}$). By only considering pre-policy preferences, the interior carbon tax is given by:

$$\tau^C = \frac{1 + \sigma(1 - q_0) - (1 + \sigma)\bar{C}}{(1 + \sigma)\bar{C}}$$

whereas if considering post-policy preferences, we have that

$$\tau^{C'} = \frac{1 + \sigma[1 - q^{SS}(\tau^{C'})] - (1 + \sigma)\bar{C}}{(1 + \sigma)\bar{C}}$$

The optimal tax rate will generally differ in these two cases. However, as compared to the preference satisfaction approach, the planner now does not face a normative dilemma between which weight to give to ex-ante and ex-post preferences. All she must do is setting the tax to guarantee that the target is met. Nevertheless, as compared to the preference satisfaction approach presented in previous sections, in this case failing to internalize that preferences are endogenous leads to policies that are unambiguously inefficient. In contrast, when relying on the preference satisfaction approach, failing to internalize the endogeneity of preferences cannot be considered inefficient in a strict sense, since the policy chosen would be efficient according to pre-policy preferences.

Also note that while relying on exogenous targets considerable simplifies policy evaluation and design, it comes at the potential cost of choosing policies that do not respect the Pareto principle, as it can make both preference types worse

³⁸As I we have mentioned, the different versions of the welfarist approach proposed imply some degree of paternalism, as the planner is indirectly choosing preferences (whether she likes it or not). However, to be fully paternalistic one could argue that the substantive criteria should directly appear in the objective function of the government.

off under some circumstances (Kaplow and Shavell (2001)). To see this in our simple example, note when the target \bar{C} is very stringent, τ^C and $\tau^{C'}$ can be higher than the bliss point of environmentalist $\tilde{\tau}^e$. This situation is clearly Pareto inefficient, as both types of agents could improve their utility by setting a lower tax rate.

We can think of other purely substantive criteria. For instance, given the inevitability of some degree of paternalism when preferences are endogenous, we could decide to be explicitly paternalistic and choose policy to promote environment-friendly preferences. In this case, we would just choose the biggest possible tax rate ($\tau = \bar{\tau}$) in order to maximize the relative size of this group.

Of course, we can also combine the preference satisfaction criteria with some of the substantive goals discussed in this subsection. For instance, an interesting hybrid approach would specify a social welfare function in which environmental preferences get a higher weight than the one implied by the relative size of this group. Or we could also maximize individual utilities subject to a constraint on total aggregate consumption. By construction, these possibilities would have some of the benefits of both approaches (e.g., respect for individual utilities, promotion of moral virtues) as well their shortcomings (e.g., potential violations of the Pareto principle, choice of ex-ante or ex-post preferences).

3.5 Conclusion

That preferences are endogenous is a claim that very few people would deny. Yet, the vast literature on environmental cost-benefit analysis has overwhelmingly ignored this possibility. From our previous discussion we can hypothesize that the main reason behind the reluctance to incorporate preference endogeneity probably lies in the conceptual difficulties it poses for normative work. This analytical convenience would justify keeping the exogeneity assumption if either this premise is not particularly incorrect, or its consequences are mild. Unfortunately, this is not the case. Given that refraining from doing normative analysis is not an option when it comes to environmental policy, this essay has discussed the pros and cons of different welfare-theoretic approaches for dealing with preference endogeneity.

The analysis suggests that in these circumstances we do not need to reject the standard preference satisfaction approach. There are a variety of possibilities

to adjust this framework that keep some of its most desirable properties, such as its focus on the consequences of policies and its respect for individual autonomy. The modification that seems particularly promising considers both pre-policy and post-policy preferences, and weights them by relying on falsifiable dynamic models of endogenous preferences and on empirically measurable parameters. Moreover, given that some degree of paternalism is inevitable when policies determine preferences, the weights could also partially reflect some other substantive criteria based on moral or scientific principles. On the negative side, by means of a simple theory model we have illustrated that the specific welfare criteria that we adopt has substantial implications for optimal policy design. Therefore, any normative analysis must check the sensitivity of the results to different welfare perspectives. A satisfactory policy recommendation would arise when they point in similar directions, while designing optimal policy would get complicated when recommendations differ widely.

Future work on this topic could refine and formalize some of the concepts discussed in this essay and propose new ones based on sound philosophical principles. Another interesting line of research would develop and test positive models of preference formation that could guide the development of normative frameworks, in the spirit of the one proposed in this paper. Finally, empirical work could attempt to identify fundamental preferences and to guide the choice of pre-policy and post-policy weights.

Appendix A

A.1 Mathematical Appendix

A.1.1 Appendix for Section 1.4

Proof of Proposition 1

The proof of Proposition 1 simply follows Spence and Starrett (1975). The main result is given by the following proposition.

Proposition 15 *The government problem (1.8) is equivalent to a problem with flow utility given by a function $G(q_t)$; that is,*

$$\begin{aligned} \max_{\delta_t \in [0,1] \forall t \geq 0} & \int_0^{\infty} e^{-\rho t} W(\delta_t, q_t) dt \\ \text{s.t.} & \dot{q}_t = r q_t (1 - q_t) (1 - \delta_t - q_t) \\ & q(0) = q_0, q_t \in [0, 1], \end{aligned}$$

is equivalent to problem

$$\begin{aligned} \max_{\delta_t \in [0,1] \forall t \geq 0} & \int_0^{\infty} e^{-\rho t} G(q_t) dt \\ \text{s.t.} & \dot{q}_t = r q_t (1 - q_t) (1 - \delta_t - q_t) \\ & q(0) = q_0, q_t \in [0, 1]. \end{aligned} \tag{A.1}$$

With the previous result we prove Proposition 1.

Proposition 1 *The optimal path for problem (1.8) is a most rapid approach (MRAP). That is, the optimal solution approaches as fast as possible a steady state in which per-period welfare is maximized.*

Proof. Given the linearity of $W(q_t, \delta_t)$ and $\dot{q}_t(q_t, \delta_t)$ in δ_t , the only possible optimal policies are given by

$$\delta_t^* = 0 \qquad \text{or} \qquad \delta_t^* = 1. \qquad (\text{A.2})$$

For sufficiency, observe that problem (1.8) can be equivalently rewritten as (A.1). Then, the optimal solution is characterized by reaching a state that maximizes flow utility, $G(q)$, as fast as possible by setting $\delta_t^* = 1$ or $\delta_t^* = 0$ appropriately. ■

Shape of $G(q)$

From Proposition 1 we know that the optimal steady-states are characterized by the local maxima of $G(q)$, and thus those maxima need to be found. The following proposition shows where $G(q)$ is maximized:

Proposition 16 *For any value of parameters, one has that*

$$\lim_{q \rightarrow 0} G(q) = \infty \qquad \lim_{q \rightarrow 1} G(q) = \infty, \qquad (\text{A.3})$$

and a unique local minimum exists on the interval $(0,1)$.

Proof. Observe that from the law of motion, we can express δ as a linear function of $\frac{\dot{q}}{rq(1-q)}$ and q

$$d(q, \dot{q}) = 1 - q - \frac{\dot{q}}{rq(1-q)} = A_1 + A_2q + A_3 \frac{\dot{q}}{q(1-q)},$$

with $A_2 \leq 0$ and $A_3 \leq 0$.

Substituting this expression inside the government's flow utility $W(q, \delta(q); \omega)$, we get

$$W(q, \dot{q}) = B_0 + B_1q + B_2q^2 + B_3 \frac{\dot{q}}{(1-q)} + B_4 \frac{\dot{q}}{q}.$$

If we write $W(q, \dot{q}) = M(q) + N(q)\dot{q}$, we can easily check that

$$M(q) = B_0 + B_1q + B_2q^2 \qquad N(q) = B_3 \frac{1}{(1-q)} + B_4 \frac{1}{q},$$

with $B_3 \geq 0$ and $B_4 \leq 0$ from $A_3 \leq 0$ and the functional form of W . Hence, for $S(q)$, we have that

$$S(q) = \int_{q_0}^q N(s)ds = \int_{q_0}^q \left(B_3 \frac{1}{(1-s)} + B_4 \frac{1}{s} \right) ds = C - B_3 \ln(1-q) + B_4 \ln(q),$$

which is well defined as long as $q_0 \in (0, 1)$.

Putting everything together

$$G(q) = C_0 + C_1q + C_2q^2 + C_3 \ln(1-q) + C_4 \ln(q),$$

with $C_2 \geq 0$, because of the functional form of W and $A_2 \leq 0$, as well as $C_3 \leq 0$, and $C_4 \leq 0$. The first derivative of $G(q)$ is given by

$$G'(q) = C_1 + 2C_2q - C_3 \frac{1}{1-q} + C_4 \frac{1}{q},$$

Moreover, the second derivative is always positive

$$G''(q) = 2C_2 - C_3 \frac{1}{(1-q)^2} - C_4 \frac{1}{q^2} > 0,$$

because $C_2 \geq 0, C_3, C_4 \leq 0$ for $q \in (0, 1)$. Hence, $G(q)$ is strictly convex with a local minimum in $(0, 1)$ and

$$\lim_{q \rightarrow 0} G(q; \omega) = \infty \qquad \lim_{q \rightarrow 1} G(q; \omega) = \infty,$$

■

Proof of Theorem 1

According to Spence and Starrett (1975), the optimal trajectory of problem (1.8) consists of reaching as fast as possible the value of the state where $G(q)$ is locally maximized. In our case, $G(q)$ has no local maxima, so we cannot directly apply their approach. However, Theorem 1 modifies the results of Spence and Starrett (1975) to show that the solution to problem (1.8) is still characterized by an MRAP.

Theorem 1 *For any value of $\omega \in \Omega$ and any initial condition $q_0 \in (0, 1)$, optimal policies set either $\delta = 0$ or $\delta = 1$ forever with no switch in policies.*

Furthermore, there is no interior steady state so optimal paths approach as fast as possible one of the extreme stationary points of the state variable:

$$\lim_{t \rightarrow \infty} q_t^* = 0 \qquad \text{or} \qquad \lim_{t \rightarrow \infty} q_t^* = 1.$$

where q_t^* is the path of the state variable under the optimal policy δ^* .

Proof. From Proposition 1, we know that any optimal policy will set $\delta_t^* = 0$ or $\delta_t^* = 1$. Assume initial conditions q_0 . Recall the law of cultural transmission:

$$\dot{q}_t = r q_t (1 - q_t) (1 - \delta_t - q_t).$$

Observe, from the law of motion, that the state variable q_t will stay in $[0, 1]$ for any choice of the control variable δ_t . Moreover, it is easy to see from the previous equation that by setting δ_t appropriately, the government can reach any point $q \in (0, 1)$ in finite time, and that it can also choose to stay at any given point $q_t \in (0, 1)$ just by setting $\delta_t = 1 - q_t$ forever.

Now we prove the first part of the theorem. Observe that a swinging path covering any interval $I \subset [0, 1]$ is not optimal. Clearly, a swinging path is dominated by staying at

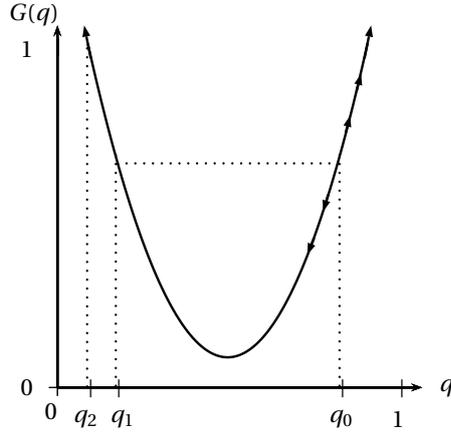
$$q^* \in \arg \max_{q \in I} G(q)$$

Assume *wlog* that $G(q_0) \geq \min_{q \in [0, 1]} G(q)$ and that $q_0 \geq \arg \min_{q \in [0, 1]} G(q)$, as in Figure A.1.

Now we prove the second part. Because continuity holds, we can only go either left or right with no jumps as shown by the arrows on the graph. Assume first that we go left. Any optimal path must reach q_1 , because if not staying at q_0 forever strictly dominates that path, contradicting that the path is optimal.

Because of the first part of the proposition, we see that once we have reached q_1 , a swinging path is not optimal on the interval $[q_1, q_0]$. In this case, any swinging path is clearly dominated by, for example, staying at q_0 the corresponding amount of time.

With a similar argument, we proceed to prove the second part of the theorem. Consider an arbitrary point q_2 to the left of q_1 . The government is better off by reaching q_2 as fast as possible and staying there forever. Hence, we can rule out paths that never reach any $q_2 \in (0, q_1]$. The case of going right under an

Figure A.1 – Optimal path on $G(q)$


optimal path starting from q_0 follows the same arguments. ■

Proof of Theorem 2

Theorem 1 shows that there are no interior steady states and that the long-run optimal is homogeneous. Theorem 2 characterizes the optimal policy and therefore how these steady-states are reached.

Theorem 2 *For any parameter values, it exists a unique \bar{q}_0 such that the government is indifferent between setting $\delta^* = 0$ or $\delta^* = 1$ forever. That is, the optimal policy is characterized by threshold \bar{q}_0 as follows:*

$$\delta^*(q) = \begin{cases} 1 & \text{if } q \leq \bar{q}_0 \\ 0 & \text{if } q \geq \bar{q}_0 \end{cases}$$

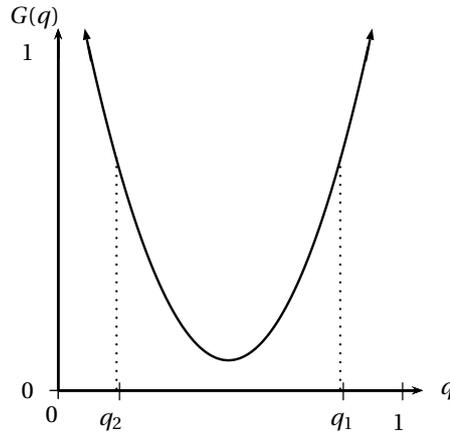
Proof. Define

$$F(q_0) = \int_0^\infty e^{-\rho t} [W(q_t^1(q_0), \delta = 1)] dt - \int_0^\infty e^{-\rho t} [W(q_t^0(q_0), \delta = 0)] dt \quad (\text{A.4})$$

$$\text{s.t.} \quad (\text{A.5})$$

$$\dot{q}^\delta = rq_t(1 - q_t)(1 - \delta_t - q_t), q(0) = q_0 \text{ with } \delta \in \{0, 1\}, \quad (\text{A.6})$$

where q^δ is the path for q when the government sets $\delta \in \{0, 1\}$. Denote by $q^\delta(q_0)$ the solution to the differential equation $\dot{q} = q(1 - q)(1 - \delta - q)$ with

Figure A.2 – Uniqueness


initial condition q_0 . That is

$$\dot{q}^0 = rq(1 - q)^2 \qquad \dot{q}^1 = -rq^2(1 - q), \qquad (\text{A.7})$$

Because setting $\delta = 1$ or $\delta = 0$ forever are the only possible optimal policies, F represents the welfare difference between the only optimal policies at every initial point q_0 . Therefore, when $F(q_0) > 0$, it is optimal to set $\delta^* = 1$ and, on the contrary, when $F(q_0) < 0$, $\delta^* = 0$ is optimal.

Evaluating $F(q_0)$ at $q_0 = 0$ and at $q_0 = 1$ yields:

$$F(0) = \int_0^\infty e^{-\rho t} r (2 - (\alpha + \beta)) dt = \frac{r}{\rho} (2 - (\alpha + \beta)) > 0, \qquad (\text{A.8})$$

$$F(1) = \int_0^\infty e^{-\rho t} r (-(\alpha + \beta)) dt = -\frac{r}{\rho} (\alpha + \beta) < 0. \qquad (\text{A.9})$$

By continuity of $F(q_0)$ in q_0 and applying the intermediate value theorem, a \bar{q}_0 exists such that $F(\bar{q}_0) = 0$. At that point, the government is indifferent between setting $\delta = 0$ or $\delta = 1$ for all $t \geq 0$.

It remains to show that the threshold \bar{q}_0 is unique. Assume two points exist for which we have indifference, say, q_1 and q_2 as depicted in Figure A.2. Without loss of generality assume that $G(q_1) \geq G(q_2)$. Hence, at initial point q_1 , the government can go to q_2 and then, because at q_2 it is indifferent between moving left or right, then government can go from q_2 to 0. But observe that that path is dominated if we go directly from q_1 to 1. Hence, going from q_1 to q_2 and then from q_2 to 0 cannot be optimal. Thus, q_1 and q_2 cannot be both

indifference points. ■

A.1.2 Appendix for Section 1.5

First, let us denote by $V(q)$ the government discounted utility under the optimal policy

$$\begin{aligned} V(q) &= \max_{\delta_t \in [0,1]} \int_0^\infty -e^{-\rho t} W(q_t, \delta^*(q_t)) dt \\ \text{s.t. } \quad \dot{q}_t &= r q_t (1 - q_t) (1 - \delta^*(q_t) - q_t) \text{ and } q(0) = q, \end{aligned} \quad (\text{A.10})$$

also known as the value function. Observe that the value function is well-defined and continuous on $[0, 1]$ because

$$||W(q, \delta)|| < M < \infty,$$

for some $M < \infty$, and for all $q, \delta \in [0, 1]$.

Proof of Theorem 3

Theorem 3 *Assume that $\delta^*(q)$ is a solution to the following optimal control problem:*

$$\begin{aligned} \max_{\delta_t \in \Delta} \int_0^\infty e^{-\rho t} W(q_t, \delta_t) dt \\ \text{s.t. } \quad \dot{q}_t = g(\delta_t, q_t) \text{ and } q_0 = q. \end{aligned} \quad (\text{A.11})$$

Denote by $\delta^S(q)$ the stationary policy function

$$\dot{q} = g(q, \delta^S(q)) = 0,$$

and define function

$$H(q) = W(q, \delta^S(q)).$$

If for some interior \tilde{q} , we have that $\delta^*(\tilde{q}) = \delta^S(\tilde{q}) \in \Delta^\circ$, such that $\{g(\tilde{q}, \delta) | \delta \in [0, 1]\}$ is an open neighborhood of 0, then \tilde{q} is a local maximum of $H(q)$.

Proof. We prove it by contradiction. Assume that optimal policy is given by $\delta^*(q)$ and that for some \tilde{q}

$$\delta^*(\tilde{q}) = \delta^S(\tilde{q})$$

such that $\{g(\tilde{q}, \delta) | \delta \in [0, 1]\}$ is an open neighborhood of 0 so we can move in any

direction at \tilde{q} , but \tilde{q} is not a local maximum of $H(q)$. Without loss of generality assume $H'(\tilde{q}) > 0$. Pick any function $\epsilon(q)$ such that $\dot{\tilde{q}} = g(\tilde{q}, \delta^S(\tilde{q}) + \epsilon(\tilde{q})) > 0$ with policy $\delta^S(q) + \epsilon(q)$ feasible for an open neighborhood of \tilde{q} . Define \bar{q}_τ such that

$$\bar{q}_\tau = q_\tau,$$

where q follows the path defined by policy $\delta^S(q) + \epsilon(q)$.¹ For every $\tau > 0$, construct policy

$$\delta'_\tau(q) = \begin{cases} \delta^S(q) & \text{if } q = \bar{q}_\tau \\ \delta^S(q) + \epsilon(q) & \text{if } q \neq \bar{q}_\tau. \end{cases}$$

The intuition for policy δ'_τ is to move away from δ^S at an $\epsilon(q)$ rate until hitting \bar{q}_τ , and stay at \bar{q}_τ forever afterwards. For example, if $\tau = 0$, then $\delta'(q) = \delta^S(q)$ for all q .

We want to compare δ^* with δ'_τ starting at \tilde{q} .² Given their definitions, the discounted utility for each policy is given by:

$$J(\tilde{q}, \delta^*) = \int_0^\infty e^{-\rho t} H(\tilde{q}) dt \quad (\text{A.12})$$

$$J(\tilde{q}, \delta'_\tau) = \int_0^\tau e^{-\rho t} W(q_t, \delta'_\tau(q_t)) dt + e^{-\rho \tau} (H(q_\tau)). \quad (\text{A.13})$$

Define $F(\tau)$ as the surplus difference between policy δ'_τ and δ^* :

$$F(\tau) \equiv J(\tilde{q}, \delta'_\tau) - J(\tilde{q}, \delta^*) = \int_0^\tau e^{-\rho t} (W(q_t, \delta'_\tau(q_t)) - H(\tilde{q})) dt + e^{-\rho \tau} (H(q_\tau) - H(\tilde{q})). \quad (\text{A.14})$$

Because δ^* is the optimal policy, it must be the case that $F(\tau) \leq 0$. However, observe that

$$F(\tau) = F'(0)\tau + o(\tau^2),$$

¹We omit the dependence on the initial value q_0 for the ease of exposition.

² $\delta'_\tau(q)$ might not be feasible for all τ and all initial values q_0 , but since $\delta^*(\tilde{q}) = \delta^S(\tilde{q}) \in \Delta^o$, $\delta'_\tau(q)$ will be well-defined for sufficiently small τ in a neighborhood of \tilde{q} .

because $F(0) = 0$. Taking derivatives from A.14 with respect to τ^3

$$F'(\tau) = e^{-\rho\tau} (H(q_\tau) - H(\tilde{q})) - \rho e^{-\rho\tau} (W(q_\tau, \delta'(q_\tau)) - H(\tilde{q})) + e^{-\rho\tau} H'(q_\tau) \dot{q}_\tau, \quad (\text{A.15})$$

and evaluating at $\tau = 0$

$$F'(0) = H'(\tilde{q}) \dot{\tilde{q}} > 0,$$

which contradicts the fact that δ^* is the optimal policy.⁴ ■

Proof of Theorem 5

We proceed in steps. First, consider the corresponding HJB equation of problem 1.10:

$$\begin{aligned} \rho V(q) = \max_{\delta \in [0,1]} & \psi^N q + \alpha q (f(1-r) + r(1-\delta)) + (1-\alpha)(1-q)((f(1-r) + r\delta) \\ & - r^2(\beta q \delta^2 + (1-\beta)(1-q)(1-\delta)^2) \\ & + r q (1-q)(1-\delta-q)V'(q). \end{aligned}$$

Taking derivatives with respect to δ , we obtain

$$r((1-\alpha)(1-q) - \alpha q) - r^2 2(\beta q \delta - (1-\beta)(1-q)(1-\delta)) - V'(q) r q (1-q). \quad (\text{A.16})$$

Hence, for an interior solution of $\delta^*(q)$, we can write

$$\delta^*(q) = \frac{1}{\beta q + (1-\beta)(1-q)} \left((1-\beta)(1-q) + \frac{(1-\alpha)(1-q) - \alpha q - q(1-q)V'(q)}{2r} \right). \quad (\text{A.17})$$

Because there is no interior steady state, it must hold that $\delta(q) \neq 1-q$ for all $q \in (0, 1)$. This implies the following result:

³Recall that δ'_τ implicitly depends on τ and the term $\int_0^\tau \frac{\partial}{\partial \tau} \delta'(q) e^{\rho t} (W(q_t, \delta'(q_t)) - H(\tilde{q})) dt$ should be included in the derivative too. However see Supplementary Appendix A to see that

$$\int_0^\tau \frac{\partial}{\partial \tau} \delta'(q) e^{\rho t} (W(q_t, \delta'(q_t)) - H(\tilde{q})) dt = 0.$$

⁴Even when $H'(\tilde{q}) = 0$ the result still holds. If this is the case, we can use the second order Taylor approximation of the function $F(\tau)$. For more details see Supplementary Appendix A.

Proposition 17 *For any interior state $q \in (0, 1)$, it follows that*

$$V'(q) \neq s(q) \equiv 2r(1 - 2\beta) + \frac{1 - \alpha}{q} - \frac{\alpha}{1 - q}.$$

Proof. We know that $\delta(q) \neq 1 - q$ for all $q \in [0, 1]$. Set $\delta = 1 - q$ in equation (A.16) and solve for $V'(q)$. ■

From the previous proposition it follows that $V(q)$ is not differentiable at some $\bar{q} \in (0, 1)$, that is, $V(q)$ has a kink at some $\bar{q} \in (0, 1)$. This result is shown in the following proposition:

Proposition 18 *$V'(q)$ is not continuous on $(0, 1)$. Moreover, in a neighborhood of $q = 1$ it must be the case that $V'(q) > s(q)$ and therefore $\delta^*(q) < 1 - q$ and so $\dot{q} > 0$. Similarly, in a neighborhood of $q = 0$, $V'(q) < 1 - q$ and $\delta^*(q) > 1 - q$ and so $\dot{q} < 0$.*

Proof. We know that

$$V'(q) \neq s(q),$$

for all $q \in (0, 1)$. The only way that $V'(q)$ can be continuous is to have either $V'(q)$ always above, or always below that function:

$$V'(q) > s(q) \text{ or } V'(q) < s(q),$$

for all $q \in (0, 1)$. Without loss of generality let's assume that $V'(q)$ is continuous in $(0, 1)$ with

$$V'(q) < s(q) = 2r(1 - 2\beta) + \frac{1 - \alpha}{q} - \frac{\alpha}{1 - q}.$$

This means that $V'(q) < s(q)$ in a neighborhood of $q = 1$. By continuity of $V'(q)$ in $(0, 1)$ it follows

$$\lim_{q \rightarrow 1^-} V'(q) = -\infty,$$

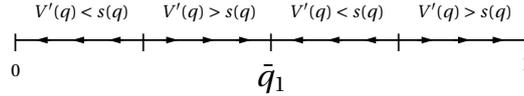
which contradicts the continuity of $V(q)$ at $q = 1$. Therefore it must be the case that $V'(q) > s(q)$ in a neighborhood of $q = 1$. Following a similar argument, $V'(q) < s(q)$ in a neighborhood of $q = 0$. Hence, $V'(q)$ is discontinuous at some $q \in (0, 1)$. ■

Proposition 19 *There is a threshold \bar{q}_0 such that*

$$V'(q) > s(q) \iff q > \bar{q}_0.$$

Proof. We know from Proposition 18 a q in $(0, 1)$ exists at which $V'(q)$ jumps function $s(q)$. We also know $V'(q) < s(q)$ near $q = 0$ as well as $V'(q) > s(q)$ near $q = 1$. Assume for a contradiction there is more than one jump. Then, there is a point \bar{q}_1 such that for some $\epsilon > 0$, $\dot{q} < 0$ for all $q \in (\bar{q}_1, \bar{q}_1 + \epsilon)$ and $\dot{q} > 0$ for all $q \in (\bar{q}_1 - \epsilon, \bar{q}_1)$ as shown in the example of Figure A.3. This would imply that \bar{q}_1 is an interior stationary point, which is a contradiction.⁵ ■

Figure A.3 – Example of many thresholds



Proposition 20 $\delta^*(q) = 0$ in a open neighborhood of $q = 1$. $\delta^*(q) = 1$ in an open neighborhood of $q = 0$.

Proof. Assume for a contradiction that $\delta^*(q) > 0$ near $q = 1$. We know from Proposition 18 that $\delta^*(q) < 1 - q < 1$ in some open neighborhood of $q = 1$, say $\mathcal{O}(1)$. Therefore for $q \in \mathcal{O}(1)$, $\delta^*(q)$ is interior, $\delta^*(q) \in (0, 1)$, and defined by the solution to the first order condition

$$1 > \delta^*(q) = \frac{1}{\beta q + (1 - \beta)(1 - q)} \left((1 - \beta)(1 - q) + \frac{(1 - \alpha)(1 - q) - \alpha q - q(1 - q)V'(q)}{2r} \right) > 0,$$

which implies

$$\frac{1 - \alpha}{q} - \frac{\alpha}{1 - q} + \frac{2r(1 - \beta)}{q} > V'(q) > \frac{1 - \alpha}{q} - \frac{\alpha}{1 - q} - \frac{2r\beta}{1 - q}$$

By continuity of $V'(q)$ in $\mathcal{O}(1) \cap (0, 1)$ it must be the case that

$$\lim_{q \rightarrow 1^-} V'(q) = -\infty,$$

which contradicts again that $V(q)$ is continuous at $q = 1$. ■

Proposition 21 The optimal policy $\delta^*(q)$ is continuous on $[0, \bar{q}_0) \cup (\bar{q}_0, 1]$.

Proof. It is easy to see that the optimal control satisfies

$$\delta^*(0) = 1 \quad \delta^*(1) = 0,$$

which combined with continuity of $V'(q)$ on $[0, \bar{q}_0) \cup (\bar{q}_0, 1]$ and Proposition 20 delivers the result. ■

⁵Observe limit cycles are ruled out for dynamic autonomous systems or single variable with discount factor $\rho > 0$, since the Jacobian of the canonical system, J , satisfies $\text{tr}(J) = \rho > 0$. For a more detailed discussion see Grass et al. (2008), proposition 3.83

Corollary 8 *The long-run steady state of q_t are characterized as follows*

$$\lim_{t \rightarrow \infty} q_t = \begin{cases} 0 & \text{if } q_0 \leq \bar{q}_0 \\ 1 & \text{if } q_0 \geq \bar{q}_0, \end{cases}$$

with the property that \bar{q}_0 is a Skiba point of the dynamic system, that is, the government is indifferent between converging to 0, $\lim_{t \rightarrow \infty} q_t = 0$ or converging to 1, $\lim_{t \rightarrow \infty} q_t = 1$ at \bar{q}_0 .

Proof. It follows directly from the the threshold characteristic of $\delta^*(q)$

$$\delta^*(q) = \begin{cases} > 1 - q & \text{if } q_0 \leq \bar{q}_0 \\ < 1 - q & \text{if } q_0 \geq \bar{q}_0. \end{cases}$$

■

Putting all the previous results together, we are able to prove theorem 5:

Theorem 5 *There exists $\bar{q}_0 \in (0, 1)$ such that*

$$\delta^*(q) > 1 - q \quad \text{if } q \leq \bar{q}_0, \quad \delta^*(q) < 1 - q \quad \text{if } q \geq \bar{q}_0.$$

Moreover $\delta^*(q)$ continuous on $[0, \bar{q}_0) \cup (\bar{q}_0, 1]$ and there two open neighborhoods of $q = 0$ and $q = 1$, say $\mathcal{O}(0)$ and $\mathcal{O}(1)$ in $[0, 1]$, such that

$$\delta^*(q) = 1 \forall q \in \mathcal{O}(0) \quad \delta^*(q) = 0 \forall q \in \mathcal{O}(1).$$

A.1.3 Appendix for Section 1.6

For details about the microfoundation of the objective function in the electoral competition game see Supplementary Appendix A.

Observe that because $\|W(q, \delta^i, \delta^{-i})\| = \|\psi^N q + p_i(\delta^i, \delta^{-i})\| \leq \psi^N + 1 < \infty$, we know that the value functions

$$\begin{aligned} V^i(q_0) &= \left\{ \mathbb{E}_0 \int_0^\infty e^{-\rho t} W^i(q_t, \delta^{*i}(q_t), \delta^{*-i}(q_t)) dt \right\} \\ V^{-i}(q_0) &= \left\{ \mathbb{E}_0 \int_0^\infty e^{-\rho t} W^{-i}(q_t, \delta^{*i}(q_t), \delta^{*-i}(q_t)) dt \right\}, \end{aligned}$$

are well-defined and continuous. The corresponding HJB equations are

$$\begin{aligned}\rho V^i(q) &= \max_{\delta \in [0,1]} W^i(q, \delta^i, \delta^{*-i}) + \mathbb{E}_q[g(q, \delta, \delta^{*-i})V_q^i(q)], \\ \rho V^{-i}(q) &= \max_{\delta \in [0,1]} W^i(q, \delta^{i*}, \delta) + \mathbb{E}_q[g(q, \delta^{i*}, \delta)V_q^{-i}(q)],\end{aligned}$$

where E_q is the expectation conditional on current state q . Given that two parties solve identical problems, we can restrict our attention to symmetric equilibria. Under symmetric equilibria, it follows that $V^i = V^{-i} = V$. Furthermore, it also holds

$$\mathbb{E}_q[g(q, \delta^i, \delta^{-i})V_q(q)] = rq(1-q) \left(1 - q - (p^i(\delta^i, \delta^{-i})\delta_i + (1 - p^i(\delta^i, \delta^{-i}))\delta^{-i}) \right) V_q(q), \quad (\text{A.18})$$

First, we prove some results that are used in the proof of Theorem 6:

Claim 1 *Let $q^0(t)$ and $q^h(t)$ two paths defined by dynamics*

$$\dot{q} = rq(1-q)(1-q-\delta(q)),$$

with initial conditions q_0 and $q_0 + h$, respectively. Then,

$$q^h(t) > q^0(t) \iff h > 0, \forall t \geq 0.$$

Proof. Observe both functions $q^0(t)$ and $q^h(t)$ are differentiable, and therefore continuous. We prove the claim by contradiction. Suppose it does not hold. Because $q^0(t)$ and $q^h(t)$ are continuous with $q^0(0) = q_0 < q^h(0) = q_0 + h$, if the claim does not hold, a \tilde{t} exists such that

$$q^0(t) < q^h(t) \quad \forall t < \tilde{t} \quad \text{and} \quad q^0(\tilde{t}) = q^h(\tilde{t}) = \tilde{q}.$$

The previous inequality implies

$$\dot{q}^0(\tilde{t}) = \lim_{\Delta \rightarrow 0} \frac{q^0(\tilde{t}) - q^0(\tilde{t} - \Delta)}{\Delta} > \lim_{\Delta \rightarrow 0} \frac{q^h(\tilde{t}) - q^h(\tilde{t} - \Delta)}{\Delta} = \dot{q}^h(\tilde{t}),$$

which is a contradiction because at \tilde{t} both dynamics are defined as

$$\dot{q}^0(\tilde{t}) = \dot{q}^h(\tilde{t}) = \tilde{q}(1-\tilde{q})(1-\tilde{q}-\delta(\tilde{q})).$$

■

Using the previous claim, we have the following result:

Claim 2 *At the points where $V(q)$ is differentiable, for any symmetric equilibrium of the nation-building electoral competition problem it must hold*

$$V_q(q) \geq 0$$

Proof. Under any symmetric equilibrium, it follows

$$p^i(q, \delta^*, \delta^*) = \frac{1}{2}.$$

Hence

$$V(q) = \int_0^\infty e^{-\rho t} \left(\psi^N q_t + \frac{1}{2} \right) dt.$$

Therefore

$$V(q+h) - V(q) = \int_0^\infty e^{-\rho t} (q_t^h - q_t^0) dt > 0,$$

which implies $V_q(q) \geq 0$ for any symmetric equilibrium $\delta^*(q)$ because $q_t^h > q_t^0$ for any $h > 0$. ■

Second, we show that equilibrium strategies cannot be interior, $\delta^*(q) \notin (0, 1)$.

Proposition 22 *For all $q \in [0, 1]$, there is a unique symmetric equilibrium announcement that is not interior*

$$\delta^*(q) \in \{0, 1\}.$$

Proof. We start by characterizing the best-response function δ^{*i} for given opponent's strategy δ^{-i} by solving the following problem:

$$\begin{aligned} \rho V(q) &= \max_{\delta \in [0,1]} \psi^N q + p^i(\delta, \delta^{-i}, q) + \mathbb{E}_q [g(q, \delta, \delta^{-i}) V_q(q)] \\ &= \max_{\delta \in [0,1]} \psi^N q + \frac{1}{2} + \Phi(q)(\delta - \delta^{-i}) \\ &\quad + r q (1 - q) \left(1 - q - \left[\frac{1}{2}(\delta + \delta^{-i}) + \Phi(q)(\delta - \delta^{-i})^2 \right] \right) V_q(q), \quad (\text{A.19}) \end{aligned}$$

where we have used $p^i(\delta, \delta^{-i}, q) = \frac{1}{2} + \Phi(q)(\delta - \delta^{-i})$ with

$$\Phi(q) = \frac{(1 - q)\phi^R - q\phi^N}{q\phi^N + (1 - q)\phi^R},$$

and equation A.18. We check for equilibrium strategies by invoking the *one-shot deviation principle*. Observe that in equation A.19, tomorrow's payoff is included in the continuation value through the term $\mathbb{E}_q [g(q, \delta, \delta^{-i}) V_q(q)]$.

That is, we check if the strategy of player i is a best-response to δ^{-i} assuming that in the future players keep playing the equilibrium with associated payoff $V(q)$.

The first-order conditions of equation A.19, which are given by

$$\Phi(q) - rq(1-q)V_q(q) \left(\frac{1}{2} + 2\Phi(q)(\delta^i - \delta^{-i}) \right), \quad (\text{A.20})$$

with second-order conditions given by

$$-rq(1-q)V_q(q)2\Phi(q).$$

We consider different cases, depending on the sign of $\Phi(q)$. Observe that the sign of $\Phi(q)$ determines the equilibrium announcement of a game without forward-looking parties that only care about winning the elections.

Case 1: $\Phi(q) < 0$

When $q \in \Phi^- \equiv \{q | \Phi(q) < 0\}$ and taking the other player strategy as given δ^{-i} , the probability of winning the election for party i is maximized at $\delta^i = 0, \forall \delta^{-i}$. Since the continuation value $\mathbb{E}_q[g(q, \delta, \delta^{-i})V_q(q)]$ is also maximized at $\delta^i = 0$ given that $V_q(q) \geq 0$ for all δ^{-i} , we have that $\delta^i(q) = 0$ is a dominant strategy for $i = A, B$. Therefore, for $q \in \Phi^-$, there is a unique equilibrium with $\delta^A(q) = \delta^B(q) = 0$.

Given the previous equilibrium announcements, we can solve for the value function in the subspace Φ^- . Substituting equilibrium strategies inside the HJB equation, we obtain

$$\rho V(q) = \psi^N q + \frac{1}{2} + rq(1-q)^2 V_q(q).$$

When we substitute $H(q) = \rho V(q) - (\psi^N q + \frac{1}{2})$, the resulting differential equation for $H(q)$ is given by

$$H'(q) - \frac{\rho}{rq(1-q)^2} H(q) = -\frac{\psi^N}{\rho}.$$

We obtain the integrating factor of the previous ODE, $m(q)$, by solving $m_q(q) = -\frac{\rho}{rq(1-q)^2} m(q)$. The solution is given by

$$m(q) = C e^{-\frac{\rho}{r}(\ln(q) - \ln(1-q) + \frac{1}{1-q})}.$$

Hence, the solution to the original equation is

$$m(1)H(1) - m(q)H(q) = - \int_q^1 \frac{\psi^N}{\rho} m(q) dq.$$

But since $m(1) = 0$, the solution for $H(q)$ is

$$H(q) = \frac{1}{m(q)} \int_q^1 \frac{\psi^N}{\rho} m^0(q) dq,$$

and so

$$V(q) = \frac{\psi^N q}{\rho} + \frac{1}{2\rho} + \frac{1}{m(q)} \int_q^1 \frac{\psi^N}{\rho} m(q) dq, \quad (\text{A.21})$$

for all $q \in \Phi^-$.

Case 2: $\Phi(q) = 0$

Observe that when $\Phi(q) = 0$ it must be $q = \hat{q}_S$. In this case, the probability of winning is always $\frac{1}{2}$ and so it is independent of the announcements. Then, if $V_q(\hat{q}_S) > 0$, then $\delta^*(\hat{q}_S) = 0$ is also a dominant strategy and if $V_q(\hat{q}_S) = 0$, then the whole interval $[0, 1]$ is a dominant strategy.

It is easy to prove that we can only have $V_q(\hat{q}_S) > 0$. Assume not, that is $V_q(\hat{q}_S) = 0$. Substituting $V_q(\hat{q}_S) = 0$ in equation A.19, it follows

$$\rho V(\hat{q}_S) = \psi^N \hat{q}_S + \frac{1}{2}.$$

By continuity of $V(q)$, it must hold

$$\lim_{q \rightarrow \hat{q}_S^+} \rho V(q) = \rho V(\hat{q}_S).$$

Using the solution of the value function given in A.21, we have

$$\lim_{q \rightarrow \hat{q}_S^+} \rho V(q) = \psi^N \hat{q}_S + \frac{1}{2} + \frac{1}{m(\hat{q}_S)} \int_{\hat{q}_S}^1 \frac{\psi^N}{\rho} m(q) dq,$$

and so

$$\frac{1}{m(\hat{q}_S)} \int_{\hat{q}_S}^1 \frac{\psi^N}{\rho} m(q) dq = 0.$$

This is a contradiction because $m(q) > 0$ for all $q \in (0, 1)$ with $m(\hat{q}_S) < \infty$ where $\hat{q}_S = \frac{\phi^N}{\phi^N + \phi^R} \in (0, 1)$.

Therefore, we conclude equilibrium announcements satisfy:

$$\delta^*(q) = 0, \quad \forall \quad \Phi(q) \leq 0$$

Case 3: $\Phi(q) > 0$

When $q \in \Phi^+ \equiv \{q | \Phi(q) > 0\}$, there is a trade-off between increasing the probability of winning elections and announcing a policy such that q increases in the next period.

The best-response function to opponent's announcement δ^{-i} is characterized by the first order conditions

$$\Phi(q) - rq(1-q)V_q(q) \left(\frac{1}{2} + 2\Phi(q)(\delta^i - \delta^{-i}) \right), \quad (\text{A.22})$$

because the second-order conditions

$$-rq(1-q)V_q(q)2\Phi(q) \leq 0,$$

given that $\Phi(q) > 0$ and $V_q(q) \geq 0$. We proceed in cases:

- **Case** $V_q(q) = 0$: If $V_q(q) = 0$, we are solving the static problem for which we know $\delta(q) = 1$ is a dominant strategy. Therefore, if $V_q(q) = 0$, the only equilibrium is $\delta^*(q) = 1$.
- **Case** $V_q(q) > 0$: We prove that there is no interior equilibrium by contradiction. Assume that for some $q \in \Phi^+$, there is such equilibrium with an interior announcement, say $\delta^*(q) \in (0, 1)$. Because $V_q(q) > 0$, the solution to the first order conditions is a maximum. This solution, given opponent's strategy δ^{-i} , is

$$\delta^i = \delta^{-i} + \frac{1}{2\Phi(q)} \left[\frac{\Phi(q)}{rq(1-q)V_q(q)} - \frac{1}{2} \right] = \delta^{-i} + \Theta(q),$$

where $\Theta(q) \equiv \frac{1}{2\Phi(q)} \left[\frac{\Phi(q)}{rq(1-q)V_q(q)} - \frac{1}{2} \right]$. Setting δ^{-i} equal to the best response of the player $-i$, we have that at an interior equilibrium announcement it must hold:

$$\delta^i = \delta^i + 2\Theta(q) \Rightarrow \Theta(q) = 0 \Rightarrow V_q(q) = \frac{2\Phi(q)}{rq(1-q)}. \quad (\text{A.23})$$

If δ^* is an equilibrium announcement, it must also hold

$$\begin{aligned} \psi^N q + \frac{1}{2} + rq(1-q)(1-q-\delta^*)V_q(q) &\geq \\ \psi^N q + \frac{1}{2} + \Phi(q)(\delta - \delta^*) + rq(1-q) \left(1 - q - \left[\frac{1}{2}(\delta + \delta^*) + \Phi(q)(\delta - \delta^*)^2 \right] \right) V_q(q), \end{aligned}$$

for all $\delta \in [0, 1]$, which implies

$$rq(1-q)V_q(q) \left(\frac{1}{2} + \Phi(q)(\delta - \delta^*) \right) (\delta - \delta^*) \geq \Phi(q)(\delta - \delta^*)$$

In particular, for $\delta \neq \delta^*$, we have that

$$rq(1-q)V_q(q) \left(\frac{1}{2} + \Phi(q)(\delta - \delta^*) \right) = \Phi(q) + rq(1-q)V_q(q)\Phi(q)(\delta - \delta^*) \geq \Phi(q),$$

where we have substituted $\frac{1}{2}rq(1-q)V_q(q) = \Phi(q)$. Simplifying the previous expression it must hold

$$rq(1-q)V_q(q)\Phi(q)(\delta - \delta^*) \geq 0,$$

for all $\delta \in [0, 1], \delta \neq \delta^*$. Given that $V_q(q) > 0$ and $\Phi(q) > 0$, which implies $q < 1$, for $q > 0$ the previous statement can only be true for all $\delta \in [0, 1]$ if and only if $\delta^* = 0$, which is not interior, a contradiction. For $q = 0$, it is straightforward to see $\delta^* = 1$ is the only equilibrium.

Hence because also $\delta^*(q) \notin (0, 1)$ for any $q \in \Phi^+$, we finally conclude $\delta^*(q) \in \{0, 1\}$ for all $q \in [0, 1]$.

We are left to prove that for all $q \in [0, 1]$ there is a unique symmetric equilibrium announcement. From the previous discussion we already know the only equilibrium announcement for q such that $\Phi(q) \leq 0$ is $\delta^*(q) = 0$. Let's look at q such that $\Phi(q) > 0$. If $V_q(q) = 0$, it is straightforward to see that $\delta = 1$ is a dominant strategy and therefore $\delta^*(q) = 1$ is the unique equilibrium announcement. Now suppose that $V_q(q) > 0$, so first-order conditions characterize a maximum and the best-response function is given by

$$BR(\delta) = \begin{cases} 1 & \text{if } \delta + \Theta(q) > 1 \\ \delta + \Theta(q) & \text{if } \delta + \Theta(q) \in (0, 1) \\ 0 & \text{if } \delta + \Theta(q) < 0, \end{cases}$$

$$\text{where } \Theta(q) = \frac{1}{2\Phi(q)} \left[\frac{\Phi(q)}{rq(1-q)V_q(q)} - \frac{1}{2} \right].$$

Observe we can characterize equilibrium announcements by the fixed-point of the opponent's best response to a player's best response. In a symmetric equilibrium, the best-response functions of both players are identical, so equilibrium announcements are characterized as the fixed point of the best-response composition

$$\delta = BR(BR(\delta)) = \begin{cases} 1 & \text{if } \delta + 2\Theta(q) > 1 \\ \delta + 2\Theta(q) & \text{if } \delta + 2\Theta(q) \in (0, 1) \\ 0 & \text{if } \delta + 2\Theta(q) < 0. \end{cases}$$

It immediately follows:

- If $\Theta(q) > 0$, the only fixed point of the best-response composition is $\delta = 1$. In this case, the unique equilibrium is $\delta^*(q) = 1$.
- If $\Theta(q) < 0$, the only fixed point of the best-response function is $\delta = 0$. In this case, the unique equilibrium is $\delta^*(q) = 0$.
- If $\Theta(q) = 0$, we use the previous discussion to show that there are no interior solutions and that the unique equilibrium announcement is $\delta^*(q) = 0$.

■

With the following proposition we show that the equilibrium policy is also characterized by a bang-bang as in the baseline model, with the property that the threshold of the dynamic electoral competition game, \tilde{q}_D is always below the threshold of the static electoral competition game, \tilde{q}_S .

Proposition 23 *A $q' < \tilde{q}_S$ exists such that $\delta^*(q) = 0$ is the equilibrium announcement for $q > q'$.*

Proof.

Let's verify the announcement $\delta = 0$ can be supported inside Φ^+ in a neighborhood around \tilde{q}_S as an equilibrium. If that is the case, equation A.19 must satisfy

$$\rho V(q) = \psi^N q + \frac{1}{2} + rq(1-q)^2 V_q(q).$$

Recall the solution of this differential equation is given by

$$V(q) = \frac{\psi^N q}{\rho} + \frac{1}{2\rho} + \frac{1}{m(q)} \int_q^1 \frac{\psi^N}{\rho} m(q) dq,$$

where $m(q)$ is the corresponding integrating factor.

From Proposition 22 we know that $\delta^*(q) = 0$ can be supported as an equilibrium announcement if $\Theta(q) < 0$, or equivalently $V_q(q) > \frac{2\Phi(q)}{rq(1-q)}$ for $q \in \Phi^+$. We proceed to check this condition for the previous solution of the value function $V(q)$.

Taking derivatives with respect to q from the previous expression, we obtain that

$$V_q(q) = \psi^N \left[\frac{2}{\rho} + \frac{1}{rq(1-q)^2 m(q)} \int_q^1 m(x) dx \right].$$

where we have used the fact that $m_q(q) = -\frac{\rho}{rq(1-q)^2} m(q)$. Observe that at \tilde{q}_S , the previous condition is satisfied for any $\psi^N \geq 0$ as

$$\psi^N \left[\frac{2}{\rho} + \frac{1}{\tilde{q}_S(1-\tilde{q}_S)^2 m(\tilde{q}_S)} \int_{\tilde{q}_S}^1 m(x) dx \right] > \psi^N \frac{2}{\rho} \geq 0 = 2 \frac{2\Phi(\tilde{q}_S)}{\tilde{q}_S(1-\tilde{q}_S)}.$$

Moreover, if $\psi^N > 0$ by continuity of $\frac{2\Phi(q)}{q(1-q)}$ around \tilde{q}_S , this condition is satisfied in an open neighborhood of \tilde{q}_S , say $\mathcal{O}(\tilde{q}_S)$. That is, there exist $q' \in \Phi^+$ with $q' < \tilde{q}_S$, such that $\delta^*(q) = 0$ can be sustained as an equilibrium announcement for $q > q'$. ■

This final proposition completes the characterization:

Proposition 24 $\delta^*(q)$ is decreasing, with at most one jump from 1 to 0 at some \tilde{q}_D . Moreover, $\delta^*(\tilde{q}_D) = 1$.

Proof. Assume there is more than one jump, that is there is \hat{q} such that we have a jump from 0 to 1. It must be the case $\hat{q} < \tilde{q}_S$ because we know that $\delta^*(q) = 0$ for all $q \geq \tilde{q}_S$, and so $\hat{q} \in \Phi^+$. By continuity of the value function, at any discontinuity of $\delta^*(q)$ it must hold

$$\lim_{q \rightarrow \hat{q}^-} \psi^N q + \frac{1}{2} + q(1-q)^2 V_q(q) = \lim_{q \rightarrow \hat{q}^+} \psi^N q + \frac{1}{2} - q^2(1-q)^2 V_q(q)$$

which implies $V_q(\hat{q}) = 0$. In this case, for any discontinuity $q \in \Phi^+$, we know $\delta = 1$ is a dominant strategy because $V_q(q) = 0$.

It also follows,

$$\lim_{q \rightarrow \hat{q}} q(1-q)V_q(q) = 0.$$

But we can find an $\epsilon > 0$ such that $\delta^*(q) = 0$ for all with $(\hat{q} - \epsilon, \hat{q}) \subseteq \Phi^+$. Therefore, from Proposition 22 it must hold

$$\Theta(q) \leq 0 \implies 0 < 2\Phi(q) \leq q(1-q)V_q(q),$$

for $q \in (\hat{q} - \epsilon, \hat{q})$. Taking limits $q \rightarrow \hat{q}^-$

$$0 \leq 2\Phi(\hat{q}) \leq \hat{q}(1 - \hat{q})V_q(\hat{q}) = 0$$

which implies $\Phi(\hat{q}) = 0$, and so $\hat{q} = \tilde{q}_S$, a contradiction. The last discussion proves that only one discontinuity of $\delta^*(q)$ exists, which we denote by \tilde{q}_D . Therefore it must hold that $V_q(\tilde{q}_D) = 0$. ■

Putting all results together we finally obtain the result of Theorem 6 which shows that the equilibrium policy is also defined as a threshold policy:

Theorem 6 *There is a unique equilibrium in symmetric strategies of the dynamic electoral competition game with nation building motives. The equilibrium strategies are described as a threshold policy given by \tilde{q}_D such that:*

$$\delta^{A^*}(q) = \delta^{B^*}(q) = \begin{cases} 1 & \text{if } q \leq \tilde{q}_D \\ 0 & \text{if } q > \tilde{q}_D. \end{cases}$$

with $0 < \tilde{q}_D < \tilde{q}_S$, where \tilde{q}_S defines the threshold of the symmetric equilibrium for the static electoral competition game, and is given by:

$$\tilde{q}_S = \frac{\phi^R}{\phi^N + \phi^R}$$

Proof. By Proposition 22, we know that either $\delta^*(q) = 0$ or $\delta^*(q) = 1$. We also know $\delta^*(0) = 1$ and the function $\delta^*(q)$ is decreasing with only one discontinuity at some \tilde{q}_D , such that $\delta^*(\tilde{q}_D) = 1$ from Proposition 24. Moreover, from Proposition 23, for any $\psi^N > 0$ we know a $q < \tilde{q}_S$ exists such that $\delta^*(q) = 0$. This implies that $\tilde{q}_D < \tilde{q}_S$. Then, the equilibrium policy is given by

$$\delta^*(q) = \begin{cases} 1 & \text{if } q \leq \tilde{q}_D \\ 0 & \text{if } q > \tilde{q}_D. \end{cases}$$

■

The previous result delivers a very striking property of the optimal equilibrium path: Because no equilibrium announcement is interior, the only long-run steady states are $q = 0$ or $q = 1$.

Proposition 25 *The long-run steady states of the nation-building electoral-competition game are located at 0 and 1.*

Proof. Equilibrium announcements are never interior. Therefore, for every

$q \in (0, 1)$, $\delta^*(q) \neq 1 - q \Rightarrow \dot{q} \neq 0$. Therefore no interior steady state exists under equilibrium strategies. ■

Proof of Proposition 6

Proposition 6 *In the case that parties only care about office, $\psi^N = 0$, it holds*

$$\tilde{q}_D = \tilde{q}_S.$$

On the contrary, if parties only care about nation-building, $O = 0$, it holds

$$\tilde{q}_D = 0.$$

Proof. First we find the symmetric equilibrium for the special cases $\psi^N = 0$ and $O = 0$. Consider first that there is no nation-building motive, that is $\psi^N = 0$. In every period both parties compete to split a pie of size 1 and therefore the sum of payoffs is constant for every period. It immediately follows that the sum of discounted payoffs is also constant and therefore the previous game is a constant-sum game for any starting point $q \in [0, 1]$. As in any constant-sum game, all equilibria are payoff-equivalent, with payoff $V(q) = \frac{1}{2\rho}$ for all q . Therefore, continuation values are independent of future q and equilibrium strategies must maximize the per-period probability of winning elections. More concretely, because all equilibrium payoffs are given by $V(q) = \frac{1}{2\rho}$ for all $q \in [0, 1]$, then $V_q(q) = 0$ and the HJB equation of the dynamic problem collapses to the problem of the static game. Then, the equilibrium strategies are given by:

$$\delta^{i*}(q) = \arg \max_{\delta^i} p^i(\delta^i, \delta_{-i}, q) = \begin{cases} 1 & \text{if } \Phi(q) > 0 \\ [0, 1] & \text{if } \Phi(q) = 0 \\ 0 & \text{if } \Phi(q) < 0, \end{cases}$$

which is independent of $\delta^{*-i}(q)$. Moreover, because $\frac{\partial \Phi(q)}{\partial q} < 0$ and dynamics satisfy $\dot{q}_t > 0$ if $\delta = 1$, $\dot{q}_t < 0$ if $\delta = 0$, we have that on equilibrium Φ^- and Φ^+ are invariant sets, meaning that under equilibrium strategies

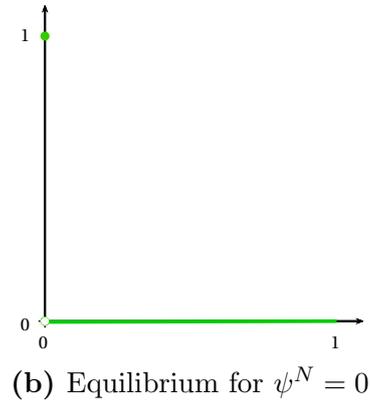
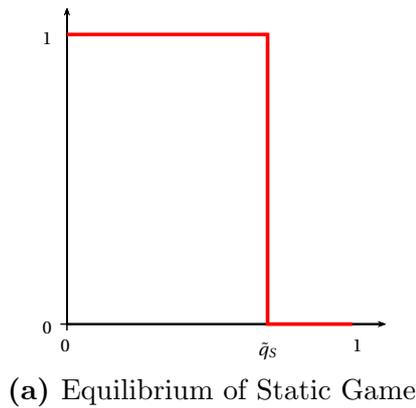
$$q_t \in \Phi^- \iff q_0 \in \Phi^-, \quad \forall t \geq 0.$$

Therefore, for given initial q_0 , the equilibrium path for $\{\delta_t^*\}_{t \geq 0}$ is fully defined by:

$$\delta_t^* = \begin{cases} 1 & \text{if } q_0 < \tilde{q}_S \\ [0, 1] & \text{if } q_0 = \tilde{q}_S \\ 0 & \text{if } q_0 > \tilde{q}_S, \end{cases}$$

which is exactly the same solution as in the static electoral competition game. Moreover, as in the baseline model with a secure government, the long-run steady states are located at $q = 0$ and $q = 1$, and which one occurs is only determined by the initial q_0 .

Consider now the other limiting case, when political parties are not office motivated, i.e. $W^i(q, \delta^i, \delta^{-i}) = \psi^N q$. In this game, it can be easily checked with the HJB equations that the unique Markov-perfect equilibrium is $\delta^A(q) = \delta^B(q) = 0, \forall q$ and $\forall t \geq 0$. Intuitively, parties have aligned nation-building incentives and do not care about winning elections per se. Given that both derive benefits from increasing q , it is optimal for them to do it in the fastest way. ■



A.2 Supplementary Appendix

A.2.1 Proof of Lemma 1

Lemma 1 *Under imperfect empathy and quadratic costs*

$$C(e) = \frac{1}{2}e^2,$$

the optimal socialization efforts are given by

$$e_t^N = (1 - q_t)(1 - \delta_t)r \quad e_t^R = q_t\delta_t r,$$

and the law of motion for cultural transmission is

$$\dot{q} = q_t(1 - q_t)(e_t^N - e_t^R) = rq_t(1 - q_t)(1 - \delta_t - q_t).$$

Proof.

A parent of trait i obtains utility V^{ij} if her child holds identity j . The imperfect empathy assumption implies parents evaluate children's actions using their own utility function. We assume a children of type i derives utility from private consumption but only consumes the public good associated to her identity (as the other provides zero utility). Therefore, for each combination of $i, j \in \{N, R\}$ one has

$$\begin{aligned} V^{NN} &= f(1 - r) + (1 - \delta_t)r & V^{RR} &= f(1 - r) + \delta_t r \\ V^{NR} &= f(1 - r) & V^{RN} &= f(1 - r). \end{aligned} \quad (\text{A.24})$$

Therefore, parents do not derive any utility from seeing their children consuming the club public good associated with the other identity.

Parents socialization problem for a type i parent is given by

$$\max_{e \in [0,1]} P_t^{ii}(e)V^{ii} + (1 - P_t^{ii}(e))V^{ij} - \frac{1}{2}e^2, \quad (\text{A.25})$$

with optimal socialization efforts

$$\begin{aligned} e_t^N &= (1 - q_t)\Delta V^N = (1 - q_t)g_t^N = (1 - q_t)(1 - \delta_t)r \\ e_t^R &= q_t\Delta V^R = q_tg_t^R = q_t\delta_t r, \end{aligned} \quad (\text{A.26})$$

where $\Delta V^i = V^{ii} - V^{ij}$. Observe that for a parent of type i the optimal socialization effort depends positively on ΔV^i and negatively on q^i . The term ΔV^i is the degree of *cultural intolerance*, which increases in the level of own identity public good. Next we construct the evolutionary dynamics of cultural traits. Between t and $t + dt$, a fraction dt of the population dies and is replaced by the same number of new agents. Hence, at each point in time, type N proportion is given by the remaining parents of type N plus the fraction of newly born children inheriting trait N . Therefore, the fraction of agents with

a national trait at time $t + dt$, q_{t+dt} , is

$$q_{t+dt} = (1 - dt)q_t + dt[q_t P^{NN} + (1 - q_t)P^{RN}]. \quad (\text{A.27})$$

Recall that transition probabilities P^{NN} and P^{RN} are given by

$$P_t^{NN}(e_t^N) = e_t^N + (1 - e_t^N)q_t \quad P_t^{RN}(e_t^R) = (1 - e_t^R)q_t \quad (\text{A.28})$$

Using A.26, A.27 and A.28, and taking $dt \rightarrow 0$, we obtain the following differential equation for q_t

$$\dot{q} = q_t(1 - q_t)(e_t^N - e_t^R) = q_t(1 - q_t)(1 - \delta_t - q_t)r. \quad (\text{A.29})$$

■

A.2.2 Micro-foundations for the rates of protests

In this section we provide microfoundations for the individual decision on whether to participate in protests and we present different alternatives on how protests affect the objective function of the government.

Participation rate in protests

We rely on a stylized version of the model of political unrest developed by Passarelli and Tabellini (2017). As in their model, we assume that individuals engage in political unrest if the benefits of participating are greater than the costs. We also assume that the benefits of protesting are purely emotional rewards. That is, individuals join protests due to feelings of aggrivement and to the psychological reward that participating in protests provides to the individual. Following Passarelli and Tabellini (2017), we assume that individuals with identity i feel entitled to a particular policy $\widehat{g^i(\delta_t)}$. If this “reference” point is not implemented, individuals experiment a sense of injustice that causes them anger and frustration. The psychological reward of joining others in a protest is concomitant to this feeling of being treated unfairly. The further away actual policy is from their ideal point of a group of citizens, the more aggrieved they feel and the more they enjoy protesting.⁶

⁶One could argue that the choice to participate in a riot or a civil conflict should be based on individual expectations about how joining a protest changes the policy choices of the

Formally, the emotional benefit of protesting $B^i(\cdot)$ is a function of the distance between their ideal policy $\widehat{g^i(\delta_t)}$ and the actual policy $g^i(\delta_t)$. In principle, emotional benefits could also depend on how many members from the group participate. Therefore, individual benefits from protesting are given by

$$B^i(\widehat{g^i(\delta_t)}, g^i(\delta_t), q_t^i) = F(\text{dist}(\widehat{g^i(\delta_t)}, g^i(\delta_t)), q_t^i) = \text{dist}(\widehat{g^i(\delta_t)}, g^i(\delta_t)) \times h(q_t^i)$$

with dist defined as some distance, and $h(\cdot)$ an arbitrary function to be defined later. This specification allows for several specifications depending on the choice of $\text{dist}(\cdot)$, $h(\cdot)$ and $\widehat{g^i(\delta_t)}$.

However, joining protests is costly. Concretely, we assume that individuals in group i face cost c , independently drawn from some distribution F . These costs capture common features such as repression as well as idiosyncratic costs, such as foregone income from not working. Thus, individual j in group i participates in protests if and only if $B^i(\widehat{g^i(\delta_t)}, g^i(\delta_t), q_t) - c^{ij} \geq 0$. Hence, if $c^{ij} \sim U[0, 1]$, the individual probability of engaging in protests is given by

$$p_t^i = Pr(c^{ij} \leq B^i(\widehat{g^i(\delta_t)}, g^i(\delta_t), q_t)) = B^i(\widehat{g^i(\delta_t)}, g^i(\delta_t), q_t)$$

Therefore, the total participation rate $P^i(\delta_t, q_t)$ in protests of group i is given by

$$D^i(\delta_t, q_t) = q_t^i \times p_t^i = q_t^i \times B^i(\widehat{g^i(\delta_t)}, g^i(\delta_t), q_t)$$

Finally, as we discuss below, protests affect the objective function of the central government, either by creating a direct welfare loss for the government, or indirectly by generating dead-weight losses for citizens which in turn are internalized by a welfarist government.

central government. Although we recognize that this “instrumental” motive has its merits, we believe that it is not very relevant in our context. In a sufficiently large and heterogeneous population of potential protesters, which is generally the case in our context, the marginal impact of one more individual protesting in the decision of the government is negligible. Hence, an atomistic individual is unlikely to take this costly political action. Given that the expected change in welfare through influencing policy choices is close to zero, and in the absence of any explicit material gain of protesting, the benefit from protesting must come from psychological or social rewards. In our case, as argued by Laitin (2007), a key feature of national identities is the willingness that creates on individuals to engage in costly political actions, in order to defend their own nation for the psychological reward that provides and despite obvious material losses.

Benchmark case

In the benchmark case, we assume the following

- $dist(x, y) = |x - y|$
- $h(q_t^i) = 1$
- $\widehat{g^N}(\delta_t) = r$ and $\widehat{g^R}(\delta_t) = r$

That is, the benefits of protesting depend linearly on the distance between the ideal policy and the policy implemented, and individual emotional rewards are orthogonal to the number of individuals participating.⁷ Also, we assume an extreme polarization of preferences, in the sense that members of each group feel entitled to a level of public good equal to the total tax collection in the region i.e. the ideal δ^i for each group is $\widehat{\delta}_t = 0$ for type N and $\widehat{\delta}_t = 1$ for type R. Therefore, we have that

$$D_t^N(\delta_t, q_t) = q_t \left[\underbrace{r}_{\text{Ideal}} - \underbrace{(1 - \delta_t)r}_{\text{Real}} \right] = q_t \delta_t r$$

$$D_t^R(\delta_t, q_t) = (1 - q_t) \left[\underbrace{r}_{\text{Ideal}} - \underbrace{\delta_t r}_{\text{Real}} \right] = (1 - q_t)(1 - \delta_t)r$$

Finally, for the baseline case we assume that the government directly experiences a loss of welfare which is proportional to the participation in protests of both groups. Therefore, the utility function of the central government is

$$W(q_t, \delta_t) = \psi^N q_t + \alpha q_t U^N(\delta_t) + (1 - \alpha)(1 - q_t) U^R(\delta_t) - (\beta q_t \delta_t r + (1 - \beta)(1 - q_t)(1 - \delta_t)r),$$

where β and $1 - \beta$ capture the disruptions created by protests, which inflict a direct loss of social welfare to the central government. In this setting, β is a measure for the relative impact of protests of group N with respect to group R ,

⁷In all the specifications of the protest function we assume that the individual decision about participating in protests is independent of the number of members from her group joining the protest i.e. $h(q_t^i) = 1$. An interesting possibility is to allow for complementarities in protests. Concretely, we could assume that the individual emotional benefit increases with the number of individuals that also participate in protests i.e. $h(q_t^i) = p_t^i q_t^i$, where p_t^i is the average participation rate at time t of individuals in group i . Nevertheless, the main qualitative results of the paper are robust to these type of protests.

and it comprises factors such as how organized individuals are, the capacity of regional cultural leaders to mobilize people along identity cleavages, the physical resources they have to cause disruption, their influence on media or the support they have from international public opinion.⁸

Quadratic case

In section 1.5 we illustrate how the results of the model change when we relax the linearity assumption of the objective function. Concretely, we keep the rest of the assumptions but we have that $dist(x, y) = (x - y)^2$ instead of $dist(x, y) = |x - y|$. Therefore, protests are given by

$$D_t^N(\delta_t, q_t) = q_t \left[\underbrace{r}_{\text{Ideal}} - \underbrace{(1 - \delta_t)r}_{\text{Real}} \right]^2 = q_t \delta_t^2 r^2$$

$$D_t^R(\delta_t, q_t) = (1 - q_t) \left[\underbrace{r}_{\text{Ideal}} - \underbrace{\delta_t r}_{\text{Real}} \right]^2 = (1 - q_t)(1 - \delta_t)^2 r^2$$

Different ideal point

The previous choice of the ideal point, which is a maintained assumption throughout the paper, corresponds to a very extreme case in which individuals in both groups are entirely selfish. However, considering ideal points that involve some sharing of resources may be more reasonable for some real-world examples.⁹ Moreover, it may be that this assumption is behind the full homogenization result, as it introduces a strong conflict over resources. Nevertheless, it turns out that our homogeneity results are robust to ideal points that incorporate some fairness concerns.

To see this, consider that protests have the same structure as in the benchmark model but ideal points are defined as follows

⁸In the context of the paper, D_t^N and D_t^R can also capture the idea that political unrest above some threshold could generate violent civil conflict and a secessionist attempt in the peripheral region. Then, the participation rate can be interpreted as the probability of reaching that turning point.

⁹We thank the editor for his suggestion about checking the robustness of the results to less extreme choices of ideal points.

$$\widehat{g^i(\delta_t)} = r(\lambda^N + (1 - \lambda^N)q_t)$$

$$\widehat{g^i(\delta_t)} = r(\lambda^R + (1 - \lambda^R)(1 - q_t))$$

where a higher $\lambda^i \in [0, 1]$ implies a higher degree of selfishness of individuals in group i . Note that

$$\lim_{\lambda^i \rightarrow 1} \widehat{g^i(\delta_t)} = r$$

$$\lim_{\lambda^i \rightarrow 0} \widehat{g^i(\delta_t)} = q_t^i r$$

Therefore, the formulation of ideal points has two extreme cases: 1) the one in the paper, where citizens are entirely selfish; 2) the “perfectly fair” case, where individuals feel entitled to get in public goods a fraction of the budget equal to the size of their group in the population. The value of λ^i captures the self-serving bias of the individuals in group i , as individuals judgments combine what is fair and what is beneficial for them.

Now, consider a situation where individuals protest whenever the policy deviates from their bliss point, even if it is beneficial to them. For comparability with results in Section 5, also consider quadratic protests. The protest functions are given by

$$D_t^N(\delta_t, q_t) = q_t [\underbrace{\widehat{g^N(\delta_t)}}_{\text{Ideal}} - \underbrace{g^N(\delta_t)}_{\text{Real}}]^2 = q_t [r(\lambda^N + (1 - \lambda^N)q_t) - (1 - \delta_t)r]^2$$

$$D_t^R(\delta_t, q_t) = (1 - q_t) [\underbrace{\widehat{g^R(\delta_t)}}_{\text{Ideal}} - \underbrace{g^R(\delta_t)}_{\text{Real}}]^2 = (1 - q_t) [r(\lambda^R + (1 - \lambda^R)(1 - q_t)) - \underbrace{\delta_t r}_{\text{Real}}]^2$$

From now on, we assume that $\lambda^i = \lambda, \forall i$, as it simplifies the algebra (but the results below hold for any combination of λ^N and λ^R).

If the function $H(q)$ for this problem is strictly convex for all q , then Theorem 3 holds, so long run steady states are homogeneous. Recall that $H(q)$ gives the per-period utility derived from the policy $\delta(q)$ that keeps q unchanged. We have that

$$H(q) = \psi^N q + (\alpha q + (1 - \alpha)(1 - q))f(1 - r) + r(\alpha q^2 + (1 - \alpha)(1 - q)^2) - r^2 \left[\beta q_t [(\lambda + (1 - \lambda)q_t) - q_t]^2 + (1 - \beta)(1 - q_t) [(\lambda + (1 - \lambda)(1 - q_t)) - (1 - q_t)]^2 \right].$$

The second derivative of this function is given by

$$H''(q) = 2r + 2r^2\lambda^2[\beta(1 - q) + (1 - \beta)q - (1 - 2q)(1 - 2\beta)].$$

Observe for all $\beta, q, \lambda, r \in [0, 1]$

$$\beta(1 - q) + (1 - \beta)q - (1 - 2q)(1 - 2\beta) \geq -1.$$

Hence,

$$H''(q) \geq 2r - 2r^2\lambda^2 = 2r(1 - r\lambda^2) \geq 0.$$

Therefore, $H(q)$ does not have a maximum in $[0, 1]$ for any choice of ideal point. Hence, long-run steady states are culturally homogeneous.¹⁰

In conclusion, allowing for ideal points that involve some sharing of resources does not alter the full-homogenization result. When the two groups have closer views about what they are entitled to (lower λ), the zero-sum conflict is weakened because the government can reduce the utility losses coming from protests by choosing a value of δ close to the ideal point of both groups. However, the conflict never completely disappears as long as there is a heterogeneous distribution of identities. The reason is that it is unavoidable for the government to pick winners and losers, as a larger provision of one public good always comes at the expense of a reduction in the other public good. Therefore, the government can only avoid dealing with conflicting motives by homogenizing the population.

¹⁰Another possibility is to assume that individuals only protest when the deviation is detrimental for them. In this case, they may do nothing (zero protests) or they may show support for the government if it benefits them (“positive” protests). That is

$$D_t^i(\widehat{g^i(\delta_t)}, g^i(\delta_t), q_t^i) = q_t^i \max \{ \widehat{g^i(\delta_t)} - g^i(\delta_t), 0 \}$$

or

$$D_t^i(\widehat{g^i(\delta_t)}, g^i(\delta_t), q_t^i) = q_t^i [\widehat{g^i(\delta_t)} - g^i(\delta_t)]$$

Although we do not present it here, the same result goes through if we consider these alternative formulations of the protest function, *for any choice of ideal point*.

Alternative rationales for the objective function

One could think of alternative rationales for how protests affect the objective function of the government. One possibility is to assume that citizens experience a direct intrinsic utility loss from seeing the other group protesting, which in turn is internalized by the government, as it cares about the utilities of individuals. In the same way as protesting to defend one's identity provides an emotional reward (by singing the anthem, carrying the flag, etc...), seeing protests by the group with the oppositional identity can create feelings of anger and reductions of self and group-esteem. Let β and $1 - \beta$ be the marginal disruption created by protests of groups N and R , respectively. Then we can write

$$\begin{aligned} U^N(\cdot) &= f((1 - r)) + (1 - \delta_t)r - (1 - \beta)D^R(\delta_t, q_t) \\ U^R(\cdot) &= f((1 - r)) + \delta_t r - \beta D^N(\delta_t, q_t) \end{aligned}$$

Therefore, we have

$$\begin{aligned} W &= \psi^N q_t + \alpha q_t U_t^N(\cdot) + (1 - \alpha)(1 - q_t) U_t^R(\cdot) \\ &= \psi^N q_t + \alpha q_t [f((1 - r)) + (1 - \delta_t)r - (1 - \beta)D^R(\delta_t, q_t)] \\ &\quad + (1 - \alpha)(1 - q_t) [f((1 - r)) + \delta_t r - \beta D^N(\delta_t, q_t)] \end{aligned}$$

We can see that this objective function is similar to the previous one, with a higher order term for q on the protest side. In this case, protests of both groups are higher at intermediate values of q , which makes homogeneous steady-states more desirable.

Another alternative is to assume that, in order to keep order and counteract the disruptive costs of protests, the government uses revenue from taxes, which is taken away from the total public budget used to provide public goods. To keep comparability, we can assume that in order to repair the damage created by protests, the government needs to employ a fraction ζ and a fraction η of the public budget r to counteract protest by N and R respectively. Therefore, $g_t^N + g_t^R = r[1 - \zeta D^N(q_t, \delta_t) - \eta D^R(q_t, \delta_t)]$.¹¹ We can also assume that, in addition

¹¹We need to assume that η and ζ are sufficiently small so that $g_t^N + g_t^R \geq 0$

to the destruction of public goods, riots have an effect on disposable (after tax) income. For instance, this would be due to the shutdown of economic activity, the increase in risk premium of bonds or the destruction of physical capital needed to generate income. In both cases, we will get a very similar objective function.

These different rationales tell slightly different plausible stories about the processes of nation-building. However, the different models are formally equivalent and their qualitative results identical. In some sense, one can move from one to another by relabelling parameters, as the key results are robust to the chosen specification.

A.2.3 Proofs of Propositions 2, 3, and 4

Proof of Proposition 2

First we prove the difference in welfare of the two policies is decreasing at $q_0 = \bar{q}_0$; that is

$$\frac{\partial}{\partial q} F(\bar{q}_0) < 0.$$

First, we rule out $\frac{\partial}{\partial q} F(\bar{q}_0) = 0$. Simply observe for any variable x

$$F(\bar{q}_0(x), x) \equiv 0$$

Hence, for any variable x , it follows

$$\frac{\partial}{\partial q} F(\bar{q}_0(x), x) \frac{\partial}{\partial x} \bar{q}_0(x) + \frac{\partial}{\partial x} F(\bar{q}_0(x), x) = 0.$$

It is easy to verify $\frac{\partial}{\partial q} F(\bar{q}_0(x), x) < 0$. Assume for a contradiction $\frac{\partial}{\partial q} F(\bar{q}_0) \geq 0$. Then, because F is continuous in q_0 and $F(\bar{q}_0) = 0$ it must be the case that $\exists q'_0 > \bar{q}_0$ such that $F(q'_0) \geq 0$. Because F is continuous with $F(1) < 0$, we contradict the result that $F(\cdot)$ has a unique zero.

Now because $\frac{\partial}{\partial q} F(\bar{q}_0(x), x) \neq 0$, we can write

$$\frac{\partial}{\partial x} \bar{q}_0 = - \left(\frac{\partial}{\partial q} F(\bar{q}_0) \right)^{-1} \frac{\partial}{\partial x} F(\bar{q}_0).$$

But since $\frac{\partial}{\partial q} F(\bar{q}_0) < 0$, we have that

$$\text{sign}\left(\frac{\partial}{\partial x} \bar{q}_0(\omega)\right) = \text{sign}\left(\frac{\partial}{\partial x} F(\bar{q}_0(\omega); \omega)\right).$$

Hence, for parameter x , we only need to check the sign of

$$\frac{\partial}{\partial x} F(\bar{q}_0(x), x).$$

Recall that

$$\begin{aligned} W^\delta(q; \omega) &= \psi^N q + \\ &\quad + \alpha q (f(1-r) + (1-\delta)r) \\ &\quad + (1-\alpha)(1-q)(f(1-r) + \delta r) \\ &\quad - \beta q \delta r \\ &\quad - (1-\beta)(1-q)(1-\delta)r, \end{aligned}$$

where ω is a vector including all the parameters. Let's denote

$$\begin{aligned} S(q^1, q^0, \omega) &= W^1(q^1(q_0; \omega); \omega) - W^0(q^0(q_0; \omega); \omega) \\ &= \psi^N (q^1 - q^0) + (2\alpha - 1)f(1-r)(q_1 - q_0) + r((1-\alpha)(1-q_1) - \alpha q_0) \\ &\quad + r((1-\beta)(1-q_0) - \beta q_1) \end{aligned}$$

It follows

$$\frac{\partial}{\partial x} F(q_0; \omega) = \int_0^\infty \frac{\partial}{\partial x} e^{-\rho t} S(q_t^1(\omega), q_t^0(\omega), \omega) dt,$$

for any parameter x . Next, we do comparative statics on the parameters of the model.

- We begin with the comparative statics for ψ^N . These are as follows

$$\frac{\partial}{\partial \psi^N} F(\bar{q}_0) = \int_0^\infty e^{-\rho t} (q_t^1 - q_t^0) dt < 0,$$

because we always have that $q_t^0 > \bar{q}_0 > q_t^1$ for all $t > 0$. Therefore

$$\frac{\partial \bar{q}_0}{\partial \psi^N} < 0.$$

- Next, we do comparative statics for α . These are as follows

$$\frac{\partial}{\partial \alpha} F(\bar{q}_0) = \int_0^\infty e^{-\rho t} \{2f(1-r)(q_t^1 - q_t^0) - r(q_t^0 + (1 - q_t^1))\} dt < 0,$$

hence

$$\frac{\partial \bar{q}_0}{\partial \alpha} < 0.$$

- Clearly, for parameter β , we obtain similar results

$$\frac{\partial}{\partial \beta} F(\bar{q}_0) = \int_0^\infty e^{-\rho t} \{ -r((1 - q_t^0) + q_t^1) \} dt < 0,$$

Therefore

$$\frac{\partial \bar{q}_0}{\partial \beta} < 0.$$

- Now if utility of consumption is given by $f(1 - r) = \frac{(1-r)^{1-\sigma}}{1-\sigma}$, where $\theta \in (0, 1), \sigma > 0$, it holds

$$\begin{aligned} \frac{\partial}{\partial \theta} F(\bar{q}_0) &= \int_0^\infty e^{-\rho t} \left\{ (2\alpha - 1) \frac{(1-r)^{1-\sigma}}{1-\sigma} (q_t^1 - q_t^0) \right\} dt > 0 \iff \alpha < \frac{1}{2} \\ \frac{\partial}{\partial \sigma} F(\bar{q}_0) &= \int_0^\infty e^{-\rho t} \left\{ (2\alpha - 1) \theta \frac{(1-r)^{1-\sigma}}{(1-\sigma)^2} (1 - \ln(1-r)(1-\sigma)) (q_t^1 - q_t^0) \right\} dt > 0 \\ &\iff \alpha < \frac{1}{2} \end{aligned}$$

Hence, it follows

$$\begin{aligned} \frac{\partial \bar{q}_0}{\partial \theta} > 0 &\iff \alpha < \frac{1}{2}, \\ \frac{\partial \bar{q}_0}{\partial \sigma} > 0 &\iff \alpha < \frac{1}{2}, \end{aligned}$$

because $q_t^1 - q_t^0 < 0$, $(1-r)^{1-\sigma} > 0$, and $1 - \ln(1-r)(1-\sigma) > 0$ for all $r \in (0, 1)$.

Now we show how the comparative statics on ρ and r can go both ways.

Proof of Proposition 3

Next we do comparative statics on ρ . Taking derivatives of $F(q_0)$ with respect to ρ , we obtain the following expression

$$\begin{aligned} \frac{\partial}{\partial \rho} F(q_0) &= \int_0^\infty \frac{\partial}{\partial \rho} e^{-\rho t} S(q_t^1, q_t^0) dt = \\ &= - \int_0^\infty t e^{-\rho t} S(q_t^1, q_t^0) dt. \end{aligned} \tag{A.30}$$

It is easy to see $S(q^1, q^0)$ is bounded. Hence, an $M > 0$ exists such that $|S(q_t^1, q_t^0)| \leq M$. For example, we can pick $M = f(1-r)$ whenever $f(c) = \frac{c^{1-\sigma}}{1-\sigma}$,

$\gamma \geq 0$. Therefore,

$$\left| \frac{\partial}{\partial \rho} F(q_0; \omega) \right| \leq \int_0^\infty t e^{-\rho t} |S(q_t^1(\omega), q_t^0(\omega), \omega)| dt < \int_0^\infty t e^{-\rho t} M dt = \frac{1}{\rho^2} M < \infty,$$

and the integral A.30 is always well-defined.

Recall that the function $S(q^1, q^0)$ can be written as

$$S(q^1, q^0) = (A_1 q^1 + B_1) - (A_0 q_0 + B_0),$$

with

$$\begin{aligned} A_1 &= \psi^N + (2\alpha - 1)f(1 - r) - (1 - \alpha + \beta)r, & B_1 &= (1 - \alpha)(f(1 - r) + r) \\ A_0 &= \psi^N + (2\alpha - 1)f(1 - r) + (\alpha + 1 - \beta)r, & B_0 &= (1 - \alpha)f(1 - r) - (1 - \beta)r \end{aligned}$$

The sign of the comparative statics on ρ can go both ways as it will depend on the other parameters of the model. Hence, we analyze different cases.

- Assume, α large enough such that $A_1 > 0$, which implies $A_0 > 0$. It is easy to see that

$$S(q_t^1, q_t^0) = A_1 q_t^1 + B_1 - (A_0 q_t^0 + B_0),$$

is strictly decreasing in t with

$$\begin{aligned} \lim_{t \rightarrow \infty} S(q_t^1, q_t^0) &= B_1 - A_0 - B_0 = -\psi^N + (1 - 2\alpha)(f(1 - r) + r) \\ &= -(A_1 + (\alpha + \beta)r) < 0 \end{aligned}$$

Because $\int_0^\infty e^{-\rho t} S(q_t^1, q_t^0) S(q_t^1, q_t^0) dt = 0$ with $S(q_t^1, q_t^0)$ is strictly decreasing, and

$$\lim_{t \rightarrow \infty} S(q_t^1, q_t^0) < 0,$$

a T exists such that $S(q_T^1, q_T^0) = 0$, with $S(q_t^1, q_t^0) > 0$ for all $t \leq T$, and

$S(q_t^1, q_t^0) < 0$ for all $t \geq T$. It follows

$$\begin{aligned}
 \frac{\partial}{\partial \rho} F(\bar{q}_0) &= - \int_0^\infty t e^{-\rho t} S(q_t^1, q_t^0) dt \\
 &= - \int_0^T t e^{-\rho t} S(q_t^1, q_t^0) dt - \int_T^\infty t e^{-\rho t} S(q_t^1, q_t^0) dt \\
 &> - \int_0^T T e^{-\rho t} S(q_t^1, q_t^0) dt - \int_T^\infty ((t - T) + T) e^{-\rho t} S(q_t^1, q_t^0) dt \\
 &= -T \int_0^\infty e^{-\rho t} S(q_t^1, q_t^0) dt - \int_T^\infty (t - T) e^{-\rho t} S(q_t^1, q_t^0) dt \\
 &> 0 - \int_T^\infty (t - T) e^{-\rho t} S(q_t^1, q_t^0) dt > 0,
 \end{aligned}$$

because $t - T \geq 0$ for all $t \geq T$ and $S(q_t^1, q_t^0) < 0$ for all $t > T$. The last inequality implies

$$\frac{\partial}{\partial \rho} \bar{q}_0 > 0.$$

- Assume α small enough such that $A_0 < 0$ which implies $A_1 < 0$. Following a similar argument, we obtain

$$\frac{\partial}{\partial \rho} F(\bar{q}_0) = - \int_0^\infty t e^{-\rho t} S(q_t^1, q_t^0) dt < 0,$$

implying in turn

$$\frac{\partial}{\partial \rho} \bar{q}_0 < 0.$$

In this way, we have shown that the comparative statics on ρ can go in both directions. The following result summarizes the previous discussion

Proposition 26 *The comparative statics on ρ can go both ways and depend on the other parameters of the model*

- If $\psi^N + (2\alpha - 1)f(1 - r) - r(1 - \alpha + \beta) > 0$, then

$$\frac{\partial}{\partial \rho} \bar{q}_0 > 0.$$

- If $\psi^N + (2\alpha - 1)f(1 - r) + r(\alpha + 1 - \beta) < 0$, then

$$\frac{\partial}{\partial \rho} \bar{q}_0 < 0.$$

Proof of Proposition 4

Finally, we do comparative statics on r .

Proposition 27 *The following equality holds*

$$\frac{\partial}{\partial r} F(\bar{q}_0) = -\frac{\rho}{r} \frac{\partial}{\partial \rho} F(\bar{q}_0) + \Lambda \int_0^\infty e^{-\rho t} (q_t^0 - q_t^1) dt \quad (\text{A.31})$$

with $\Lambda = (2\alpha - 1)f'(1 - r) + \frac{1}{r}((2\alpha - 1)f(1 - r) + \psi^N)$.

Proof. Given that r enters the law of motion we have

$$\frac{d}{dr} S(q^1, q^0) = \frac{\partial}{\partial q^1} S \frac{\partial}{\partial r} q^1 + \frac{\partial}{\partial q^0} S \frac{\partial}{\partial r} q^0 + \frac{\partial}{\partial r} S, \quad (\text{A.32})$$

where the first two terms come from r entering in the law of motion and the third term comes from r entering in the function S . First, observe we can write

$$G(q_t) = rt + G(q_0)$$

where $G'(y) = \frac{1}{g(y)}$ with $\dot{q} = rg(q)$. Therefore, taking derivatives with respect to r on both sides of the previous expression

$$\frac{\partial}{\partial r} q_t = \frac{t}{G'(q_t)} = tg(q_t) = \frac{t}{r} \dot{q}_t,$$

Finally observe

$$\begin{aligned} \int_0^\infty e^{-\rho t} \frac{t}{r} \dot{q}_t dt &= \left[\frac{t}{r} e^{-\rho t} q_t \right]_0^\infty - \int_0^\infty \frac{1}{r} e^{-\rho t} (1 - \rho t) q_t dt \\ &= - \int_0^\infty \frac{1}{r} e^{-\rho t} (1 - \rho t) q_t dt, \end{aligned}$$

where we have used integration by parts.

Recall that the function $S(q^1, q^0)$ can be written as

$$S(q^1, q^0) = (A_1 q^1 + B_1) - (A_0 q^0 + B_0),$$

with

$$\begin{aligned} A_1 &= \psi^N + (2\alpha - 1)f(1 - r) - (1 - \alpha + \beta)r, & B_1 &= (1 - \alpha)(f(1 - r) + r), \\ A_0 &= \psi^N + (2\alpha - 1)f(1 - r) + (\alpha + 1 - \beta)r, & B_0 &= (1 - \alpha)f(1 - r) - (1 - \beta)r, \end{aligned}$$

so the integral of the first two terms of expression A.32 are given by

$$\begin{aligned}
 & \int_0^\infty e^{-\rho t} \left(\frac{\partial}{\partial q^1} S \frac{\partial}{\partial r} q^1 + \frac{\partial}{\partial q^0} S \frac{\partial}{\partial r} q^0 \right) dt = \int_0^\infty e^{-\rho t} \left(A_1 \frac{t}{r} \dot{q}_t^1 - A_0 \frac{t}{r} \dot{q}_t^0 \right) dt \\
 & = - \int_0^\infty e^{-\rho t} \frac{1}{r} (1 - \rho t) \{ S(q_t^1, q_t^0) - (B_1 - B_0) \} dt \\
 & = \frac{\rho}{r} \int_0^\infty e^{-\rho t} t S(q_t^1, q_t^0) dt - \frac{1}{r} (B_1 - B_0) \int_0^\infty e^{-\rho t} (1 - \rho t) dt \\
 & = \frac{\rho}{r} \int_0^\infty e^{-\rho t} t S(q_t^1, q_t^0) dt - r (B_1 - B_0) \left[t e^{-\rho t} \right]_0^\infty \\
 & = \frac{\rho}{r} \int_0^\infty e^{-\rho t} t S(q_t^1, q_t^0) dt \tag{A.33}
 \end{aligned}$$

$$= - \frac{\rho}{r} \frac{\partial}{\partial \rho} F(\bar{q}_0) \tag{A.34}$$

The third term in expression A.32 is given by

$$\begin{aligned}
 & (-(2\alpha - 1)f'(1 - r) - (1 - \alpha + \beta))q^1 + (1 - \alpha)(-f'(1 - r) + 1) \\
 & - (-(2\alpha - 1)f'(1 - r) + (\alpha + 1 - \beta))q^0 + (1 - \alpha)f'(1 - r) + (1 - \beta) = \\
 & \left((2\alpha - 1)f'(1 - r) + \frac{1}{r} ((2\alpha - 1)f(1 - r) + \psi^N) \right) (q^0 - q^1) + \frac{1}{r} S(q^1, q^0). \tag{A.35}
 \end{aligned}$$

Denote

$$\Lambda \equiv \left((2\alpha - 1)f'(1 - r) + \frac{1}{r} ((2\alpha - 1)f(1 - r) + \psi^N) \right).$$

Integrating expression A.35

$$\begin{aligned}
 \int_0^\infty e^{-\rho t} \frac{\partial}{\partial r} S(q_t^1, q_t^0) dt & = \int_0^\infty e^{-\rho t} \left\{ \Lambda (q_t^0 - q_t^1) + \frac{1}{r} S(q_t^1, q_t^0) \right\} dt \\
 & = \Lambda \int_0^\infty e^{-\rho t} (q_t^0 - q_t^1) dt. \tag{A.36}
 \end{aligned}$$

Combining A.34 and A.36, we obtain

$$\begin{aligned}
 \frac{\partial}{\partial r} F(\bar{q}_0) & = \int_0^\infty \frac{t}{r} \rho e^{-\rho t} S(q_t^1, q_t^0) dt + \Lambda \int_0^\infty e^{-\rho t} (q_t^0 - q_t^1) dt \\
 & = - \frac{\rho}{r} \frac{\partial}{\partial \rho} F(\bar{q}_0) + \Lambda \int_0^\infty e^{-\rho t} (q_t^0 - q_t^1) dt, \tag{A.37}
 \end{aligned}$$

■

The first term of expression A.37 captures the fact that ρ and r play opposite roles in our model: an increase in r makes dynamics faster, so it is effectively equal to moving any future point closer to the present, or equivalently, putting

more weight into the future. Hence, an increase in r can be equivalently seen as a reduction in ρ . Besides the effect that r has on the dynamics, it also has an effect on individual utilities and protests, which is captured by the second term in A.37.

Finally, using the last proposition we see that the comparative statics on r can also go both ways because they depend on the other parameters of the model.

Proposition 28 *The comparative statics on r can go both ways and depend on the other parameters of the model:*

- For small α , and sufficiently large ψ^N , it follows

$$\frac{\partial}{\partial r} \bar{q}_0 < 0.$$

- On the other hand, for large α , and sufficiently small ψ^N , it follows

$$\frac{\partial}{\partial r} \bar{q}_0 > 0.$$

Proof. Take small α and sufficiently large ψ^N such that

$$A_1 > 0 > \Lambda.$$

Using Proposition 27 we see

$$\frac{\partial}{\partial \rho} F(\bar{q}_0) > 0.$$

Combining the previous inequality with 27 and $\Lambda < 0$, we obtain

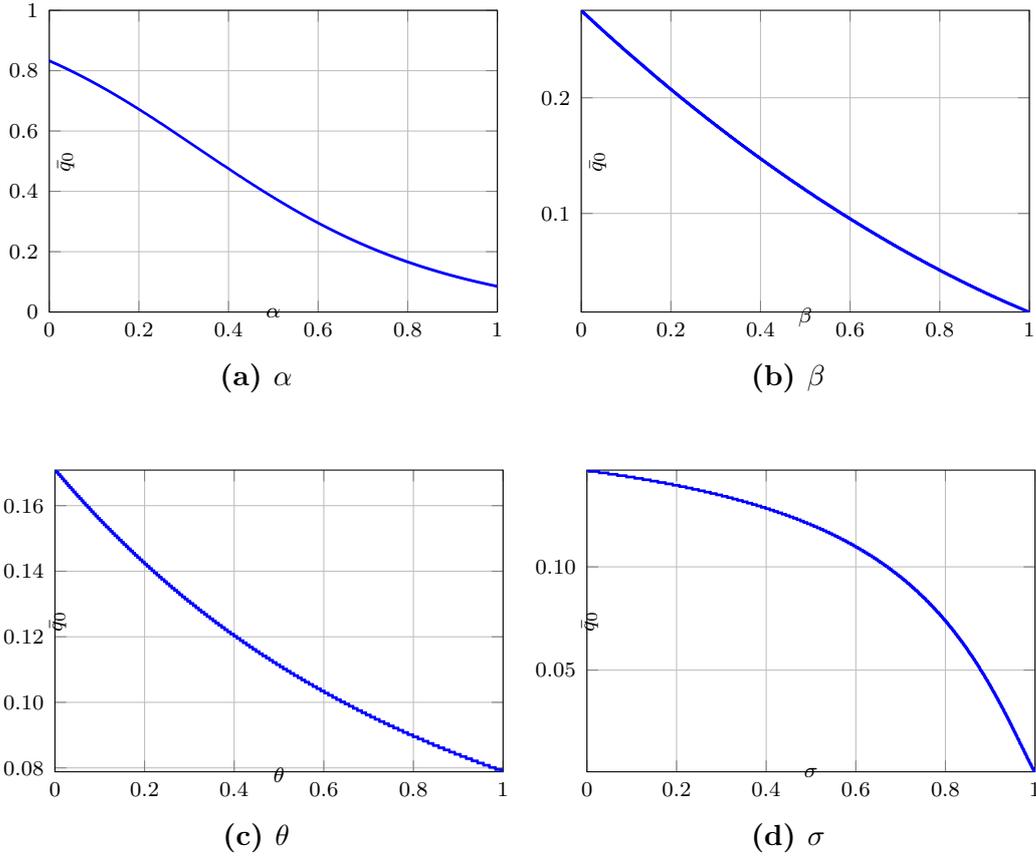
$$\frac{\partial}{\partial r} F(\bar{q}_0) = -\frac{\rho}{r} \frac{\partial}{\partial \rho} F(\bar{q}_0) + \int_0^\infty e^{-\rho t} (q_t^0 - q_t^1) dt < 0,$$

which proves the first part of the proposition. The second part is proved similarly. ■

To complement our analysis, the following graphs show numerical solutions for the threshold \bar{q}_0 . We fix the other parameters at $\alpha = 0.9, \beta = 0.5, \theta = 0.4, r = 0.3, \sigma = 0.5$, and $\rho = 0.5$ and let the corresponding parameter run over some range.

For r and ρ , we show a case with $A_1 > 0$, where we choose $\alpha = 0.9, \beta = 0.5, \theta = 0.4, \sigma = 0.5, r = 0.3, \rho = 0.5$ as baseline parameters and plot the

Baseline values $\psi^N = 0.5, \alpha = 0.9, \beta = 0.5, \theta = 0.3, \sigma = 0.5, r = 0.3, \rho = 0.5$



region in which the condition is satisfied.

A.2.4 Proofs and extra material for Section 1.5

Technical details for proof of Theorem 3

First, we prove $\frac{\partial}{\partial \tau}(W(q_t, \delta'_\tau(q_t)) - H(\tilde{q})) = 0$. Observe

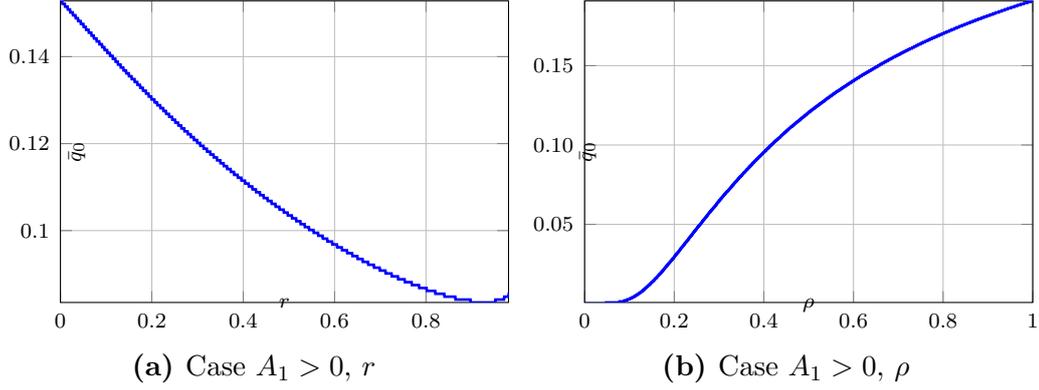
$$\frac{\partial}{\partial \tau}(W(q_t, \delta'_\tau(q_t)) - H(\tilde{q})) = \frac{\partial}{\partial \delta} W(q_t, \delta'_\tau(q_t)) \frac{\partial}{\partial \tau} \delta'_\tau(q_t).$$

It holds

$$\delta'_\tau(q) = \mathbb{1}\{t < \tau\}(\delta^S(q) + \epsilon(q)) + \mathbb{1}\{t \geq \tau\}\delta^S(q),$$

therefore

$$\frac{\partial}{\partial \tau} \delta'_\tau(q) = \Delta(\tau)(\delta^S(q) + \epsilon(q)) - \Delta(\tau)\delta^S(q) = -\Delta(\tau)\epsilon(q),$$



where $\Delta(\tau)$ is the Dirac delta function

$$\Delta(\tau) = \begin{cases} 1 & \text{if } t = \tau \\ 0 & \text{if } t \neq \tau. \end{cases}$$

Integrating

$$\int_0^\tau \frac{\partial}{\partial \tau} \delta'_\tau(q) e^{-\rho t} (W(q_t, \delta'_\tau(q_t)) - H(\tilde{q})) dt = \int_0^\tau -\frac{\partial}{\partial \delta} W(q_t, \delta'_\tau(q_t)) \Delta(\tau) \epsilon(q_t) dt = 0,$$

for all $\tau > 0$.

Second, we prove Theorem 3 still holds when $H'(\tilde{q}) = 0$ using a second order Taylor expansion for $F(\tau)$. We take derivatives with respect to τ from expression A.15 to obtain an expression for $F''(\tau)$

$$\begin{aligned} F''(\tau) &= -\rho e^{-\rho\tau} (H(q_\tau) - H(\tilde{q})) + e^{-\rho\tau} H'(q_\tau) \dot{q}_\tau \\ &\quad + \rho^2 e^{-\rho\tau} (W(q_\tau, \delta'(q_\tau)) - H(\tilde{q})) \\ &\quad - \rho e^{-\rho\tau} \left(\frac{\partial}{\partial q} W(q_\tau, \delta'(q_\tau)) + \frac{\partial}{\partial \delta} W(q_\tau, \delta'(q_\tau)) \frac{\partial}{\partial q} \delta'(q_\tau) \right) \dot{q}_\tau \\ &\quad - \rho e^{-\rho\tau} H'(q_\tau) \dot{q}_\tau + e^{-\rho\tau} H''(q_\tau) (\dot{q}_\tau)^2. \end{aligned}$$

Observe $\frac{\partial}{\partial q} W(q_\tau, \delta'(q_\tau)) + \frac{\partial}{\partial \delta} W(q_\tau, \delta'(q_\tau)) \frac{\partial}{\partial q} \delta'(q_\tau) = \frac{\partial}{\partial q} H(q_\tau) = H'(q_\tau)$, hence the previous expression simplifies to

$$\begin{aligned} F''(\tau) &= -\rho e^{-\rho\tau} (H(q_\tau) - H(\tilde{q})) + e^{-\rho\tau} (1 - 2\rho) H'(q_\tau) \dot{q}_\tau \\ &\quad + \rho^2 e^{-\rho\tau} (W(q_\tau, \delta'(q_\tau)) - H(\tilde{q})) \\ &\quad + e^{-\rho\tau} H''(q_\tau) (\dot{q}_\tau)^2. \end{aligned}$$

Evaluating at $\tau = 0$

$$F''(0) = \frac{1}{\rho} H''(\tilde{q})(\dot{\tilde{q}})^2 > 0,$$

because \tilde{q} is not local maximum of $H(q)$, with $H'(\tilde{q}) = 0$. Therefore, it must hold $H''(\tilde{q}) > 0$. Then

$$J(\tilde{q}, \delta'_\tau(\tilde{q})) - J(\tilde{q}, \delta^*(\tilde{q})) = F(\tau) > 0,$$

a contradiction.

A.2.5 Quadratic Protests

When protests enter as quadratic costs we have

$$\begin{aligned} H(q) &= \psi^N q + (\alpha q + (1 - \alpha)(1 - q))f(1 - r) \\ &\quad + r(\alpha q^2 + (1 - \alpha)(1 - q)^2) - r^2 q(1 - q)(\beta(1 - q) + (1 - \beta)q). \end{aligned}$$

The second derivative of this function is given by

$$H''(q) = 2r + 2r^2(\beta(1 - q) + (1 - \beta)q - (1 - 2q)(1 - 2\beta)).$$

Observe for all $\beta, q \in [0, 1]$

$$\beta(1 - q) + (1 - \beta)q - (1 - 2q)(1 - 2\beta) \geq -1.$$

Hence,

$$H''(q) \geq 2r - 2r^2 = 2r(1 - r) \geq 0.$$

Therefore, $H(q)$ does not have a maximum in $[0, 1]$.

Comparative statics in the quadratic case

Proposition 29 *When protests are quadratic, comparative statics for threshold \bar{q}_0 are as follows*

- It holds

$$\frac{\partial}{\partial \psi^N} \bar{q}_0 < 0, \quad \frac{\partial}{\partial \alpha} \bar{q}_0 < 0, \quad \frac{\partial}{\partial \beta} \bar{q}_0 < 0.$$

- $\alpha \geq \frac{1}{2}$ if and only if

$$\frac{\partial}{\partial \theta} \bar{q}_0 \leq 0, \quad \frac{\partial}{\partial \sigma} \bar{q}_0 \leq 0.$$

- If $\psi^N + (2\alpha - 1)f(1 - r) - r^2(1 - \alpha + \beta) > 0$

$$\frac{\partial}{\partial \rho} \bar{q}_0 >$$

- If $\psi^N + (2\alpha - 1)f(1 - r) + r^2(\alpha + 1 - \beta) < 0$, then

$$\frac{\partial}{\partial \rho} \bar{q}_0 < 0$$

Proof. It follows the same argument as in the comparative statics for the linear case. See Supplementary Appendix A.2.3. ■

A.2.6 Technical details for Section 1.6

Microfoundations of political parties' objective function

We follow the probabilistic voting model with majority voting and aggregate uncertainty proposed by Persson and Tabellini (2000) based on Lindbeck and Weibull (1987). Recall that parties A and B make simultaneous announcements δ^A and δ^B in every period, with full commitment. Voters are myopic, in the sense that they only value policies according to their utility in period t .¹² Voter j in group i votes for A if

$$U^i(\delta^A) > U^i(\delta^B) + \sigma^{ij} + \mu,$$

where σ^{ij} measures ideological idiosyncratic preference toward party B . σ^{ij} is i.i.d. and drawn from a uniform distribution $\mathcal{U}\left[\frac{-1}{2\phi^i}, \frac{1}{2\phi^i}\right]$. Note the distributions have density ϕ^i and neither group is biased on average toward one of the parties. We could think about this parameter as reflecting another policy dimension orthogonal to policy δ_t , for which political parties cannot make credible commitments but on which they implement some policy after the election in accordance with their ideology. In a sense, it is a measure of ideological

¹²Concretely, voters do not internalize the effect of their choices on the dynamics of identities.

homogeneity within the group that translates into political strength. μ captures average relative popularity of party B , drawn i.i.d. from $\mathcal{U}\left[\frac{-1}{2}, \frac{1}{2}\right]$. Note that without introducing aggregate uncertainty (given by the value of μ), the probability of winning that we define below is not continuous on the announcement, and the model collapses to a modified version of the Downsian model in which all that matters are the preferences of the swing voter. In that case, any forward-looking motive will have no bite, as any party deviating from the preferences of the swing voter loses the elections with probability 1. In that case, the only possible equilibrium is to play the optimal strategy of the static game.¹³

The probability that a randomly drawn voter of group i votes for A is given by

$$Pr(\sigma^{ij} < U^i(\delta^A) - U^i(\delta^B) - \mu) = F^i\left(U^i(\delta^A) - U^i(\delta^B) - \mu\right) = \frac{1}{2} + \phi^i[U^i(\delta^A) - U^i(\delta^B) - \mu].$$

Hence, the vote share for party A for policy announcements δ^A and δ^B for given q at time t is

$$\pi^A(\delta^A, \delta^B, q) = \frac{1}{2} + q\phi^N\left[U^N(\delta^A) - U^N(\delta^B) - \mu\right] + (1 - q)\phi^R\left[U^R(\delta^A) - U^R(\delta^B) - \mu\right].$$

We assume a majority voting electoral rule, so party A wins the election at time t if $\pi^A > \frac{1}{2}$. Because at the time announcements are made the popularity shock μ is unknown, π^A is a random variable and therefore party A wins the election with probability p^A given by

$$\begin{aligned} p^A(\delta^A, \delta^B, q_t) &= Pr\left(\pi^A > \frac{1}{2}\right) \\ &= \frac{1}{2} + \frac{q_t\phi^N[U^N(\delta^A) - U^N(\delta^B)] + (1 - q_t)\phi^R[U^R(\delta^A) - U^R(\delta^B)]}{q_t\phi^N + (1 - q_t)\phi^R} \\ &= \frac{1}{2} + \Phi(q_t)(\delta^A - \delta^B), \end{aligned}$$

where $\Phi(q) = \frac{(1 - q)\phi^R - q\phi^N}{q\phi^N + (1 - q)\phi^R}$. It follows that party B wins the elections with probability $p^B(\delta^A, \delta^B, q_t) = 1 - p^A(\delta^A, \delta^B, q_t)$.

¹³The results of this section remain if instead of introducing aggregate uncertainty and majority voting we assume that there is no aggregate uncertainty but: a) the benefits from office for each party are proportional to its vote share and; b) the policy implemented is a weighted average of the announcements. This specification yields an equivalent game and it allows to discuss how the degree of proportionality of the electoral system (i.e. how vote shares translate into power shares) affects nation-building prospects.

A.2.7 Static electoral competition

First, we consider parties that are myopic, in the sense that they do not internalize identity dynamics.¹⁴ Therefore, in each period they solve the static political-economy game, i.e., they maximize the objective function taking what the other party does as given. The Nash equilibria of the static electoral-competition game are characterized by

$$\begin{aligned}\delta^{*i}(q) &= \arg \max_{\delta \in [0,1]} \psi^N q + p^i(\delta, \delta^{*-i}) = \arg \max_{\delta \in [0,1]} \psi^N q + \frac{1}{2} + \Phi(q)(\delta - \delta^{*-i}), \\ \delta^{*-i}(q) &= \arg \max_{\delta \in [0,1]} \psi^N q + p^{-i}(\delta^{*i}, \delta) = \arg \max_{\delta \in [0,1]} \psi^N q + \frac{1}{2} + \Phi(q)(\delta - \delta^{*i}).\end{aligned}$$

It is easy to see that for given q , the symmetric Nash equilibrium is characterized by

$$\delta^{i*}(q) = \arg \max_{\delta} \Phi(q)(\delta - \delta^{*-i}) = \begin{cases} 1 & \text{if } \Phi(q) > 0 \\ [0, 1] & \text{if } \Phi(q) = 0 \\ 0 & \text{if } \Phi(q) < 0. \end{cases}$$

Because $\Phi(q)$ is strictly decreasing in q , the previous equilibrium strategy can be equivalently defined as

$$\delta^{i*}(q) = \begin{cases} 1 & \text{if } q_0 < \tilde{q}_S \\ [0, 1] & \text{if } q_0 = \tilde{q}_S \\ 0 & \text{if } q_0 > \tilde{q}_S, \end{cases}$$

where \tilde{q}_S is given by $\Phi(\tilde{q}_S) = 0$, that is $\tilde{q}_S = \frac{\phi^R}{\phi^R + \phi^N} \in [0, 1]$. Given these equilibrium policies, if $q_0 < \tilde{q}_S$, q_t decreases over time converging to $q = 0$. Alternatively, if $q_0 > \tilde{q}_S$, q_t increases over time converging to $q = 1$. When a group of voters is more concerned about policy δ , in the sense that they are more responsive to changes in the announcement (i.e. higher value of ϕ^i), they are more likely to win elections and, eventually, become the only group in society. Therefore, as in the dynamic game, the survival of regional identities is more likely when the regionalist are demographically big, when the peripheral region is sufficiently pivotal, and when citizens in the regionalist group are ideologically motivated toward identity policy δ with respect to other policy

¹⁴Equivalently, we can have political parties that only live for one period.

dimensions.

A.2.8 Parties with opposite nation-building motives

In the electoral competition game we have assumed that both parties want to promote the same national identity. However, as the recent histories of some countries in Africa and Asia show, there are several cases that are better modelled as a game between two *forward-looking* parties that are biased in opposite directions. Unfortunately, characterizing the solution to this differential game is technically intractable with the tools developed in this paper, because we cannot restrict our attention to symmetric equilibria. Solving it is a very interesting venue for future research, and it may potentially generate persistent conflict and diversity *as an equilibrium outcome*. However, we believe that also in this case it would be unlikely to obtain either cycles or heterogeneous steady-states, because the two key ingredients for long-run homogeneity under electoral competition (a strong conflict over scarce resources as well as a policy implementation that favours the majority) remain valid in the case where parties have opposite nation-building motives.

In order to sketch how the results could change with parties that represent only the interests of their own groups, we can analyze an example in which parties are shortsighted. However, note that in the shortsighted case the nation-building motive plays no role, so whether parties are biased towards increasing the size of the group with the national or the regional identity is irrelevant. Therefore, in order to have some action, we need to consider political parties that are ideologically motivated to implement some policy. For this, consider a simple modification to the current model, where party A chooses $\delta_t^A = 0$ whenever it wins elections and party B chooses $\delta_t^A = 1$.¹⁵ As in the benchmark model, we assume that there is an idiosyncratic shock and a common shock to party popularity, but the latter is now distributed as a uniform $\mathcal{U}[-1, 1]$.¹⁶ Therefore, following the steps above, the probability that party A wins the election at time t , when parties announce $\delta_t^A = 0$ and $\delta_t^B = 1$, is given by

¹⁵This example corresponds to a situation where parties are ideologically motivated and cannot commit to implement other policies once they are in office. Despite its simplicity, this assumption captures well the situation of countries such as Nigeria and Kenya, where parties are generally shortsighted, represent different ethnic groups and take turns in power to loot the country.

¹⁶This change is just to make probabilities bounded between 0 and 1.

$$p_t^A = \frac{q\phi^N}{q\phi^N + (1-q)\phi^R}$$

Hence, for ideologically motivated parties the probability of winning elections is increasing in the size of the group that it favors with its policies. Recall that dynamics are given by

$$\dot{q}_t = g(q_t) = \begin{cases} q_t(1-q_t)^2 & \text{with prob. } p_t^A \\ -q_t^2(1-q_t) & \text{with prob. } 1-p_t^A \end{cases}$$

As compared to our model of electoral competition, the policy announcements of candidates do not converge because of their extreme ideological bias. As a result, q does not always move in the same direction and the system does not necessarily reach a homogeneous steady state. However, if enough time passes, we should expect q eventually moving in the same direction. The reason is that the biggest group has a higher probability of winning elections and, as a result, get its desired policy. This increases the size of this group through the cultural evolution mechanism, which in turn makes them more likely to win elections again. Therefore, even with shortsighted and ideologically motivated candidates, homogeneous populations are the most likely long run outcome, because majority groups tend to become larger over time.

A.2.9 Voters in the central region

In this subsection we show that introducing voters in the central region does not qualitatively change the results of the electoral-competition game. The reason is that including these voters only changes the function $\Phi(q)$. Therefore, the key properties of the objective function of the central government remain similar and the key features needed for the proof go through. The main qualitative difference comes from the fact that, for some regions of parameters, some trivial cases might arise in which $\Phi(q)$ is lower than zero for all q . We illustrate this last point by means of an example.

Assume the central government is democratically elected each period by people of the central and peripheral regions. The country as a whole has a population of size 1, out of which a fraction $\lambda \in (0, 1)$ lives in the peripheral region and a fraction $1 - \lambda$ lives in the central region. Within the peripheral region, a fraction q belong to group N and a fraction $1 - q$ belong to group R . Utilities

are given by¹⁷

$$\begin{aligned} U^N(\delta) &= g^N = 1 - \delta \\ U^R(\delta) &= g^R = \delta \\ U^C(\delta) &= g^N = 1 - \delta. \end{aligned}$$

Here, we have made the simple but natural assumption that voters in the central region have the same preferences as nationalist individuals in the peripheral region. This aims to capture the idea that citizens in the central region are socialized to the national identity and enjoy the nationalist public good in the same way as nationalist individuals in the peripheral region. As before, voter j in group i votes for A if

$$U^i(\delta^A) > U^i(\delta^B) + \sigma^{ij} + \mu.$$

Assuming majority voting as before

$$p^A(\delta^A, \delta^B, q) = \frac{1}{2} + \Phi(q)(\delta^A - \delta^B),$$

where $\Phi(q)$ is now given by

$$\Phi(q) = \frac{\phi^R \lambda (1 - q) - \phi^N \lambda q - \phi^C (1 - \lambda)}{\phi^N \lambda q + \phi^R \lambda (1 - q) + \phi^C (1 - \lambda)}.$$

When $\phi^R \lambda - \phi^C (1 - \lambda) < 0$, we have $\Phi(q) < 0, \forall q \in [0, 1]$. This is the main difference with respect to the previous case, where for any value of the parameters it was guaranteed that $\Phi(q)$ always took positive and negative value in q between 0 and 1. Now, if $\phi^R \lambda - \phi^C (1 - \lambda) < 0$, the only equilibrium is given by $\delta = 0$ for any initial q_0 . This corresponds to the case in which the regionalist group is not sufficiently pivotal in national elections, which happens when the region is sufficiently small (low λ) or when citizens in the central region are relatively more ideologically concerned with policy δ than the regionalists (ϕ^C relatively large with respect to ϕ^R). In the non-trivial case when $\phi^R \lambda - \phi^C (1 - \lambda) > 0$, we are back to a similar case where voters in the central region are not introduced. Therefore, the equilibrium of the dynamic game is analogous, with the minor difference that the thresholds \tilde{q}_S and \tilde{q}_D are additionally affected by the parameters λ and ϕ^C .

One could think of more realistic and detailed specifications that would yield

¹⁷To simplify on parameters, we assume that the total budget of the government is of size 1.

more interesting comparative statics with respect to the two thresholds, without changing the method of the proof for the results in Section 1.6. For instance, we could have specified that citizens in both regions experiment disutility from protests. This could lead to a case where some citizens in the central region might vote for a policy that favors regionalist individuals, because their desire to reduce conflict might offset their nationalist sentiment. In this case, the persistence of regional identities in democracies would be a function of the complex interaction between the ideological concerns about identity policies of the three groups (ϕ^i), the size/pivotality of the peripheral region (λ) and parameters capturing the impact of protests.

Proof of Proposition 7

Proposition 7 *The threshold \tilde{q}_D is decreasing in ψ^N*

$$\frac{\partial}{\partial \psi^N} \tilde{q}_D \leq 0,$$

with limiting cases

$$\lim_{\psi^N \rightarrow 0} \tilde{q}_D = \tilde{q}_S, \quad \lim_{\psi^N \rightarrow \infty} \tilde{q}_D = 0.$$

On the contrary, \tilde{q}_D is increasing in ρ :

$$\frac{\partial}{\partial \rho} \tilde{q}_D \geq 0.$$

Proof. We first show the comparative statics on the parameter ψ^N . Simply recall the derivative of the recovered value function when $\delta^*(q) = 0$

$$V_q(q) = \psi^N \left[\frac{2}{\rho} + \frac{1}{q(1-q)^2 m(q)} \int_q^1 m(x) dx \right]$$

Observe that $V_q(q)$ is strictly increasing in ψ^N , hence if we have $\psi^N < \psi^{N'}$ and for some $q < \tilde{q}_S$ we have that

$$V_q(q; \psi^N) > 2 \frac{\Phi(q)}{q(1-q)},$$

then it also follows that

$$V_q(q; \psi^{N'}) > V_q(q; \psi^N) > 2 \frac{\Phi(q)}{q(1-q)}.$$

This means that if $q < \tilde{q}_D(\psi^N)$, then $q < \tilde{q}_D(\psi^{N'})$ too, and hence

$$\tilde{q}_D(\psi^{N'}) \leq \tilde{q}_D(\psi^N)$$

Moreover, it is easy to see that

$$\lim_{\psi^N \rightarrow 0} \tilde{q}_D(\psi^N) = \tilde{q}_S, \quad \lim_{\psi^N \rightarrow \infty} \tilde{q}_D(\psi^N) = 0.$$

For the discount factor ρ , the comparative statics are proven using a similar argument. ■

A.2.10 Endogenous tax rate

We modify the baseline model such that the government is able to choose the tax rate $\{r_t\}_{t \geq 0}$ as well as the relative provision of each type of public good $\{\delta_t\}_{t \geq 0}$. The resulting government's problem has two control variables and is given by

$$\begin{aligned} \max_{r_t, \delta_t, \in [0,1], \forall t \geq 0} \int_0^\infty e^{-\rho t} W(q_t, \delta_t, r_t) dt \\ \text{s.t. } \quad \dot{q}_t = r_t q_t (1 - q_t) (1 - \delta_t - q_t) \\ q(0) = q_0, q_t \in [0, 1]. \end{aligned} \quad (\text{A.38})$$

with corresponding HJB equation given by

$$\rho V(q) = \max_{r, \delta} W(q, \delta, r) + g(q, r, \delta) V'(q), \quad (\text{A.39})$$

where

$$\begin{aligned} W(q, \delta, r) &= \psi^N q_t + \alpha q (f(1 - r) + (1 - \delta)r) + (1 - \alpha)(1 - q)(f(1 - r) + \delta r) \\ &\quad - r(\beta q \delta + (1 - \beta)(1 - q)(1 - \delta)) \\ g(q, \delta, r) &= r q (1 - q)(1 - q - \delta). \end{aligned}$$

The following proposition holds:

Proposition 30 *Assume utility from private consumption is $f(x) = \theta \frac{x^{1-\sigma}}{1-\sigma}$ with $\theta, \sigma \in (0, 1)$. Then, open neighborhoods of $q = 0$ and $q = 1$ in $[0, 1]$ exist, say, $\mathcal{O}(0)$ and $\mathcal{O}(1)$, such that*

$$r^*(q) > 0, \text{ with } \delta^*(q) = 1 \quad \forall q \in \mathcal{O}(0), \quad \text{and} \quad r^*(q) > 0, \text{ with } \delta^*(q) = 0 \quad \forall q \in \mathcal{O}(1).$$

Proof. From A.39, the optimal tax-rate for $q = 0$ and $q = 1$ is given by

$$r^*(0) = r^*(1) = 1 - \theta^{\frac{1}{\sigma}} \in (0, 1),$$

with corresponding value function

$$\rho V(0) = (1 - \alpha) \left(1 + \theta^{\frac{1}{\sigma}} \frac{\sigma}{1 - \sigma}\right) \quad \rho V(1) = \psi^N + \alpha \left(1 + \theta^{\frac{1}{\sigma}} \frac{\sigma}{1 - \sigma}\right).$$

First, we prove that $r^*(q)$ is continuous at $q = 1$ and at $q = 0$, by contradiction. Assume not, so $\lim_{q \rightarrow 1} r^*(q) = c \neq 1 - \theta^{\frac{1}{\sigma}}$. From Theorem 1 we know $\delta^*(q)$ is continuous at $q = 0$ and at $q = 1$. Then, by continuity of $V(q)$ it must hold

$$\lim_{q \rightarrow 1} \rho V(q) = \psi^N + \alpha(f(1 - c) + c) = \psi^N + \alpha \left(1 + \theta^{\frac{1}{\sigma}} \frac{\sigma}{1 - \sigma}\right) = \rho V(1),$$

which implies

$$f(1 - c) + c = f(\theta^{\frac{1}{\sigma}}) + 1 - \theta^{\frac{1}{\sigma}} = \max_x f(1 - x) + x.$$

Observe the function $f(1 - x) + x$ is strictly concave, so the only solution of the previous equation is precisely $c = 1 - \theta^{\frac{1}{\sigma}}$, and therefore

$$\lim_{q \rightarrow 1} r^*(q) = 1 - \theta^{\frac{1}{\sigma}} = r^*(1),$$

which proves $r^*(q)$ is continuous at $q = 1$. Similarly for $q = 0$. Because $r^*(1) = r^*(0) > 0$, by continuity of $r^*(q)$ open neighborhoods in $[0, 1]$ of $q = 0$ and $q = 1$ exist such that $r^*(q) > 0$ for all q in those neighborhoods.

For the second part of the proposition, we use continuity of $\delta^*(q)$ at $q = 0$ and at $q = 1$, which follows from Theorem 1. By continuity of $r^*(q)$ at $q = 0$ and $q = 1$, we can find open neighborhoods of $q = 1$ and $q = 0$, $\mathcal{O}(0)$ and $\mathcal{O}(1)$ respectively, such that $r^*(q)$ and $\delta^*(q)$ are continuous inside them. Also, from Theorem 1, either $\delta^*(q) = 0$ or $\delta^*(q) = 1$, with $\delta^*(0) = 1$ and $\delta^*(1) = 0$. By continuity of $\delta^*(q)$ in $\mathcal{O}(0)$ and $\mathcal{O}(1)$, it follows

$$r^*(q) > 0, \delta^*(q) = 1 \quad \forall q \in \mathcal{O}(0), \quad \text{and} \quad r^*(q) > 0, \delta^*(q) = 0 \quad \forall q \in \mathcal{O}(1).$$

■

The previous result implies that when the population is largely homogeneous, it is better for the government to collect taxes, provide public goods, and homogenize toward the prevailing identity, because at those states the participation rate in protests of the minority group is small and it is optimal to

pursue full homogenization.

Toward a general solution

Unfortunately, finding a closed-form solution of the optimal tax rate r is analytically intractable given the cubic law of motion of the state variable q . However, in this section we outline the steps toward a full solution of problem A.38.

First, we show that the solution to problem A.38 is equivalent to a sequential maximization problem. From Theorem 1, we know that the solution $\delta^*(r, q)$ for any r and q is given by

$$\delta^*(r, q) = \arg \max_{\delta} W(q, \delta, r) + g(q, r, \delta)V'(q) = \begin{cases} 1 & \text{if } q \leq \bar{q}_0(r) \\ 0 & \text{if } q \geq \bar{q}_0(r). \end{cases}$$

for any given r , and q . That is, for any r , including the optimal tax-rate $r^*(q)$, we know that $\delta^*(r^*(q), q)$ can only take two values, i.e. $\delta^*(r^*(q), q) \in \{0, 1\}$ for all $q \in [0, 1]$. The previous result greatly simplifies problem A.38, to

$$\rho V(q) = \max_{\delta \in \{0,1\}} \left\{ \max_{r \in [0,1]} W(q, \delta, r) + g(q, \delta, r)V'(q), \max_{r \in [0,1]} W(q, 1, r) + g(q, 1, r)V'(q) \right\} \quad (\text{A.40})$$

Next, to find interior solutions $r^*(q) \in (0, 1)$ we could solve each sub-problem in problem A.40 by solving the corresponding ODE obtained from the envelope and first order conditions of the HJB equation. However, there are no analytic solutions to those ODEs. To illustrate this point, we can look at the solution for low values of q , for which we know $\delta^*(q) = 1$, and hence the corresponding ODE for $r^*(q)$ is given by

$$r_q^1 = \frac{1}{rf''(1-r)(\alpha q + (1-\alpha)(1-q))} \left\{ \psi^N + (2\alpha - 1)(f'(1-r)r + f(1-r)) \right. \\ \left. - \frac{\rho}{q^2(1-q)} \left((\alpha q + (1-\alpha)(1-q))f'(1-r) \right. \right. \\ \left. \left. - (1-\alpha)(1-q) + \beta q \right) \right\}.$$

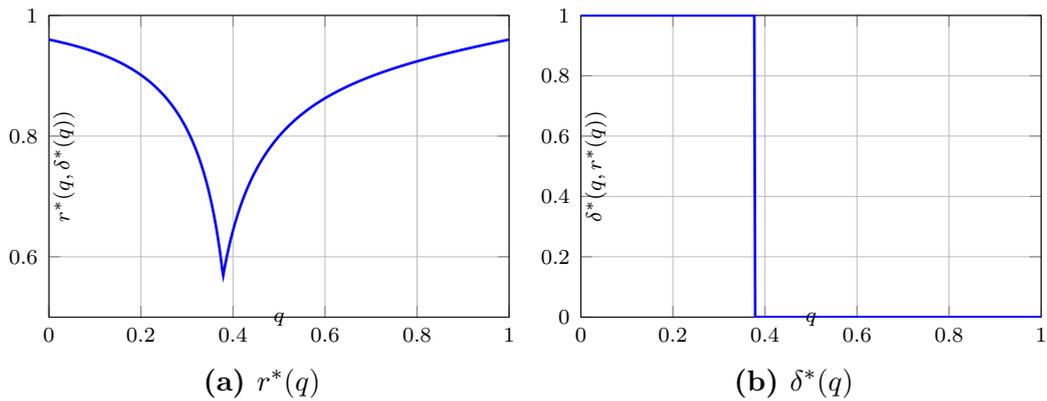
However, the previous ODE does not have a closed-form analytic solution, even after choosing specific values of the parameters α , β , and ρ and σ . Therefore,

obtaining a full complete characterization of $r^*(q)$ is analytically intractable. Similarly, the corresponding ODE for large values of q is given by

$$r_q^0 = \frac{1}{r f''(1-r)(\alpha q + (1-\alpha)(1-q))} \left\{ \begin{aligned} &\psi^N + (2\alpha - 1)(f'(1-r)r + f(1-r)) \\ &+ \frac{\rho}{q(1-q)^2} \left(-(\alpha q + (1-\alpha)(1-q))f'(1-r) \right. \\ &\left. + (\alpha q - (1-\beta)(1-q)) \right) \end{aligned} \right\}$$

Furthermore, observe how general results about the monotonicity of $r^*(q)$ are difficult to obtain because the sign of the previous ODEs depend on the other parameters of the model. Hence different parameter combinations will lead to different results.

Parameter values: $\psi^N = 0.5, \alpha = 0.5, \beta = 0.5, \theta = 0.2, \sigma = 0.5, \rho = 0.5$



For illustrative purposes, we numerically solve problem A.40 using the numerical methods proposed in Achdou et al. (2017). Figures A.7a and A.7b illustrate the solution for $r^*(q)$ and $\delta^*(q)$ for some parameters. We can see that whenever $r^*(q) > 0$ for all $q \in [0, 1]$, the optimal solution for $\delta^*(q)$ resembles the bang-bang nature of the baseline model; that is, a threshold \bar{q}_0 exists such that $\delta_t = 1 \forall t$, with q_t converging to $q = 0$ whenever $q \leq \bar{q}_0$ and vice versa. All numerical examples display similar qualitative results. Importantly, even when r is chosen optimally (and conditional on being strictly positive), convergence to an extreme steady state still occurs, showing softer budget constraints do not eliminate the overall conflict.¹⁸

¹⁸For some parameter choices, we found cases in which $r^*(q) = 0$ for intermediate values of q . This finding corresponds to cases in which the government is sufficiently welfarist and individuals' marginal utility of private consumption is relatively big (sufficiently high σ and θ).

Moreover, the numerical solution suggests the optimal tax rate $r^*(q)$ is higher for more homogeneous populations and reaches a minimum at the indifference threshold \bar{q}_0 , as a result of a *static* trade-off present in the choice of r . On the one hand, an increase in r reduces the private consumption of both groups. On the other hand, it increases the resources available to provide one of the two public goods. For intermediate values of q , the negative effect dominates because all citizens are affected by the tax collection, but only one group benefits from public-good provision. However, as the government comes closer to the homogeneous states, the positive effects dominate because the benefits from the public-good provision are larger. Moreover, we can see r increases sharply at early stages and at diminishing rate afterwards. This behavior results from the *dynamic* effect of changing r and directly affects the law of motion: By increasing r , the government can move faster in any direction. Therefore, for intermediate values of q , the government wants to change r sharply in order to rapidly reduce the size of the group that pushes welfare down.

Appendix B

B.1 Proof of Lemma 2

At the production stage, the problem of the social planner is to choose $q_S(\theta)$ and $q_B(\theta)$ to maximize total welfare W , taking capacity $K > 0$ and the demand distribution $G(\theta)$ as given. Therefore, we look for the solution to the following problem:

$$\begin{aligned} \max_{q_S(\theta), q_B(\theta), \forall \theta} W(q_S(\theta), q_B(\theta)) &= \int_{\underline{\theta}}^{\bar{\theta}} \left[v\theta - \frac{(\theta - q_S(\theta) + q_B(\theta))^2}{2} \right] g(\theta) d\theta \\ \text{s.t. } h_1(q_S(\theta), q_B(\theta)) &= \int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} q_S(\theta) g(\theta) d\theta \geq 0 \\ h_2(q_B(\theta)) &= K - \int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta \geq 0 \\ h_3(q_S(\theta)) &= q_S(\theta) \geq 0, \forall \theta \\ h_4(q_B(\theta)) &= q_B(\theta) \geq 0, \forall \theta \end{aligned}$$

with $K > 0$. We can define the constraint set of the problem as:

$$C := \{q_S(\theta), q_B(\theta) \in X : h_j(q_S(\theta); q_B(\theta)) \geq 0, j = \{1, 2, 3, 4\}\}$$

The set $X = (0, +\infty)^2$ is open and convex because it is Cartesian product of open intervals which are open, convex sets. Note that the objective function $W(\cdot)$ and the constraints are continuously differentiable functions. Moreover, C is closed, bounded and compact, so the solution set to the problem is non-empty. Moreover, $W(\cdot)$ is strictly concave in $q_B(\theta)$ and $q_S(\theta)$. The constraints are (weakly) concave, so the solution to the problem is unique.

The Lagrangian of the problem is:

$$\begin{aligned} \mathbb{L}(q_B(\theta), q_S(\theta), \eta_S(\theta), \eta_B(\theta), \lambda, \mu) = & \\ & \int_{\underline{\theta}}^{\bar{\theta}} \left[v\theta - \frac{(\theta - q_S(\theta) + q_B(\theta))^2}{2} \right] g(\theta) d\theta + \int_{\underline{\theta}}^{\bar{\theta}} \eta_S(\theta) q_S(\theta) g(\theta) d\theta \\ & + \int_{\underline{\theta}}^{\bar{\theta}} \eta_B(\theta) q_B(\theta) g(\theta) d\theta + \lambda \left[\int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} q_S(\theta) g(\theta) d\theta \right] + \mu \left[K - \int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta \right] \end{aligned}$$

where $\lambda, \mu, \eta_S(\theta)$ and $\eta_B(\theta)$ are the multipliers associated with their respective constraints $h_1(\cdot), h_2(\cdot), h_3(\cdot), h_4(\cdot) \geq 0$. The Karush-Kuhn-Tucker (KKT) conditions are:

$$\theta - q_S(\theta) + q_B(\theta) - \lambda + \eta_S(\theta) = 0, \forall \theta \quad (\text{B.1})$$

$$\theta - q_S(\theta) + q_B(\theta) - \lambda + \mu - \eta_B(\theta) = 0, \forall \theta \quad (\text{B.2})$$

$$\eta_i(\theta) \geq 0, \forall \theta, i = \{S, B\}$$

$$q_i(\theta) \geq 0, \forall \theta, i = \{S, B\}$$

$$\eta_i(\theta) q_i(\theta) = 0, \forall \theta, i = \{S, B\}$$

$$\int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} q_S(\theta) g(\theta) d\theta \geq 0 \quad (\text{B.3})$$

$$\lambda \geq 0 \quad (\text{B.4})$$

$$\lambda \left[\int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} q_S(\theta) g(\theta) d\theta = 0 \right] = 0 \quad (\text{B.5})$$

$$\mu \geq 0 \quad (\text{B.6})$$

$$K - \int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta \geq 0 \quad (\text{B.7})$$

$$\mu \left[K - \int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta \right] = 0 \quad (\text{B.8})$$

These conditions are necessary and sufficient, due the concavity of the objective function and the constraints. Without loss of generality, we can focus attention on cases in which for any $\theta \in [\underline{\theta}, \bar{\theta}]$, $q_B(\theta) > 0 \rightarrow q_S(\theta) = 0$ & $q_S(\theta) > 0 \rightarrow q_B(\theta) = 0$. We conjecture that there exists $\theta_1 \in [\underline{\theta}, \bar{\theta}]$ and $\theta_2 \in [\underline{\theta}, \bar{\theta}]$, with $\theta_1 \leq \theta_2$, such that:

$$\begin{cases} q_B(\theta) > 0 & \text{if } \theta < \theta_1 \\ q_B(\theta) = 0 & \text{if } \theta \geq \theta_1 \end{cases} \quad \text{and} \quad \begin{cases} q_S(\theta) = 0 & \text{if } \theta \leq \theta_2 \\ q_S(\theta) > 0 & \text{if } \theta > \theta_2 \end{cases} .$$

We proceed by finding the expressions for $q_B(\theta)$, $q_S(\theta)$, θ_1 and θ_2 implied by this conjecture that satisfy all the KKT conditions. Note that $\lambda > 0$ must be satisfied in every possible solution of this problem, as if the associated constraint holds with inequality, one can always increase the value of the program by increasing $q_S(\theta)$ or reducing $q_B(\theta)$.

From condition (B.1):

$$q_S(\theta) = \theta - \lambda, \forall \theta > \theta_2$$

and from condition (B.2):

$$q_B(\theta) = \lambda - \mu - \theta, \forall \theta < \theta_1$$

By continuity:

$$\begin{aligned} q_S(\theta_2) = 0 &\Rightarrow \theta_2 = \lambda \Rightarrow q_S^{FB}(\theta) = \theta - \theta_2, \forall \theta > \theta_2 \\ q_B(\theta_1) = 0 &\Rightarrow \theta_1 = \lambda - \mu \Rightarrow q_B^{FB}(\theta) = \theta_1 - \theta, \forall \theta < \theta_1 \end{aligned}$$

From condition B.3:

$$\int_{\underline{\theta}}^{\theta_1} (\theta_1 - \theta)g(\theta)d\theta = \int_{\theta_2}^{\bar{\theta}} (\theta - \theta_2)g(\theta)d\theta. \quad (\text{B.9})$$

We have two possible cases depending on the value of the exogenous parameter K . When K is binding, $\mu > 0$ and $\theta_2 - \theta_1 = \mu > 0$. Define $x = \theta_2 - \theta_1$. By symmetry of $G(\theta)$, equation (B.9) implies that θ_2 and θ_1 must be symmetric around the mean, i.e.,

$$\theta_1 = \mathbb{E}(\theta) - \frac{x}{2} \Rightarrow \theta_1^{FB} = \mathbb{E}(\theta) - \frac{\mu^{FB}}{2} \quad (\text{B.10})$$

$$\theta_2 = \mathbb{E}(\theta) + \frac{x}{2} \Rightarrow \theta_2^{FB} = \mathbb{E}(\theta) + \frac{\mu^{FB}}{2} \quad (\text{B.11})$$

with μ^{FB} implicitly given by:

$$\int_{\underline{\theta}}^{\theta_1(\mu^{FB})} (\theta_1(\mu^{FB}) - \theta)g(\theta)d\theta = \int_{\theta_2(\mu^{FB})}^{\bar{\theta}} (\theta - \theta_2(\mu^{FB}))g(\theta)d\theta = K. \quad (\text{B.12})$$

When K is not binding, so that $\mu = 0$, from equations (B.10) and (B.11) it is straightforward to establish that $\theta_1^{FB} = \theta_2^{FB} = \mathbb{E}(\theta)$. Therefore, the unique

solution in this case is:

$$\left\{ \begin{array}{ll} q_B(\theta) = \mathbb{E}(\theta) - \theta & \text{if } \theta < \mathbb{E}(\theta) \\ q_B(\theta) = 0 & \text{if } \theta \geq \mathbb{E}(\theta) \end{array} \right. \text{ and } \left\{ \begin{array}{ll} q_S(\theta) = 0 & \text{if } \theta \leq \mathbb{E}(\theta) \\ q_S(\theta) = \theta - \mathbb{E}(\theta) & \text{if } \theta > \mathbb{E}(\theta) \end{array} \right.$$

B.2 Proof of Proposition 8

Now we turn to the problem of choosing optimal K at the investment stage. The problem of the social planner at the investment stage is to maximize total welfare (which is a function of K alone) given the optimal operation of storage at the production stage.

Let $V(K)$ be the value function after substituting the optimal solutions $q_S^*(\theta, K)$ and $q_B^*(\theta, K)$ at the production stage. Thus, the problem of the social planner at the investment stage is

$$\max_K W(q_S^*(\theta, K), q_B^*(\theta, K), K) - C(K) = V(K) - C(K)$$

By the envelope theorem, we have that:

$$\frac{dV(K)}{dK} = \frac{\partial \mathbb{L}(q_B(\theta), q_S(\theta), \eta_S(\theta), \eta_B(\theta), \lambda, \mu)}{\partial K} = \mu^{FB}.$$

Therefore, the unique interior solution K^{FB} is given by:

$$\frac{\partial W}{\partial K} = 0 \Leftrightarrow \mu^{FB} - C'(K) = \theta_2(K^{FB}) - \theta_1(K^{FB}) - C'(K^{FB}) = 0 \quad (\text{B.13})$$

with

$$\int_{\underline{\theta}}^{\theta_1(\mu(K))} (\theta_1(\mu(K)) - \theta)g(\theta)d\theta = \int_{\theta_2(K^{FB})}^{\bar{\theta}} (\theta - \theta_2(K^{FB}))g(\theta)d\theta = K^{FB}.$$

B.3 Proof of Lemma 3

The problem of the competitive fringe is:

$$\max_{q(\theta)} \left\{ p(\theta)q(\theta) - \frac{q^2(\theta)}{2(1-\alpha)} \right\}.$$

The first order condition, which is both necessary and sufficient, is:

$$p(\theta) - \frac{q^*(\theta)}{1-\alpha} = 0, \Leftrightarrow q_F^*(\theta) = (1-\alpha)p(\theta), \forall \theta$$

The dominant producer chooses its output in order to maximize its profits,

$$\max_{q_D(\theta)} \pi_D(\theta) = \frac{\theta - q_S(\theta) + q_B(\theta) - q_D(\theta)}{1 - \alpha} q_D(\theta) - \frac{[q_D(\theta)]^2}{2\alpha}$$

Hence, the first order condition of problem is:

$$\frac{\partial \pi_D(\theta)}{\partial q_D(\theta)} = 0 \Leftrightarrow \frac{\theta - q_S(\theta) + q_B(\theta) - 2q_D(\theta)}{1 - \alpha} - \frac{q_D(\theta)}{\alpha} = 0 \Leftrightarrow q_D(\theta) = \frac{\alpha}{1 + \alpha} (\theta - q_S(\theta) + q_B(\theta)), \forall \theta \quad (\text{B.14})$$

with second order condition satisfied. Note that the above implies

$$q_F(\theta) = \frac{1}{1 + \alpha} (\theta - q_S(\theta) + q_B(\theta)).$$

The equilibrium price is

$$p(\theta) = \frac{1}{1 - \alpha^2} (\theta - q_S(\theta) + q_B(\theta)).$$

B.4 Proof of Lemma 4

At the production stage, the problem of the social planner is to solve problem

$$\max_{q_B(\theta), q_S(\theta)} W = \int_{\underline{\theta}}^{\bar{\theta}} v \theta g(\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} \left[\frac{[q_D(\theta)]^2}{2\alpha} + \frac{[\theta - q_S(\theta) + q_B(\theta) - q_D(\theta)]^2}{2(1 - \alpha)} \right] g(\theta) d\theta,$$

subject to constraints (2.2) and (2.1). The structure of the functional optimization problem is identical to the one in the Proof of Proposition 2 in Appendix B.1, with concavity and compactness assumptions satisfied, so a unique solution to the problem exists. Without loss of generality, we can focus attention on cases in which for any $\theta \in [\underline{\theta}, \bar{\theta}]$, $q_B(\theta) > 0 \rightarrow q_S(\theta) = 0$ & $q_S(\theta) > 0 \rightarrow q_B(\theta) = 0$. The KKT conditions are:

$$\frac{\theta - q_S(\theta) + q_B(\theta) - q_D(\theta)}{1 - \alpha} - \lambda = 0, \forall \theta \geq \theta_2 \quad (\text{B.15})$$

$$\frac{\theta - q_S(\theta) + q_B(\theta) - q_D(\theta)}{1 - \alpha} - \lambda < 0, \forall \theta < \theta_2$$

$$\frac{\theta - q_S(\theta) + q_B(\theta) - q_D(\theta)}{1 - \alpha} - \lambda + \mu = 0, \forall \theta \leq \theta_1 \quad (\text{B.16})$$

$$\frac{\theta - q_S(\theta) + q_B(\theta) - q_D(\theta)}{1 - \alpha} - \lambda + \mu > 0, \forall \theta > \theta_1$$

$$\int_{\underline{\theta}}^{\theta_1} q_B(\theta) g(\theta) d\theta = \int_{\theta_2}^{\bar{\theta}} q_S(\theta) g(\theta) d\theta \quad (\text{B.17})$$

with $\theta_1 \leq \theta_2$ and the complementary slackness conditions identical to equations (B.4)-(B.8) of the problem in Appendix B.1. Note that condition (B.17) already incorporates the fact that we must have $\lambda > 0$ in any optimal solution to the problem. Note that these conditions are necessary and sufficient, due to the concavity of both the objective function and the constraints.

From conditions (B.15) and (B.16), and using the best response of the dominant firm, equation (B.14):

$$(\theta - q_S(\theta) + q_B(\theta)) \frac{1}{1 - \alpha^2} = \lambda, \forall \theta > \theta_2$$

and from condition (B.16):

$$(\theta - q_S(\theta) + q_B(\theta)) \frac{1}{1 - \alpha^2} = \lambda - \mu, \forall \theta < \theta_1$$

By continuity:

$$\begin{aligned} q_S(\theta_2) = 0 &\Rightarrow \theta_2 = \lambda(1 - \alpha^2) \Rightarrow q_S^{SB}(\theta) = \theta - \theta_2, \forall \theta > \theta_2 \\ q_B(\theta_1) = 0 &\Rightarrow \theta_1 = (\lambda - \mu)(1 - \alpha^2) \Rightarrow q_B^{SB}(\theta) = \theta_1 - \theta, \forall \theta < \theta_1 \end{aligned}$$

From condition (B.17):

$$\int_{\underline{\theta}}^{\theta_1} (\theta_1 - \theta)g(\theta)d\theta = \int_{\theta_2}^{\bar{\theta}} (\theta - \theta_2)g(\theta)d\theta. \quad (\text{B.18})$$

We have two possible cases depending on the value of the exogenous parameter K . When K is binding, $\mu > 0$ and $\theta_2 - \theta_1 = \mu(1 - \alpha^2) > 0$. Define $x = \theta_2 - \theta_1$. By symmetry of $G(\theta)$, equation (B.27) implies that θ_2 and θ_1 must be symmetric around the mean, i.e.,

$$\theta_1 = \mathbb{E}(\theta) - \frac{x}{2} \Rightarrow \theta_1^{SB} = \mathbb{E}(\theta) - \frac{\mu^{SB}}{2}(1 - \alpha^2) \quad (\text{B.19})$$

$$\theta_2 = \mathbb{E}(\theta) + \frac{x}{2} \Rightarrow \theta_2^{SB} = \mathbb{E}(\theta) + \frac{\mu^{SB}}{2}(1 - \alpha^2) \quad (\text{B.20})$$

with μ^{SB} implicitly given by:

$$\int_{\underline{\theta}}^{\theta_1(\mu^{SB})} (\theta_1(\mu^{SB}) - \theta)g(\theta)d\theta = \int_{\theta_2(\mu^{SB})}^{\bar{\theta}} (\theta - \theta_2(\mu))g(\theta)d\theta = K. \quad (\text{B.21})$$

Note that when K is not binding, so that $\mu = 0$, from equations (B.19) and

(B.20) it is straightforward to establish that $\theta_1^{SB} = \theta_2^{SB} = \mathbb{E}(\theta)$.

B.5 Proof of Proposition 9

The problem of the constrained social planner is to choose K to maximize total welfare, conditional on optimal behavior of all agents at the production stage. Thus, the problem is: Second best:

$$\begin{aligned} \max_{q_B(\theta), q_S(\theta)} (K, q_S^*(\theta, K), q_B^*(\theta, K)) &= vE[\theta] - \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{2\alpha} (q_D(\theta))^2 g(\theta) d\theta \\ &\quad + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{2(1-\alpha)} (\theta - q_S(\theta) + q_B(\theta) - q_D(\theta))^2 g(\theta) d\theta \end{aligned}$$

By the envelope theorem,

$$\frac{dW}{dK} = \frac{\partial W}{\partial K} + \int_{\underline{\theta}}^{\bar{\theta}} \frac{\partial W}{\partial q_D(\theta)} \frac{\partial q_D^*(\theta)}{\partial K} g(\theta) d\theta.$$

The first term is a direct effect and it equals μ^{SB} . The second term is a strategic effect which results from the impact of K on the dominant firm's output decision. Focusing on it,

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} \frac{\partial W}{\partial q_D(\theta)} \frac{\partial q_D^*(\theta)}{\partial K} g(\theta) d\theta &= - \int_{\underline{\theta}}^{\bar{\theta}} \left[\frac{\partial q_D^*(\theta)}{\partial K} \frac{q_D^*(\theta) - \alpha(\theta - q_S^*(\theta) + q_B^*(\theta))}{\alpha(1-\alpha)} \right] g(\theta) d\theta \\ &= \frac{\alpha}{(1-\alpha^2)} \int_{\underline{\theta}}^{\bar{\theta}} \left[\frac{\partial q_D^*(\theta)}{\partial K} (\theta - q_S^*(\theta) + q_B^*(\theta)) \right] g(\theta) d\theta, \end{aligned}$$

where the second line follows from using the expression for $q_D^*(\theta)$.

Since

$$q_B(\theta) = \max\{\theta_1(\mu) - \theta, 0\} \quad \text{and} \quad q_S(\theta) = \max\{\theta - \theta_2(\mu), 0\},$$

we can write,

$$\int_{\underline{\theta}}^{\bar{\theta}} \frac{\partial W}{\partial q_D(\theta)} \frac{\partial q_D^*(\theta)}{\partial K} g(\theta) d\theta = \frac{\alpha}{(1-\alpha^2)} \left(\int_{\underline{\theta}}^{\theta_1} \frac{\partial q_D^*(\theta)}{\partial K} \theta_1 g(\theta) d\theta + \int_{\theta_2}^{\bar{\theta}} \frac{\partial q_D^*(\theta)}{\partial K} \theta_2 g(\theta) d\theta \right).$$

For $\theta \in (\underline{\theta}, \theta_1)$,

$$q_D^*(\theta) = \frac{\alpha}{1+\alpha} \theta_1 \Rightarrow \frac{\partial q_D^*(\theta)}{\partial K} = \frac{\alpha}{1+\alpha} \frac{\partial \theta_1}{\partial K} = \frac{\alpha}{1+\alpha} \frac{1}{G(\theta_1)}.$$

And for $\theta \in (\theta_2, \bar{\theta})$,

$$q_D^*(\theta) = \frac{\alpha}{1+\alpha}\theta_2 \Rightarrow \frac{\partial q_D^*(\theta)}{\partial K} = \frac{\alpha}{1+\alpha} \frac{\partial \theta_2}{\partial K} = -\frac{\alpha}{1+\alpha} \frac{1}{1-G(\theta_2)}$$

Hence, the strategic effect is

$$\int_{\underline{\theta}}^{\bar{\theta}} \frac{\partial W}{\partial q_D(\theta)} \frac{\partial q_D^*(\theta)}{\partial K} g(\theta) d\theta = -\frac{\alpha^2}{(1-\alpha^2)(1+\alpha)} (\theta_2 - \theta_1) < 0$$

Note that the strategic effect disappears if $\alpha = 0$. Furthermore, it is negative, and its absolute value is increasing in α .

Putting the direct and the strategic effects together

$$\begin{aligned} \frac{dW}{dK} &= -\frac{\alpha^2}{(1-\alpha^2)(1+\alpha)} (\theta_2 - \theta_1) + \frac{1}{(1-\alpha^2)} (\theta_2 - \theta_1) - C'(K) \\ &= \frac{\alpha - \alpha^2 + 1}{(1-\alpha)(\alpha+1)^2} (\theta_2 - \theta_1) - C'(K). \end{aligned}$$

Note that $\frac{\alpha - \alpha^2 + 1}{(1-\alpha)(\alpha+1)^2}$ is increasing in α , and it equals 1 for $\alpha = 0$. It follows that $K^{SB} > K^{FB}$. Note that the solution to the problem implies that the capacity constraint must be binding in the second stage. Otherwise, idle storage capacity would imply that the marginal benefit of storage capacity is lower than its marginal cost, violating the optimality condition at the first stage.

B.6 Proof of Lemma 5

At the production stage, the problem of the competitive storage operator is to solve problem (2.10) subject to constraints (2.2) and (2.1). The structure of the functional optimization problem is identical to the one in the Proof of Proposition 2 in Appendix B.1, with concavity and compactness assumptions satisfied, so a unique solution to the problem exists. The Lagrangian of the problem, omitting the non-negativity constraints, is:

$$\begin{aligned} \mathbb{L}(q_B(\theta), q_S(\theta), \lambda, \mu) &= \int_{\underline{\theta}}^{\bar{\theta}} p(\theta) [q_S(\theta) - q_B(\theta)] g(\theta) d\theta + \lambda \left[\int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} q_S(\theta) g(\theta) d\theta \right] \\ &\quad + \mu \left[K - \int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta \right] \end{aligned}$$

where λ and μ are the Lagrange multipliers. Without loss of generality, we can focus attention on cases in which for any $\theta \in [\underline{\theta}, \bar{\theta}]$, $q_B(\theta) > 0 \rightarrow q_S(\theta) = 0$

& $q_S(\theta) > 0 \rightarrow q_B(\theta) = 0$. The KKT conditions are:

$$p(\theta) - \lambda = 0, \forall \theta \geq \theta_2 \quad (\text{B.22})$$

$$p(\theta) - \lambda < 0, \forall \theta < \theta_2 \quad (\text{B.23})$$

$$p(\theta) - \lambda + \mu = 0, \forall \theta \leq \theta_1 \quad (\text{B.24})$$

$$p(\theta) - \lambda + \mu > 0, \forall \theta > \theta_1 \quad (\text{B.25})$$

$$\int_{\underline{\theta}}^{\theta_1} q_B(\theta)g(\theta)d\theta = \int_{\theta_2}^{\bar{\theta}} q_S(\theta)g(\theta)d\theta \quad (\text{B.26})$$

with $\theta_1 \leq \theta_2$ and the complementary slackness conditions identical to equations (B.4)-(B.8) of the problem in Appendix B.1. Note that condition (B.26) already incorporates the fact that we must have $\lambda > 0$ in any optimal solution to the problem. Note that these conditions are necessary and sufficient, due to the concavity of both the objective function and the constraints.

From condition (B.22):

$$p(\theta) = \lambda, \forall \theta > \theta_2$$

and from condition (B.24):

$$p(\theta) = \lambda - \mu, \forall \theta < \theta_1$$

Since $p(\theta)$ is the best response of the dominant firm,

$$\begin{aligned} \lambda = p(\theta) &= \frac{\theta - q_S(\theta)}{1 - \alpha^2} \Leftrightarrow q_S(\theta) = \theta - (1 - \alpha^2)\lambda, \forall \theta > \theta_2 \\ \lambda - \mu = p(\theta) &= \frac{\theta + q_B(\theta)}{1 - \alpha^2} \Leftrightarrow q_S(\theta) = (1 - \alpha^2)(\lambda - \mu) - \theta, \forall \theta < \theta_1 \end{aligned}$$

By continuity:

$$\begin{aligned} q_S(\theta_2) = 0 &\Rightarrow \theta_2 = (1 - \alpha^2)\lambda \Rightarrow q_S^C(\theta) = \theta - \theta_2, \forall \theta > \theta_2 \\ q_B(\theta_1) = 0 &\Rightarrow \theta_1 = (1 - \alpha^2)(\lambda - \mu) \Rightarrow q_B^C(\theta) = \theta_1 - \theta, \forall \theta < \theta_1 \end{aligned}$$

From condition (B.26):

$$\int_{\underline{\theta}}^{\theta_1} (\theta_1 - \theta)g(\theta)d\theta = \int_{\theta_2}^{\bar{\theta}} (\theta - \theta_2)g(\theta)d\theta \quad (\text{B.27})$$

We have two possible cases depending on the value of the exogenous parameter K . When K is binding, $\mu > 0$ and $\theta_2 - \theta_1 = \mu(1 - \alpha^2) > 0$. Define $x = \theta_2 - \theta_1$

and assume that $G(\theta)$ has a well-defined mean given by $\mathbb{E}(\theta)$. By symmetry of $G(\theta)$, equation (B.27) implies that θ_2 and θ_1 must be symmetric around the mean i.e.

$$\theta_1 = \mathbb{E}(\theta) - \frac{x}{2} \Rightarrow \theta_1^C = \mathbb{E}(\theta) - \frac{\mu^C(1 - \alpha^2)}{2} \quad (\text{B.28})$$

$$\theta_2 = \mathbb{E}(\theta) + \frac{x}{2} \Rightarrow \theta_2^C = \mathbb{E}(\theta) + \frac{\mu^C(1 - \alpha^2)}{2} \quad (\text{B.29})$$

with μ^C implicitly given by:

$$\int_{\underline{\theta}}^{\theta_1(\mu^C)} (\theta_1(\mu^C) - \theta)g(\theta)d\theta = \int_{\theta_2(\mu)}^{\bar{\theta}} (\theta - \theta_2(\mu))g(\theta)d\theta = K. \quad (\text{B.30})$$

Note that when K is not binding, so that $\mu = 0$, from equations (B.28) and (B.29) it is straightforward to establish that $\theta_1^C = \theta_2^C = \mathbb{E}(\theta)$.

B.7 Proof of Proposition 10

The free-entry and perfect competition assumptions imply that entry/investments take place until expected profits are zero, conditional on operating the storage facilities optimally, i.e., $\pi_S(q_S^C(\theta), q_B^C(\theta)) = 0$. Profits of the storage operator at the investment stage are:

$$\begin{aligned} \pi_S(K, \mu(K)) &= \int_{\underline{\theta}}^{\bar{\theta}} p^C(\theta) [q_S^C(\theta) - q_B^C(\theta)] g(\theta) d\theta - C(K) \\ &= \mu^C(K)K - C(K) \end{aligned} \quad (\text{B.31})$$

with $\mu^{SB}(K)$ implicitly given by equation (B.30).

Thus, under the zero-profit condition, the equilibrium investment $K = K^C$ is the unique solution to

$$\pi_S(q_S^C(\theta), q_B^C(\theta)) = 0 \Leftrightarrow \frac{C(K)}{K} = \mu^{SB}(K) = \frac{\theta_2^C - \theta_1^C}{1 - \alpha^2}. \quad (\text{B.32})$$

Note that the solution to the problem implies that the capacity constraint must be binding in the second stage, for the same reasons as in the first-best problem.

Now, we show that $K^C > K^{FB}$. Assume on the contrary that $K^C \leq K^{FB}$.

From equations (B.12) and (B.30), this implies that:

$$\begin{aligned} \int_{\underline{\theta}}^{\theta_1^C} (\theta_1^C - \theta)g(\theta)d\theta &\leq \int_{\underline{\theta}}^{\theta_1^{FB}} (\theta_1^{FB} - \theta)g(\theta)d\theta \rightarrow \theta_1^C \leq \theta_1^{FB} \\ \int_{\theta_2^C}^{\bar{\theta}} (\theta - \theta_2^C)g(\theta)d\theta &\leq \int_{\theta_2^{FB}}^{\bar{\theta}} (\theta - \theta_2^{FB})g(\theta)d\theta \rightarrow \theta_1^C \geq \theta_1^{FB} \end{aligned}$$

Thus, from the optimal solutions (2.4) and (2.14),

$$\theta_2^C - \theta_1^C \geq \theta_2^{FB} - \theta_1^{FB} \Rightarrow \mu^C(1 - \alpha^2) \geq \mu^{FB} \Rightarrow \frac{C(K)}{K}(1 - \alpha^2) \geq C'(K).$$

A contradiction, by strict convexity of the cost function $C(K)$ and $\alpha > 0$.

On the other hand, $K^C > K^{SB}$ follows directly from the strict convexity of $C(K)$.

B.8 Proof of Lemma 6

At the production stage, the problem of the storage monopolist is:

$$\max_{q^S(\theta), q^B(\theta)} \int_{\underline{\theta}}^{\bar{\theta}} \frac{\theta - q^S(\theta) + q^B(\theta) - q^D(\theta)}{1 - \alpha} (q^S(\theta) - q^B(\theta))g(\theta) d\theta,$$

subject to constraints (2.1) and (2.2). The structure of the functional optimization problem is identical to the one in the Proof of Proposition 2, with concavity and compactness assumptions satisfied, so a unique solution to the problem exists. The Lagrangian of the problem, omitting the non-negativity constraints, is given by:

$$\begin{aligned} \mathcal{L} &= \frac{1}{1 - \alpha} \int_{\underline{\theta}}^{\bar{\theta}} [\theta - q^S(\theta) + q^B(\theta) - q^D(\theta)] [q^S(\theta) - q^B(\theta)] g(\theta) d\theta \\ &+ \lambda \left[\int_{\underline{\theta}}^{\bar{\theta}} q^B(\theta)g(\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} q^S(\theta)g(\theta) d\theta \right] + \mu \left[K - \int_{\underline{\theta}}^{\bar{\theta}} q^B(\theta)g(\theta) d\theta \right] \end{aligned}$$

where λ and μ are the Lagrange multipliers. Without loss of generality, we can focus attention on cases in which for any $\theta \in [\underline{\theta}, \bar{\theta}]$, $q_B(\theta) > 0 \rightarrow q_S(\theta) = 0$

& $q_S(\theta) > 0 \rightarrow q_B(\theta) = 0$. The KKT conditions are:

$$\begin{aligned} \frac{\theta - 2q^S(\theta) - q^D(\theta)}{1 - \alpha} - \lambda &= 0, \forall \theta \in (\theta_2, \bar{\theta}) \\ \frac{\theta}{1 - \alpha} - \lambda &< 0, \forall \theta < \theta_2 \\ \frac{\theta + 2q^B(\theta) - q^D(\theta)}{1 - \alpha} - \lambda + \mu &= 0, \forall \theta \in (0, \theta_1) \\ -\frac{\theta + 2q^B(\theta) - q^D(\theta)}{1 - \alpha} + \lambda - \mu &< 0, \forall \theta > \theta_1 \\ \int_{\underline{\theta}}^{\theta_1} q^B(\theta)g(\theta) d\theta &= \int_{\theta_2}^{\bar{\theta}} q^S(\theta)g(\theta) d\theta = K \end{aligned}$$

with $\theta_1 \leq \theta_2$ and the complementary slackness conditions identical to equations (B.4)-(B.8). These conditions are necessary and sufficient, due the concavity of both the objective function and the constraints.

Pointwise optimality implies that the system of reaction functions is:

$$\begin{aligned} q^S(\theta) &= \frac{(\theta - q^D(\theta)) - \lambda(1 - \alpha)}{2} \\ q^B(\theta) &= \frac{(\lambda - \mu)(1 - \alpha) - (\theta - q^D(\theta))}{2} \\ q^D(\theta) &= (\theta - q^S(\theta) + q^B(\theta)) \frac{\alpha}{1 + \alpha} \end{aligned}$$

Thus, we have that:

$$\begin{aligned} q_S(\theta) &= \frac{\theta - (1 - \alpha^2)\lambda}{\alpha + 2}, \forall \theta > \theta_2 \\ q_B(\theta) &= \frac{(1 - \alpha^2)(\lambda - \mu) - \theta}{\alpha + 2}, \forall \theta < \theta_1 \end{aligned}$$

By continuity:

$$\begin{aligned} q_S(\theta_2) = 0 &\Rightarrow \theta_2 = (1 - \alpha^2)\lambda \Rightarrow q^S(\theta) = \frac{\theta - \theta_2}{2 + \alpha}, \forall \theta > \theta_2 \\ q_B(\theta_1) = 0 &\Rightarrow \theta_1 = (1 - \alpha^2)(\lambda - \mu) \Rightarrow q^B(\theta) = \frac{\theta_1 - \theta}{2 + \alpha}, \forall \theta < \theta_1 \end{aligned}$$

As for the market price,

$$\begin{aligned} p(\theta) &= \frac{\theta - q^S(\theta) + q^B(\theta) - q^D(\theta)}{1 - \alpha} \\ &= \frac{1}{(1 - \alpha^2)} (\theta - q^S(\theta) + q^B(\theta)) \end{aligned}$$

Using the expressions for $q^S(\theta)$ and $q^B(\theta)$,

$$p(\theta) = \frac{\theta(1+\alpha) + \theta_2}{(\alpha+2)(1-\alpha^2)}, \quad \forall \theta > \theta_2$$

$$p(\theta) = \frac{\theta(1+\alpha) + \theta_1}{(\alpha+2)(1-\alpha^2)}, \quad \forall \theta < \theta_1$$

Note that

$$p(\theta) - p(\theta_2) = \frac{\theta - \theta_2}{(\alpha+2)(1-\alpha^2)}, \quad \forall \theta > \theta_2$$

$$p(\theta_1) - p(\theta) = \frac{\theta_1 - \theta}{(\alpha+2)(1-\alpha^2)}, \quad \forall \theta < \theta_1$$

From condition (2.2):

$$\int_{\underline{\theta}}^{\theta_1} \frac{\theta_1 - \theta}{2 + \alpha} g(\theta) d\theta = \int_{\theta_2}^{\bar{\theta}} \frac{\theta_1 - \theta}{2 + \alpha} g(\theta) d\theta = K \quad (\text{B.33})$$

We have two possible cases depending on the value of the exogenous parameter K . When K is binding, $\mu > 0$ and $\theta_2 - \theta_1 = \mu(1 - \alpha^2) > 0$. Define $x = \theta_2 - \theta_1$. By symmetry of $G(\theta)$, equation (B.33) implies that θ_2 and θ_1 must be symmetric around the mean i.e.,

$$\theta_1 = \mathbb{E}(\theta) - \frac{x}{2} \Rightarrow \theta_1^M = \mathbb{E}(\theta) - \frac{\mu(1 - \alpha^2)}{2} \quad (\text{B.34})$$

$$\theta_2 = \mathbb{E}(\theta) + \frac{x}{2} \Rightarrow \theta_2^M = \mathbb{E}(\theta) + \frac{\mu(1 - \alpha^2)}{2} \quad (\text{B.35})$$

with $\mu = \mu^M(K)$ implicitly given by:

$$\int_{\underline{\theta}}^{\theta_1(\mu^M(K))} \frac{\theta_1(\mu^M(K)) - \theta}{2 + \alpha} g(\theta) d\theta = \int_{\theta_2(\mu^M(K))}^{\bar{\theta}} \frac{\theta - \theta_2(\mu^M(K))}{2 + \alpha} g(\theta) d\theta = K \quad (\text{B.36})$$

Note that when K is not binding, so that $\mu = 0$, from equations (B.34) and (B.35) it is straightforward to establish that $\theta_1^M = \theta_2^M = \mathbb{E}(\theta)$. However, as shown below, $\mu = 0$ cannot be a solution to the first stage problem.

B.9 Proof of Proposition 11

(i) The problem of the storage firm is to maximize profits, conditional on

operating the storage facilities optimally. Thus, the problem is:

$$\begin{aligned}
 \max_K \pi_S(K, \mu(K)) &= \int_{\underline{\theta}}^{\bar{\theta}} p^M(\theta) [q_S^M(\theta) - q_B^M(\theta)] g(\theta) d\theta - C(K) \\
 &= \frac{1}{(1 - \alpha^2)(2 + \alpha)^2} \int_{\theta_2(\mu(K))}^{\bar{\theta}} [\theta(1 + \alpha) + \theta_2^M(\mu(K))] [\theta - \theta_2^M(\mu(K))] g(\theta) d\theta \\
 &\quad - \frac{1}{(1 - \alpha^2)(2 + \alpha)^2} \int_{\underline{\theta}}^{\theta_1(\mu(K))} [\theta(1 + \alpha) + \theta_1^M(\mu(K))] [\theta_1^M(\mu(K)) - \theta] g(\theta) d\theta \\
 &\quad - C(K)
 \end{aligned} \tag{B.37}$$

with $\mu^M(K)$ implicitly given by equation (B.36).

Taking the derivative with respect to K we have:

$$\frac{\partial \pi_S}{\partial K} = \frac{-1}{(1 - \alpha^2)(2 + \alpha)^2} \left[\int_{\theta_2}^{\bar{\theta}} [\alpha\theta + 2\theta_2] \frac{\partial \theta_2}{\partial K} g(\theta) d\theta + \int_{\underline{\theta}}^{\theta_1} [\alpha\theta + 2\theta_1] \frac{\partial \theta_1}{\partial K} g(\theta) d\theta \right] - C'(K). \tag{B.38}$$

Applying the implicit function theorem to equation (B.36), we can obtain:

$$\begin{aligned}
 \frac{\partial \theta_1}{\partial K} &= \frac{\partial \theta_1}{\partial \mu} \frac{\partial \mu}{\partial K} = \frac{2 + \alpha}{G(\theta_1)} \\
 \frac{\partial \theta_2}{\partial K} &= \frac{\partial \theta_2}{\partial \mu} \frac{\partial \mu}{\partial K} = -\frac{2 + \alpha}{1 - G(\theta_2)}
 \end{aligned}$$

Thus:

$$\begin{aligned}
 \frac{\partial \pi_S}{\partial K} &= \frac{1}{(1 - \alpha^2)(2 + \alpha)G(\theta_1)} \left[2G(\theta_1)(\theta_2 - \theta_1) + \alpha \left(\int_{\theta_2}^{\bar{\theta}} \theta g(\theta) d\theta - \int_{\underline{\theta}}^{\theta_1} \theta g(\theta) d\theta \right) \right] - C'(K) \\
 &= \frac{2\alpha K}{(1 - \alpha^2)G(\theta_1)} + \frac{\theta_2 - \theta_1}{1 - \alpha^2} - C'(K).
 \end{aligned}$$

Therefore,

$$\frac{\partial \pi_S}{\partial K} = 0 \Leftrightarrow \frac{2\alpha K}{(1 - \alpha^2)G(\theta_1)} + \frac{\theta_2 - \theta_1}{1 - \alpha^2} = C'(K). \tag{B.39}$$

(ii) Optimality conditions when $\alpha = 0$ are given by:

$$\begin{aligned}
 C'(K) &= \theta_2^{FB}(K) - \theta_1^{FB}(K) \\
 C'(K) &= \theta_2^{SB}(K) - \theta_1^{SB}(K) \\
 C'(K) &= \theta_2^M(K, \alpha = 0) - \theta_1^M(K, \alpha = 0)
 \end{aligned}$$

Equations (B.12), (B.21) and (B.36) imply that

$$\theta_2^{FB}(K) - \theta_1^{FB}(K) = \theta_2^{SB}(K) - \theta_1^{SB}(K) > \theta_2^M(K, \alpha = 0) - \theta_1^M(K, \alpha = 0)$$

Thus, $K^M < K^{SB} = K^{FB}$ for $\alpha = 0$.

(iii) We just need to show that $K^M < K^{FB}$, as $K^{FB} < K^{SB}$ is always true for $\alpha > 0$. With demand uniformly distributed on $[\underline{\theta}, \bar{\theta}]$, optimal investment in storage capacity is given by:

$$\begin{aligned} \theta_2^{FB}(K^{FB}) - \theta_1^{FB}(K^{FB}) &= C'(K^{FB}) \\ \bar{\theta} - \underline{\theta} - 2\sqrt{2(\bar{\theta} - \underline{\theta})K^{FB}} &= K^{FB} \\ K^{FB} &= [5 - 2\sqrt{6}](\bar{\theta} - \underline{\theta}) \end{aligned}$$

Note that, for θ uniformly distributed on $[\underline{\theta}, \bar{\theta}]$, the marginal investment revenue for the monopolist $MR^M(K)$ is:

$$\begin{aligned} MR^M(K) &= \frac{\theta_2^M - \theta_1^M(K)}{1 - \alpha^2} + \frac{2\alpha K}{(1 - \alpha^2)G(\theta_1^M(K))} \\ &= \frac{(\bar{\theta} - \underline{\theta}) - 2\sqrt{2(2 + \alpha)(\bar{\theta} - \underline{\theta})K}}{1 - \alpha^2} + \frac{\alpha 2K(\bar{\theta} - \underline{\theta})}{(1 - \alpha^2)\sqrt{2(2 + \alpha)(\bar{\theta} - \underline{\theta})K}} \\ &= \frac{1}{(1 - \alpha^2)\sqrt{2 + \alpha}} \left(\sqrt{2 + \alpha}(\bar{\theta} - \underline{\theta}) - (4 + \alpha)\sqrt{2(\bar{\theta} - \underline{\theta})K} \right) \end{aligned}$$

Evaluated at K^{FB} :

$$MR^M(K = K^{FB}) = \frac{\bar{\theta} - \underline{\theta}}{(1 - \alpha^2)\sqrt{2 + \alpha}} \left(\sqrt{2 + \alpha} - (4 + \alpha)\sqrt{2(5 - 2\sqrt{6})} \right) < 0$$

Moreover,

$$\begin{aligned} \frac{\partial MR^M(K)}{\partial K} &= \frac{2}{(2+\alpha)(1-\alpha^2)} \left(\frac{\partial \theta_2^M(K)}{\partial K} - \frac{\partial \theta_1^M(K)}{\partial K} \right) \\ &\quad + \frac{\alpha}{(2+\alpha)(1-\alpha^2)} \left(-\theta_2^M g(\theta_2^M) - \theta_1^M g(\theta_1^M) \right) < 0 \end{aligned}$$

for all K , as $\frac{\partial \theta_2^M(K)}{\partial K} < 0$ and $\frac{\partial \theta_1^M(K)}{\partial K} > 0$. Thus, $K^M(K) < K^{FB}(K)$.

B.10 Proof of Lemma 7

The structure of the functional optimization problem is identical to the one in the Proof of Proposition 2, with concavity and compactness assumptions satisfied, so a unique solution to the problem exists. For a more formal treatment of the problem, we refer the reader to the characterization of the first-best. The Lagrangian of the problem, omitting the non-negativity constraints, is given by:

$$\begin{aligned} \mathbb{L}(p(\theta), q_B(\theta), q_S(\theta), \lambda, \mu) &= \int_{\underline{\theta}}^{\bar{\theta}} \left[p(\theta) D(p; \theta) - \frac{[D(p; \theta) - q_S(\theta) + q_B(\theta)]^2}{2\alpha} \right] g(\theta) d\theta \\ &\quad + \lambda \left[\int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta - \int_{\underline{\theta}}^{\bar{\theta}} q_S(\theta) g(\theta) d\theta \right] + \mu \left[K - \int_{\underline{\theta}}^{\bar{\theta}} q_B(\theta) g(\theta) d\theta \right] \end{aligned}$$

where λ and μ are the Lagrangian multipliers and $D(p; \theta) = \theta - (1-\alpha)p(\theta)$. Without loss of generality, we can focus attention on cases in which for any $\theta \in [\underline{\theta}, \bar{\theta}]$, $q_B(\theta) > 0 \rightarrow q_S(\theta) = 0$ and $q_S(\theta) > 0 \rightarrow q_B(\theta) = 0$. The KKT conditions are:

$$\theta - (1-\alpha)[q_S(\theta) - q_B(\theta)] - (1-\alpha^2)p(\theta) = 0, \forall \theta \quad (\text{B.40})$$

$$\frac{1}{\alpha} [\theta - q_S(\theta) - (1-\alpha)p(\theta)] - \lambda = 0, \forall \theta \geq \theta_2 \quad (\text{B.41})$$

$$\frac{1}{\alpha} [\theta - (1-\alpha)p(\theta)] - \lambda < 0, \forall \theta < \theta_2$$

$$\frac{1}{\alpha} [\theta + q_B(\theta) - (1-\alpha)p(\theta)] - \lambda + \mu = 0, \forall \theta \leq \theta_1 \quad (\text{B.42})$$

$$\frac{1}{\alpha} [\theta - (1-\alpha)p(\theta)] - \lambda + \mu > 0, \forall \theta > \theta_1$$

$$\int_{\underline{\theta}}^{\theta_1} q_B(\theta) g(\theta) d\theta = \int_{\theta_2}^{\bar{\theta}} q_S(\theta) g(\theta) d\theta \quad (\text{B.43})$$

with $\theta_1 \leq \theta_2$ and the complementary slackness conditions identical to those in

equations (B.4)-(B.8). These conditions are necessary and sufficient, due the concavity of both the objective function and the constraints.

Combining conditions (B.40) and (B.41):

$$\begin{aligned} q_S(\theta) &= \frac{\theta - \lambda(1 + \alpha)}{2}, \forall \theta \geq \theta_2 \\ p(\theta) &= \frac{\theta + \lambda(1 - \alpha)}{2(1 - \alpha)}, \forall \theta \geq \theta_2 \end{aligned}$$

From conditions (B.40) and (B.42):

$$\begin{aligned} q_B(\theta) &= \frac{(\lambda - \mu)(1 + \alpha) - \theta}{2}, \forall \theta \leq \theta_1 \\ p(\theta) &= \frac{\theta + (\lambda - \mu)(1 - \alpha)}{2(1 - \alpha)}, \forall \theta \leq \theta_1 \end{aligned}$$

And from condition (B.40):

$$p(\theta) = \frac{\theta}{1 - \alpha^2} \text{ for } \theta_1 < \theta < \theta_2.$$

By continuity:

$$\begin{aligned} q_S(\theta_2) = 0 \Rightarrow \theta_2 = (1 + \alpha)\lambda \Rightarrow & \begin{cases} q_S^I(\theta) = \frac{\theta - \theta_2}{2}, \forall \theta > \theta_2 \\ p(\theta) = \frac{\theta}{2(1 - \alpha)} + \frac{\theta_2}{2(1 + \alpha)}, \forall \theta > \theta_2 \end{cases} \\ q_B(\theta_1) = 0 \Rightarrow \theta_1 = (1 + \alpha)(\lambda - \mu) \Rightarrow & \begin{cases} q_B^I(\theta) = \frac{\theta_1 - \theta}{2}, \forall \theta < \theta_1 \\ p(\theta) = \frac{\theta}{2(1 - \alpha)} + \frac{\theta_1}{2(1 + \alpha)}, \forall \theta < \theta_1 \end{cases} \end{aligned}$$

From condition B.43:

$$\int_{\underline{\theta}}^{\theta_1} \frac{\theta_1 - \theta}{2} g(\theta) d\theta = \int_{\theta_2}^{\bar{\theta}} \frac{\theta - \theta_2}{2} g(\theta) d\theta. \quad (\text{B.44})$$

We have two possible cases depending on the value of the exogenous parameter K . When K is binding, $\mu > 0$ and $\theta_2 - \theta_1 = \mu(1 + \alpha) > 0$. Define $x = \theta_2 - \theta_1$ and assume that $G(\theta)$ has a well-defined mean given by $\mathbb{E}(\theta)$. By symmetry of $G(\theta)$, equation (B.44) implies that θ_2 and θ_1 must be symmetric around the

mean i.e.,

$$\theta_1 = \mathbb{E}(\theta) - \frac{x}{2} \Rightarrow \theta_1^C = \mathbb{E}(\theta) - \frac{\mu(1+\alpha)}{2} \quad (\text{B.45})$$

$$\theta_2 = \mathbb{E}(\theta) + \frac{x}{2} \Rightarrow \theta_2^C = \mathbb{E}(\theta) + \frac{\mu(1+\alpha)}{2} \quad (\text{B.46})$$

with $\mu = \mu^I(K)$ implicitly given by:

$$\int_{\underline{\theta}}^{\theta_1(\mu^I(K))} \frac{\theta_1(\mu^I(K)) - \theta}{2} g(\theta) d\theta = \int_{\theta_2(\mu^I(K))}^{\bar{\theta}} \frac{\theta - \theta_2(\mu^I(K))}{2} g(\theta) d\theta = K \quad (\text{B.47})$$

Note that when K is not binding, so that $\mu = 0$, from equations (B.45) and (B.46) it is straightforward to establish that $\theta_1^I = \theta_2^I = \mathbb{E}(\theta)$. However, as shown below, $\mu = 0$ cannot be a solution to the first stage problem.

B.11 Proof of Corollary 1

Let us use $L(\theta)$ to denote the Lerner Index, i.e., the ratio between price minus marginal cost over price. Using the equilibrium storage decisions,

$$L(\theta) = \begin{cases} \frac{\theta(1+\alpha) - \theta_1^I(1-\alpha)}{\theta(1+\alpha) + \theta_1^I(1-\alpha)} & \text{if } \theta < \theta_1^I \\ \alpha & \text{if } \theta_1^I \leq \theta \leq \theta_2^I \\ \frac{\theta(1+\alpha) - \theta_2^I(1-\alpha)}{\theta(1+\alpha) + \theta_2^I(1-\alpha)} & \text{if } \theta > \theta_2^I \end{cases}$$

These mark-ups are continuous in θ . They are constant for $\theta_1^I \leq \theta \leq \theta_2^I$ and increasing in θ otherwise. Hence, for $\theta < \theta_1^I$, $L(\theta) < \alpha$. And for $\theta > \theta_2^I$, $L(\theta) > \alpha$. Since the two expressions are a mirror image of each other, while the markups for high demand levels are weighted more, the demand-weighted average mark-up is greater than α .

B.12 Proof of Proposition 12

By identical arguments to those in the proof for optimal First Best capacity investment (Section B.2), the unique interior solution K^I is given by the solution to:

$$\mu - C'(K) = \frac{\theta_2^I(K) - \theta_1^I(K)}{1+\alpha} - C'(K) = 0 \quad (\text{B.48})$$

with $\mu = \mu^I(K)$ implicitly given by the solution to

$$\int_{\underline{\theta}}^{\theta_1(\mu(K))} \frac{\theta_1(\mu(K)) - \theta}{2} g(\theta) d\theta = \frac{1}{2} \int_{\theta_2(\mu(K))}^{\bar{\theta}} \frac{\theta - \theta_2(\mu(K))}{2} g(\theta) d\theta = K.$$

Now we show that $K^I < K^{FB}$. Assume on the contrary that $K^I \geq K^{FB}$. Then, by strict convexity of the cost function:

$$K^I \geq K^{FB} \Rightarrow C'(K^I) \geq C'(K^{FB}) \quad (\text{B.49})$$

Moreover:

$$K^I \geq K^{FB} \Rightarrow \theta_2^I - \theta_1^I < \theta_2^{FB} - \theta_1^{FB} \Rightarrow (1 + \alpha)C'(K^I) < C'(K^{FB}) \quad (\text{B.50})$$

where the first implication comes from the capacity constraints in the optimal solution (equations (B.12) and (B.47)), and the second from first stage optimality conditions (2.4) and (2.17). Putting (B.49) and (B.50) together:

$$(1 + \alpha)C'(K^I) < C'(K^{FB}) \leq C'(K^I)$$

B.13 Proof of Proposition 14

From Lemma 5, we have that, for given K , equilibrium storage decisions under the second-best and under competitive storage coincide. Thus, $W^{SB}(K) = W^C(K)$ and $CS^{SB}(K) = CS^C(K)$, $\forall K$. Moreover,

$$\{q_B^{SB}(\theta, K), q_S^{SB}(\theta, K)\} = \arg \max_{q_B(\theta, K), q_S(\theta, K)} W(K)$$

Thus, as $q_B^M(\theta, K) \neq q_B^{SB}(\theta, K)$ and $q_S^I(\theta, K) \neq q_S^{SB}(\theta, K)$ for some θ and all K , we have that $W^{SB}(K) > W^M(K)$ and $W^{SB}(K) > W^I(K)$, $\forall K$.

For the second part of the proposition, we want to show that $W^M(K) > W^I(K)$, which is equivalent to $TC^M(K) > TC^I(K)$. Note that:

$$\begin{aligned} \lim_{K \rightarrow \infty} [TC^M(K) - TC^I(K)] &= V[\theta] \frac{5\alpha}{8(1-\alpha)(2+\alpha)^2} > 0 \\ \lim_{K \rightarrow \infty} [TC^M(K) - TC^I(K)] &= 0 \end{aligned}$$

Thus, $TC^M(K) < TC^I(K)$ if:

$$\frac{\partial TC^M(K)}{\partial K} < 0 \quad (\text{B.51})$$

$$\frac{\partial TC^I(K)}{\partial K} < 0 \quad (\text{B.52})$$

$$\left| \frac{\partial TC^M(K)}{\partial K} \right| \geq \left| \frac{\partial TC^I(K)}{\partial K} \right| \quad (\text{B.53})$$

for all $K < K^M(max)$, where:

$$K^M(max) = \int_{\underline{\theta}}^{\mathbb{E}[\theta]} \frac{\mathbb{E}[\theta] - \theta}{2 + \alpha} g(\theta) d\theta.$$

Recall that

$$TC(K) = \int_{\underline{\theta}}^{\bar{\theta}} \left(\frac{\alpha q_F^2(\theta, K)}{2} + \frac{q_F^2(\theta, K)}{2(1 - \alpha)} \right) g(\theta) d\theta$$

For the monopolist:

$$TC^M(K) = \frac{1 + \alpha - \alpha^2}{2(1 - \alpha)} \left[\int_{\underline{\theta}}^{\theta_1^M(K)} \left(\frac{\theta(1 + \alpha) + \theta_1^M(K)}{(2 + \alpha)(1 + \alpha)} \right)^2 g(\theta) d\theta + \int_{\theta_1^M(K)}^{\theta_2^M(K)} \left(\frac{\theta}{(1 + \alpha)} \right)^2 g(\theta) d\theta + \int_{\theta_2^M(K)}^{\bar{\theta}} \left(\frac{\theta(1 + \alpha) + \theta_2^M(K)}{(2 + \alpha)(1 + \alpha)} \right)^2 g(\theta) d\theta \right]$$

For the vertically integrated firm

$$TC^I(K) = \int_{\underline{\theta}}^{\theta_1^M(K)} \left(\frac{1}{2\alpha} \left(\frac{\alpha \theta_1^I(K)}{1 + \alpha} \right)^2 + \frac{1}{2(1 - \alpha)} \left(\frac{(1 + \alpha)\theta + (1 - \alpha)\theta_1^I(K)}{2(1 + \alpha)} \right)^2 \right) g(\theta) d\theta + \frac{1 + \alpha - \alpha^2}{2(1 - \alpha)} \int_{\theta_1^M(K)}^{\theta_2^M(K)} \left(\frac{\theta}{(1 + \alpha)} \right)^2 g(\theta) d\theta + \int_{\theta_2^M(K)}^{\bar{\theta}} \left(\frac{1}{2\alpha} \left(\frac{\alpha \theta_2^I(K)}{1 + \alpha} \right)^2 + \frac{1}{2(1 - \alpha)} \left(\frac{(1 + \alpha)\theta + (1 - \alpha)\theta_2^I(K)}{2(1 + \alpha)} \right)^2 \right) g(\theta) d\theta$$

For the independent monopolist, after some computations we have that:

$$\begin{aligned} \frac{\partial TC^M(K)}{\partial K} &= \frac{1 + \alpha - \alpha^2}{(1 - \alpha)(2 + \alpha)(1 + \alpha)^2 G(\theta_1^M(K))} \left[\int_{\underline{\theta}}^{\theta_1^M(K)} [\theta(1 + \alpha) + \theta_1^M(K)] g(\theta) d\theta \right. \\ &\quad \left. - \int_{\theta_2^M(K)}^{\bar{\theta}} [\theta(1 + \alpha) + \theta_2^M(K)] g(\theta) d\theta \right] < 0 \\ \left| \frac{\partial TC^M(K)}{\partial K} \right| &= \frac{1 + \alpha - \alpha^2}{(1 - \alpha)(2 + \alpha)(1 + \alpha)^2} [\theta_2^M(K) - \theta_1^M(K)] \\ &\quad + \frac{1 + \alpha - \alpha^2}{(1 - \alpha)(2 + \alpha)(1 + \alpha) G(\theta_1^M(K))} \left[\int_{\theta_2^M(K)}^{\bar{\theta}} \theta g(\theta) d\theta - \int_{\underline{\theta}}^{\theta_1^M(K)} \theta g(\theta) d\theta \right] \end{aligned}$$

and for the vertically integrated firm:

$$\begin{aligned} \frac{\partial TC^I(K)}{\partial K} &= \frac{1}{2(1 + \alpha)^2} \frac{1}{G(\theta_1^I(K))} \left[\int_{\underline{\theta}}^{\theta_1^I(K)} [(1 + \alpha)\theta + (1 + 3\alpha)\theta_1^I(K)] g(\theta) d\theta \right. \\ &\quad \left. - \int_{\theta_2^I(K)}^{\bar{\theta}} [(1 + \alpha)\theta + (1 + 3\alpha)\theta_2^I(K)] g(\theta) d\theta \right] < 0 \\ \left| \frac{\partial TC^I(K)}{\partial K} \right| &= \frac{(1 + 3\alpha)}{2(1 + \alpha)^2} [\theta_2^I(K) - \theta_1^I(K)] + \frac{1}{2(1 + \alpha)G(\theta_1^I(K))} \left[\int_{\theta_2^I(K)}^{\bar{\theta}} \theta g(\theta) d\theta - \int_{\underline{\theta}}^{\theta_1^I(K)} \theta g(\theta) d\theta \right] \end{aligned}$$

Assuming that θ is uniformly distributed on $[\underline{\theta}, \bar{\theta}]$, we have that:

$$\begin{aligned} \left| \frac{\partial TC^M(K)}{\partial K} \right| &= \frac{1 + \alpha - \alpha^2}{(1 - \alpha)(2 + \alpha)(1 + \alpha)^2} \left[(2 + \alpha)(\bar{\theta} - \underline{\theta}) - (3 + \alpha)\sqrt{2(2 + \alpha)(\bar{\theta} - \underline{\theta})K} \right] \\ \left| \frac{\partial TC^I(K)}{\partial K} \right| &= \frac{1}{(1 + \alpha)^2} \left[(1 + 2\alpha)(\bar{\theta} - \underline{\theta}) + (7\alpha + 3)\sqrt{(\bar{\theta} - \underline{\theta})K} \right] \end{aligned}$$

Thus:

$$\left| \frac{\partial TC^M(K)}{\partial K} \right| - \left| \frac{\partial TC^I(K)}{\partial K} \right| > 0 \Leftrightarrow K < \left(\frac{\alpha^2(2 + \alpha)\sqrt{\bar{\theta} - \underline{\theta}}}{(1 + \alpha - \alpha^2)(3 + \alpha)\sqrt{2(2 + \alpha)} - (7\alpha + 3)(1 - \alpha)(2 + \alpha)} \right)^2$$

Note that:

$$\left(\frac{\alpha^2(2 + \alpha)\sqrt{\bar{\theta} - \underline{\theta}}}{(1 + \alpha - \alpha^2)(3 + \alpha)\sqrt{2(2 + \alpha)} - (7\alpha + 3)(1 - \alpha)(2 + \alpha)} \right)^2 < \frac{(\mathbb{E}[\theta] - \underline{\theta})^2}{2(\bar{\theta} - \underline{\theta})(2 + \alpha)} = K^M(max)$$

Thus, $\left| \frac{\partial TC^M(K)}{\partial K} \right| > \left| \frac{\partial TC^I(K)}{\partial K} \right|$, so that $TC^M(K) < TC^I(K)$ and $W^M(K) > W^I(K)$.

For comparing consumer surplus, we follow the same approach. First note that:

$$\begin{aligned} \lim_{K \rightarrow \infty} [CS^M(K) - CS^I(K)] &= \frac{\alpha}{2 + \alpha} V[\theta] > 0 \\ \lim_{K \rightarrow \infty} [CS^M(K) - CS^I(K)] &= 0 \end{aligned}$$

Second, note that:

$$\frac{\partial CS^M(K)}{\partial K} = \frac{1}{(1 - \alpha^2)G(\theta_1^M(K))} \left[\int_{\theta_2^M(K)}^{\bar{\theta}} \theta g(\theta) d\theta - \int_{\underline{\theta}}^{\theta_1^M(K)} \theta g(\theta) d\theta \right] > 0$$

for all K and

$$\frac{\partial CS^I(K)}{\partial K} = \frac{1}{(1 + \alpha)G(\theta_1^I(K))} \left[\int_{\theta_2^I(K)}^{\bar{\theta}} \theta g(\theta) d\theta - \int_{\underline{\theta}}^{\theta_1^I(K)} \theta g(\theta) d\theta \right] > 0$$

for all K .

Finally, when θ is uniformly distributed on $[\underline{\theta}, \bar{\theta}]$, we have that:

$$\begin{aligned} \frac{\partial CS^M(K)}{\partial K} - \frac{\partial CS^I(K)}{\partial K} > 0 &\Leftrightarrow \frac{\bar{\theta} - \underline{\theta} - \sqrt{2(2 + \alpha)(\bar{\theta} - \underline{\theta})K}}{1 - \alpha^2} - \frac{\bar{\theta} - \underline{\theta} - 2\sqrt{(\bar{\theta} - \underline{\theta})K}}{1 + \alpha} > 0 \\ &\left(\frac{\alpha\sqrt{\bar{\theta} - \underline{\theta}}}{\sqrt{2(2 + \alpha) + 2\alpha - 2}} \right)^2 > K \end{aligned}$$

Note that:

$$\left(\frac{\alpha\sqrt{\bar{\theta} - \underline{\theta}}}{\sqrt{2(2 + \alpha) + 2\alpha - 2}} \right)^2 < \frac{(\mathbb{E}[\theta] - \underline{\theta})^2}{2(\bar{\theta} - \underline{\theta})(2 + \alpha)} = K^M(\max)$$

Thus, $\frac{\partial CS^M(K)}{\partial K} > \frac{\partial CS^I(K)}{\partial K}$, so that $CS^M(K) > CS^I(K)$, for all K .

Bibliography

- Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: A continuous-time approach. *National Bureau of Economic Research*, 2017.
- Abi Adams, Richard Blundell, Martin Browning, and Ian Crawford. Prices versus preferences: taste change and revealed preference. Technical report, IFS Working Papers, 2015.
- Matthew D Adler. Benefit–cost analysis and distributional weights: an overview. *Review of Environmental Economics and Policy*, 10(2):264–285, 2016.
- Philippe Aghion, Xavier Jaravel, Torsten Persson, and Dorothée Rouzet. Education and military rivalry. *Journal of the European Economic Association*, 17(2):376–412, 2018.
- George A Akerlof and Rachel E Kranton. Economics and identity. *The quarterly journal of economics*, 115(3):715–753, 2000.
- Alberto Alesina and Bryony Reich. Nation Building. *National Bureau of Economic Research*, 2 2013.
- Alberto Alesina, Bryony Reich, and Alessandro Riboni. Nation-Building, Nationalism and Wars. *National Bureau of Economic Research*, 5 2017.
- Ingela Alger and Jörgen W Weibull. Evolutionary models of preference formation. *Annual Review of Economics*, 2019.
- Blaise Allaz. Duopoly, inventories and futures markets. In *Commodity, futures and financial markets*, pages 249–271. Springer, 1991.

- Stefan Ambec and Claude Crampes. Decarbonizing electricity generation with intermittent sources of energy. *Journal of the Association of Environmental and Resource Economists*, 6(6):919–948, 2019.
- Benedict Anderson. *Imagined communities: reflections on the origin and spread of nationalism*. Verso, 1991.
- David Andrés-Cerezo. Environmental activists and the dynamics of green consumer culture and environmental policies. *Working paper*, 2021.
- Luc Arrondel. ” my father was right”: The transmission of values between generations. 2009.
- Kenneth J Arrow. A difficulty in the concept of social welfare. *Journal of political economy*, 58(4):328–346, 1950.
- Oriol Aspachs-Bracons, Irma Clots-Figueras, Joan Costa-Font, and Paolo Masella. Compulsory Language Educational Policies and Identity Formation. *Journal of the European Economic Association*, 6(2-3):434–444, 4 2008.
- Kristin M. Bakke and Erik Wibbels. Diversity, Disparity, and Civil Conflict in Federal States. *World Politics*, 59(01):1–50, 10 2006.
- Laia Balcells. Mass Schooling and Catalan Nationalism. *Nationalism and Ethnic Politics*, 19(4):467–486, 10 2013.
- Oriana Bandiera, Myra Mohnen, Imran Rasul, and Martina Viarengo. Nation-building through compulsory schooling during the age of mass migration. *The Economic Journal*, 2015.
- Joel D. Barkan. *Beyond Capitalism versus Socialism in Kenya and Tanzania*. East African Publishers, 1994.
- Gary S Becker. *Accounting for tastes*. Harvard University Press, 1998.
- Gary S Becker and Kevin M Murphy. A theory of rational addiction. *Journal of political Economy*, 96(4):675–700, 1988.
- Gary S Becker and George J Stigler. De gustibus non est disputandum. *American Economic Review*, 67(2):76–90, 1977.
- Gary S Becker et al. Habits, addictions, and traditions. Technical report, Chicago-Center for Study of Economy and State, 1991.

- Jess Benhabib and Alberto Bisin. Social construction of preferences: advertising. In *Handbook of Social Economics*, volume 1, pages 201–220. Elsevier, 2011.
- Theodore C Bergstrom. Benefit-cost in a benevolent society. *American Economic Review*, 96(1):339–351, 2006.
- B Douglas Bernheim and Antonio Rangel. Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *The Quarterly Journal of Economics*, 124(1):51–104, 2009.
- B Douglas Bernheim, Luca Braghieri, Alejandro Martínez-Marquina, and David Zuckerman. A theory of chosen preferences. *American Economic Review*, 111(2):720–54, 2021.
- Timothy Besley and Torsten Persson. *Pillars of prosperity: The political economics of development clusters*. Princeton University Press, 2011.
- Timothy Besley and Torsten Persson. Democratic values and institutions. *American Economic Review: Insights*, 1(1):59–76, 2019a.
- Timothy Besley and Torsten Persson. Jeea-fbbva lecture 2017: The dynamics of environmental politics and values. *Journal of the European Economic Association*, 17(4):993–1024, 2019b.
- Timothy Besley and Marta Reynal-Querol. The Legacy of Historical Conflict: Evidence from Africa. *American Political Science Review*, 2014.
- Emeline Bezin. A cultural model of private provision and the environment. *Journal of Environmental Economics and Management*, 71:109–124, 2015.
- Leann L Birch. Development of food preferences. *Annual review of nutrition*, 19(1):41–62, 1999.
- Alberto Bisin and Thierry Verdier. A model of cultural transmission, voting and political ideology. *European Journal of Political Economy*, 16(1):5–29, 2000.
- Alberto Bisin and Thierry Verdier. The Economics of Cultural Transmission and the Dynamics of Preferences. *Journal of Economic Theory*, 97(2):298–319, 2001.

- Alberto Bisin and Thierry Verdier. The economics of cultural transmission and socialization. In *Handbook of social economics*, volume 1, pages 339–416. Elsevier, 2011.
- BloombergNEF. New energy outlook 2020. Technical report, BloombergNEF, 2020.
- Arthur Blouin and Sharun W Mukand. Erasing ethnicity? propaganda, nation building, and identity in rwanda. *Journal of Political Economy*, 127(3):1008–1062, 2019.
- Severin Borenstein. The private and public economics of renewable electricity generation. *Journal of Economic Perspectives*, 26(1):67–92, 2012.
- Samuel Bowles. Endogenous preferences: The cultural consequences of markets and other economic institutions. *Journal of economic literature*, 36(1):75–111, 1998.
- Samuel Bowles and Sung-Ha Hwang. Social preferences and public economics: Mechanism design when social preferences depend on incentives. *Journal of public economics*, 92(8-9):1811–1820, 2008.
- Albert Breton et al. The economics of nationalism. *Journal of Political Economy*, 72:376–376, 1964.
- James M Buchanan. Natural and artifactual man. *What should economists do*, pages 93–112, 1979.
- Undral Byambadalai, Ching-to Albert Ma, and Daniel Wiesen. Changing preferences: An experiment and estimation of market-incentive effects on altruism. *Available at SSRN 3431002*, 2019.
- Ramon Caminal, Lorenzo Cappellari, and Antonio Di Paolo. Linguistic skills and the intergenerational transmission of language. *IZA Discussion Paper*, 2018.
- Davide Cantoni, Yuyu Chen, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang. Curriculum and Ideology. *Journal of Political Economy*, 125(2): 338–392, 4 2017.
- Jean-Paul Carvalho and Mark Koyama. Jewish emancipation and schism: Economic development and religious change. *Journal of Comparative Economics*, 44(3):562–584, 8 2016.

- Frank Convery, Simon McDonnell, and Susana Ferreira. The most popular tax in europe? lessons from the irish plastic bags levy. *Environmental and resource economics*, 38(1):1–11, 2007.
- Francisco Costa and François Gerard. Hysteresis and the social cost of corrective policies: evidence from a temporary energy saving program. *Available at SSRN 2666704*, 2015.
- Claude Crampes and Jean-Michel Trochet. Economics of stationary electricity storage with various charge and discharge durations. *Journal of Energy Storage*, 24:100746, 2019.
- Peter Cramton and Steven Stoft. Forward reliability markets: Less risk, less market power, more efficiency. *Utilities Policy*, 16(3):194–201, 2008.
- María-Ángeles De Frutos and Natalia Fabra. How to allocate forward contracts: The case of electricity markets. *European Economic Review*, 56(3):451–469, 2012.
- Borja De Riquer. *Escolta, Espanya: la cuestión catalana en la época liberal*. Marcial Pons Historia, 2001.
- Emilio Depetris-Chauvin, Ruben Durante, and Filipe R Campante. Building nations through shared experiences: Evidence from african football. *National Bureau of Economic Research*, 2018.
- Roberta Dessí. Collective Memory, Cultural Transmission, and Investments. *American Economic Review*, 98(1):534–560, 2 2008.
- Franz Dietrich and Christian List. Where do preferences come from? *International Journal of Game Theory*, 42(3):613–637, 2013.
- Thomas Dohmen, Armin Falk, David Huffman, and Uwe Sunde. The intergenerational transmission of risk and trust attitudes. *The Review of Economic Studies*, 79(2):645–677, 2012.
- Malte F Dold. Back to buchanan? explorations of welfare and subjectivism in behavioral economics. *Journal of Economic Methodology*, 25(2):160–178, 2018.
- Jonas Eliasson. Lessons from the stockholm congestion charging trial. *Transport Policy*, 15(6):395–404, 2008.

- J Elster. 1) "sour grapes: utilitarianism and the genesis of wants". *Utilitarianism and beyond*, Cambridge, 2, 1982.
- Mariano Fernández Enguita. Building the nation at school: Spains tables turned: Mariano fernández enguita. In *The Politics of Education*, pages 83–104. Routledge, 2012.
- EU European Commission. Eu clean energy package. Technical report, Official Journal of the European Union, 2019.
- EU European Commission. Study on energy storage. contribution to the security of the electricity supply in europe. Technical report, European Commission., 2020.
- Natalia Fabra, David Rapson, Mar Reguant, and J. Wang. Estimating the elasticity to real time pricing: Evidence from the spanish electricity market. Technical report, Mimeo, 2020.
- Marcus Fleming. A cardinal concept of welfare. *The Quarterly Journal of Economics*, 66(3):366–384, 1952.
- Marc Fleurbaey and Rossi Abi-Rafeh. The use of distributional weights in benefit–cost analysis: Insights from welfare economics. *Review of Environmental Economics and Policy*, 10(2):286–307, 2016.
- Marc Fleurbaey and Didier Blanchet. *Beyond GDP: Measuring welfare and assessing sustainability*. Oxford University Press, 2013.
- Marc Fleurbaey and Koichi Tadenuma. Universal social orderings: An integrated theory of policy evaluation, inter-society comparisons, and interpersonal comparisons. *Review of Economic Studies*, 81(3):1071–1101, 2014.
- Vasiliki Fouka. Backlash: The unintended effects of language prohibition in us schools after world war i. *Review of Economic Studies*, Forthcoming, 2019.
- Bruno S Frey and Felix Oberholzer-Gee. The cost of price incentives: An empirical analysis of motivation crowding-out. *The American economic review*, 87(4):746–755, 1997.
- Bruno S Frey, Simon Luechinger, and Alois Stutzer. The life satisfaction approach to environmental valuation. *Annu. Rev. Resour. Econ.*, 2(1):139–160, 2010.

- John Kenneth Galbraith. *The affluent society*. Houghton Mifflin Harcourt, 1998.
- Alfredo Garcia, James D Reitzes, and Ennio Stacchetti. Strategic pricing when electricity is storable. *Journal of Regulatory Economics*, 20(3):223–247, 2001.
- Ernest Gellner and John Breuilly. *Nations and nationalism*, volume 1. Cornell University Press Ithaca, NY, 1983.
- Nicola Gennaioli and Guido Tabellini. Identity, beliefs, and political conflict. *Available at SSRN 3300726*, 2018.
- Herbert Gintis. Welfare criteria with endogenous preferences: the economics of education. *International Economic Review*, pages 415–430, 1974.
- Steven Marc Goldman. Intertemporally inconsistent preferences and the rate of consumption. *Econometrica: Journal of the Econometric Society*, pages 621–626, 1979.
- Dieter Grass, Jonathan P Caulkins, Gustav Feichtinger, Gernot Tragler, and Doris A Behrens. Optimal control of nonlinear processes. *Berlino: Springer*, 2008.
- Richard J. Green and David M. Newbery. Competition in the british electricity spot market. *Journal of Political Economy*, 100(5):929–953, 1992.
- Luigi Guiso, Helios Herrera, and Massimo Morelli. Cultural Differences and Institutional Integration. *Journal of International Economics*, 99:S97–S113, 3 2016a.
- Luigi Guiso, Paola Sapienza, and Luigi Zingales. Monnets error? *Economic Policy*, 31(86):247–297, 4 2016b.
- Faruk Gul and Wolfgang Pesendorfer. Welfare without happiness. *American Economic Review*, 97(2):471–476, 2007.
- Werner Güth. An evolutionary approach to explaining cooperative behavior by reciprocal incentives. *International Journal of Game Theory*, 24(4):323–344, 1995.
- John C Harsanyi. Welfare economics of variable tastes. *The Review of Economic Studies*, 21(3):204–213, 1953.

- John C Harsanyi. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of political economy*, 63(4):309–321, 1955.
- David Harvey et al. *The condition of postmodernity*, volume 14. Blackwell Oxford, 1990.
- Esther Hauk and Giovanni Immordino. Parents, Television and Cultural Change. *The Economic Journal*, 124(579):1040–1065, 9 2014.
- Corinna Hawkes, Trenton G Smith, Jo Jewell, Jane Wardle, Ross A Hammond, Sharon Friel, Anne Marie Thow, and Juliana Kain. Smart food policies for obesity prevention. *The Lancet*, 385(9985):2410–2421, 2015.
- María José Hierro and Aina Gallego. Identities in between: Political conflict and ethnonational identities in multicultural states. *Journal of Conflict Resolution*, 62(6):1314–1339, 2018.
- Harold Hotelling. The economics of exhaustible resources. *Journal of political Economy*, 39(2):137–175, 1931.
- John Iliffe. *A modern history of Tanganyika*. Cambridge University Press, 1979.
- Electricity IRENA. Storage and renewables: costs and markets to 2030. *Abu Dhabi.: Int Renew Energy Agency*, 2017.
- Koichiro Ito, Takanori Ida, and Makoto Tanaka. Moral suasion and economic incentives: Field experimental evidence from energy demand. *American Economic Journal: Economic Policy*, 10(1):240–67, 2018.
- Noel D. Johnson. Taxes, National Identity, and Nation Building: Evidence from France. *SSRN Electronic Journal*, 4 2015.
- Paul Joskow and Jean Tirole. Transmission rights and market power on electric power networks. *The Rand Journal of Economics*, pages 450–487, 2000.
- Paul Joskow and Jean Tirole. Merchant transmission investment. *The Journal of industrial economics*, 53(2):233–264, 2005.
- Daniel Kahneman. *Thinking, fast and slow*. Macmillan, 2011.
- Louis Kaplow and Steven Shavell. Fairness versus welfare. *Harv. L. Rev.*, 114: 961, 2000.

- Louis Kaplow and Steven Shavell. Any non-welfarist method of policy assessment violates the pareto principle. *Journal of Political Economy*, 109(2): 281–286, 2001.
- Omer Karaduman. Economics of grid-scale energy storage. 2020.
- Kai A Konrad and Salmay Qari. The last refuge of a scoundrel? patriotism and tax compliance. *Economica*, 79(315):516–533, 2012.
- David D. Laitin. *Nations, states, and violence*. Oxford University Press, 2007.
- Bruno Lanz, Jules-Daniel Wnourlod, Luca Panzone, and Timothy Swanson. The behavioral effect of pigovian regulation: Evidence from a field experiment. *Journal of environmental economics and management*, 87:190–205, 2018.
- Richard Layard. *Happiness: Lessons from a new science*. Penguin UK, 2011.
- Assar Lindbeck and Jörgen W Weibull. Balanced-budget redistribution as the outcome of political competition. *Public choice*, 52(3):273–297, 1987.
- George Loewenstein and Emily Celia Haisley. The economist as therapist: Methodological ramifications of ‘light’ paternalism. *Available at SSRN 962472*, 2007.
- George Loewenstein and Drazen Prelec. Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*, 107(2):573–597, 1992.
- Andreas Madestam and David Yanagizawa-Drott. Shaping of the Nation: The Effect of Fourth of July on Political Preferences and Behavior in the United States. *Working Paper Series*, 2012.
- Anthony W Marx. *Faith in nation: exclusionary origins of nationalism*. Oxford University Press, 2005.
- Olivier Massol and Albert Banal-Estanol. Market power and spatial arbitrage between interconnected gas hubs. *Energy Journal*, 39(2):67–95, 2018.
- Linus Mattauch, Cameron Hepburn, and Nicholas Stern. Pigou pushes preferences: decarbonisation and endogenous values. 2018.
- Rory McCarthy and Dan Eager. Europe power system flexibility: The essential ingredient for decarbonisation. Technical report, Wood Mackenzie., 2020.

- David McCrone. *The sociology of nationalism : tomorrow's ancestors*. Routledge, 1998.
- John McGarry. 'Demographic engineering': the state-directed movement of ethnic groups as a technique of conflict regulation. *Ethnic and Racial Studies*, 21(4):613–638, 1 1998.
- John McLaren. Speculation on primary commodities: The effects of restricted entry. *The Review of Economic Studies*, 66(4):853–871, 1999.
- Kenneth McRoberts. *Catalonia: Nation building without a state*. Don Mills, Ont.: Oxford University Press Canada, 2001.
- Edward Miguel. Tribe or nation? nation building and public goods in kenya versus tanzania. *World politics*, 56(3):327–362, 2004.
- MITECO. Plan nacional integrado de energia y clima 2021-2030. Technical report, Ministerio de Transición Ecológica, Gobierno de España., 2020.
- Sébastien Mitraillet and Henry Thille. Speculative storage in imperfectly competitive markets. *International Journal of Industrial Organization*, 35:44–59, 2014.
- Ehsan Nasrolahpour, S Jalal Kazempour, Hamidreza Zareipour, and William D Rosehart. Strategic sizing of energy storage facilities in electricity markets. *IEEE Transactions on Sustainable Energy*, 7(4):1462–1472, 2016.
- Eric Neumayer. In defence of historical accountability for greenhouse gas emissions. *Ecological economics*, 33(2):185–192, 2000.
- David M. Newbery. Cartels, storage, and the suppression of futures markets. *European Economic Review*, 34(5):1041 – 1060, 1990.
- David MG Newbery and Joseph E Stiglitz. The theory of commodity price stabilisation rules: Welfare impacts and supply responses. *The Economic Journal*, 89(356):799–817, 1979.
- Yew-Kwang Ng. *Efficiency, equality and public policy: With a case for higher public spending*. Springer, 2000.
- Martha C Nussbaum. Creating capabilities: The human development approach and its implementation. *Hypatia*, 24(3):211–215, 2009.

- Francesco Passarelli and Guido Tabellini. Emotions and Political Unrest. *Journal of Political Economy*, 125(3):903–946, 6 2017.
- Bezalel Peleg and Menahem E Yaari. On the existence of a consistent course of action when tastes are changing. *The Review of Economic Studies*, 40(3): 391–401, 1973.
- Francisco Perez-Arce. The effect of education on time preferences. *Economics of Education Review*, 56:52–64, 2017.
- Torsten Persson and Guido Enrico Tabellini. *Political economics : explaining economic policy*. MIT Press, 2000.
- Aude Pommeret and Katheline Schubert. Energy transition with variable and intermittent renewable electricity generation. 2019.
- Fabien Prieur and Thierry Bréchet. Can education be good for both growth and the environment? *Macroeconomic Dynamics*, 17(5):1135, 2013.
- Anja Prummer and Jan-Peter Siedlarek. Community leaders and the preservation of cultural traits. *Journal of Economic Theory*, 168:143–176, 3 2017.
- Robert A Ritz. Price discrimination and limits to arbitrage: An analysis of global lng markets. *Energy Economics*, 45:324–332, 2014.
- Nicholas Rivers and Brandon Schaufele. Salience of carbon taxes in the gasoline market. *Journal of Environmental Economics and management*, 74:23–36, 2015.
- Amanda Lea Robinson. National versus ethnic identification in africa: Modernization, colonial legacy, and the origins of territorial nationalism. *World Politics*, 66(4):709–746, 2014.
- Arthur J Robson and Larry Samuelson. The evolution of time preference with aggregate uncertainty. *American Economic Review*, 99(5):1925–53, 2009.
- Stephen W Salant. Exhaustible resources and industrial structure: A nash-cournot approach to the world oil market. *Journal of Political Economy*, 84 (5):1079–1093, 1976.
- Nicholas Sambanis and Moses Shayo. Social identification and ethnic conflict. *American Political Science Review*, 107(2):294–325, 2013.

- Nicholas Sambanis, Stergios Skaperdas, and William C Wohlforth. Nation-building through war. *American Political Science Review*, 109(2):279–296, 2015.
- Wolf-Peter Schill and Claudia Kemfert. Modeling strategic electricity storage: the case of pumped hydro storage in germany. *The Energy Journal*, 32(3), 2011.
- Richard Schmalensee. On the efficiency of competitive energy storage. *Available at SSRN 3405058*, 2019.
- Tim Schwanen, David Banister, and Jillian Anable. Rethinking habits and their role in behaviour change: the case of low-carbon mobility. *Journal of transport geography*, 24:522–532, 2012.
- Dusko Sekulic, Garth Massey, and Randy Hodson. Who Were the Yugoslavs? Failed Sources of a Common Identity in the Former Yugoslavia. *American Sociological Review*, 59(1):83, 2 1994.
- Amartya Sen. Utilitarianism and welfarism. *The journal of Philosophy*, 76(9): 463–489, 1979.
- Amartya Sen et al. Commodities and capabilities. *OUP Catalogue*, 1999.
- Anton A Shardin and Michaela Szölgyenyi. Optimal control of an energy storage facility under a changing economic environment and partial information. *International Journal of Theoretical and Applied Finance*, 19(04):1650026, 2016.
- Moses Shayo. A model of social identity with an application to political economy: Nation, class, and redistribution. *American Political science review*, 103(2):147–174, 2009.
- Ramteen Sioshansi. When energy storage reduces social welfare. *Energy Economics*, 41:106–116, 2014.
- Anthony D Smith. National identity and the idea of european unity. *International affairs*, 68(1):55–76, 1992.
- Michael Spence and David Starrett. Most Rapid Approach Paths in Accumulation Problems. *International Economic Review*, 16(2):388, 6 1975.

- Enrico Spolaore. What is european integration really about? a political guide for economists. *Journal of Economic Perspectives*, 27(3):125–44, 2013.
- Bjarne Steffen and Christoph Weber. Optimal operation of pumped-hydro storage plants with continuous time-varying power prices. *European Journal of Operational Research*, 252(1):308–321, 2016.
- Alfred Stepan, Juan J Linz, and Yogendra Yadav. *Crafting state-nations: India and other multinational democracies*. JHU Press, 2011.
- Cass R. Sunstein. Endogenous preferences, environmental law. *Journal of Legal Studies*, pages 217–235, 1993.
- Guido Tabellini. The scope of cooperation: Values and incentives. *The Quarterly Journal of Economics*, 123(3):905–950, 2008.
- Richard H Thaler and Cass R Sunstein. Libertarian paternalism. *American economic review*, 93(2):175–179, 2003.
- Henry Thille. Inventories, market structure, and price volatility. *Journal of Economic Dynamics and Control*, 30(7):1081–1104, 2006.
- Thierry Verdier and Yves Zenou. Cultural leader and the dynamics of assimilation. *Journal of Economic Theory*, 175:374–414, 2018.
- Nico Voigtländer and Hans-Joachim Voth. Nazi indoctrination and anti-semitic beliefs in germany. *Proceedings of the National Academy of Sciences*, 112(26):7931–7936, 2015.
- Carl Christian Von Weizsäcker. Notes on endogenous change of tastes. *Journal of Economic Theory*, 3(4):345–372, 1971.
- Carl Christian von Weizsäcker. The welfare economics of adaptive preferences. *MPI Collective Goods Preprint*, (2005/11), 2005.
- Anke Weber. The causes of politicization of ethnicity-a comparative case study of kenya and tanzania. 2010.
- Eugen Weber. *Peasants into Frenchmen : the modernization of rural France, 1870-1914*. Stanford University Press, 1976.
- Rachel Weinberger and Frank Goetzke. Unpacking preference: How previous experience affects auto ownership in the united states. *Urban studies*, 47(10):2111–2128, 2010.

Matthew Weinzierl. The promise of positive optimal taxation: normative diversity and a role for equal sacrifice. *Journal of public Economics*, 118: 128–142, 2014.

Jeffrey C Williams and Brian D Wright. *Storage and commodity markets*. Cambridge university press, 2005.

Andreas Wimmer and Yuval Feinstein. The Rise of the Nation-State across the World, 1816 to 2001. *American Sociological Review*, 75(5):764–790, 10 2010.

Brian D Wright and Jeffrey C Williams. The welfare effects of the introduction of storage. *The Quarterly Journal of Economics*, 99(1):169–192, 1984.

Menahem E Yaari. Endogenous changes in tastes: A philosophical discussion. *Erkenntnis*, 11(1):157–196, 1977.

Yuting Yang. Electricity interconnection with intermittent renewables. 2020.