STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
5350 Master's thesis in economics
Academic year 2020-2021

The curious case of the Russo-Georgian war:
economic implications of a post-Soviet conflict

Edward Lilliehöök (23396)

Abstract. The purpose of this paper was to fill a gap in the literature on the consequences of the Russo-Georgian war of 2008. I asked: how did Russia's invasion of Georgia affect Georgian per capita GDP? I hypothesised that the effect was negative but negligible and used the synthetic control method to answer the research question. Using panel data on former eastern-bloc republics, Georgia's synthetic counterfactual was constructed. The results show that the war had a positive, causal effect on the outcome variable. By looking at individual years, the estimated effect is found to be significant at the 5 % level in 2012 and 2013. When considering the entire post-treatment period, an estimated accumulated GDP per capita effect of $4314 is found to be robust and significant at the 5 % level. This translates to an average post-war benefit of about 7.5 % per annum. The results are likely driven by the inflow of foreign aid that followed the war. Previous studies on post-war economies support this idea. The results are not free from internal validity concerns. While biases do not seem to override my overall conclusion, there are outstanding questions relating to the channels driving the effect and to the results' sensitivity to the choice of baseline donor pool.

Keywords: case study, Georgia, Russia, synthetic control method, war

JEL: C21, F34, F35, F51, F55, O19

Supervisor: Anders Olofsgård
Date submitted: 16 May 2021
Date examined: 28 May 2021
Discussants: Pratanu Mitra
Examiner: Magnus Johannesson

ACKNOWLEDGEMENTS

Tack till Anders Olofsgård och Joel Blanke.

Stadsholmen, 16 maj 2021

TABLE OF CONTENT

1 INTRODUCTION 1
2 BACKGROUND 4
2.1 Historical context and an account of the war 4
2.2 The war's aftermath 6
3 LITERATURE REVIEW 9
3.1 Destruction of capital 10
3.2 Policy and political institutions 11
3.3 (Re)allocation of resources 13
3.4 Synthetic control method literature 14
4 THEORETICAL FRAMEWORK 15
4.1 Research question and hypothesis 17
5 EMPIRICAL SPECIFICATION 17
5.1 Methodological approach 17
5.1.1 The synthetic control model 20
5.1.2 Inference and robustness checks with SCM 24
5.2 Data collection 26
5.2.1 Donor pool 26
5.2.2 Selection of variables 27
5.2.3 Time period 28
5.3 Implementation 28
5.3.1 Model specification 28
5.3.2 Donor pool 31
6 RESULTS 34
6.1 The causal effect of the Russo-Georgian war 34
6.1.1 Inference 37
6.1.2 Robustness checks 39
7 DISCUSSION 43
7.1 Internal validity 47
7.2 External validity 49
7.3 Recommendations 50
8 CONCLUSION 50
9 BIBLIOGRAPHY 52
APPENDIX 57
Appendix A Map of Georgia covering Russia's offensive in 2008 57
Appendix B Choosing country and variable weights 57
Appendix C Calculating p-values 59
Appendix D Criticism of inference methods 61
Appendix E Per capita GDP graphs 62
Appendix F In-time placebo tests 63
Appendix G Benchmark model output 64

1 INTRODUCTION

Russia's invasion of Georgia surprised many. The idea of having European borders moved by force was considered a thing of the past. This impression had not yet changed by 2008, after less than a decade of President Putin being at Russia's helm. And even though the war redrew borders and challenged the European security order, it came to have relatively minor effects on how the West related to Russia in the following years. Given the gravity of this breach of international law, little was done by the international community, and relations between the West and Russia were soon back to normal.

Today, the war of 2008 is widely held to be one of the earliest and clearest signs of the Kremlin's ambition to reappear as a major power. However, before the 2014 annexation of Crimea and the 2015 launch of a military operation in support of Syria's Assad regime, the Russo-Georgian war's part in a wider trend was not evident. Russia's invasion of Georgia in 2008 was largely forgotten, or rather happily ignored, for several years before the aggression against Ukraine came as a brutal reminder of the pre-eminence some states attribute to *strength* before *right* in international relations.

In studying the war between Russia and Georgia, it becomes clear that not only the destructive consequences most associated with war are of interest when measuring outcomes. The short, five-day war did not cause massive physical destruction. Georgia was no Chechnya. Tbilisi was no Grozny. What the war did do was to cause the displacement of hundreds of thousands of people and change Georgia's security situation. Additionally, international donors provided vast financial support in Georgia's post-war recovery while Georgia was integrated deeper into a political and economic environment conducive to liberal democracy and economic growth. Due to these opposing forces, the net economic

impact of the war is unclear. This ambiguity makes it interesting to estimate the war's aggregate economic impact.

To my knowledge, the aggregate causal impact of the Russo-Georgian war has not been measured before. That the Russo-Georgian war was short, limited, and soon overlooked by the international community may be reasons why there has been little research on its economic consequences. This is regrettable, especially if one considers the prophetic, political importance later ascribed to the conflict. The purpose of this paper is to fill this gap in the literature on the Russo-Georgian war and thereby also contribute to the wider understanding of the relationship between conflicts and economy.

To fulfil this paper's purpose, the causal effect of the conflict is to be quantified. This requires knowledge of Georgia's counterfactual economic development. However, it is often difficult to find a valid comparison unit in case studies at the aggregate level because a comparison of countries is likely to suffer from endogeneity. Countries are seldom similar enough to be used as each other's counterfactual. If comparison units are bad, differences to outcomes may result from differences to their characteristics rather than from a treatment effect. In this paper, this problem is addressed by using the synthetic control method (SCM). Through SCM, the per capita gross domestic product (GDP) trend that Georgia would have followed in absence of the war is replicated, allowing for a comparison between Georgia and its counterfactual.

This paper finds that the war seems to have caused an immediate economic dip when it broke out in 2008. After 2008, Georgia surpasses its counterfactual, which implies a positive causal effect of the war in the remainder of the sampled period. When looking at effects for each *individual* post-treatment year, the estimated effects are statistically

significant in 2012-2013. When looking at the *entire* post-treatment period, the effect is found to be positive and significant at the 5 % level. It amounts to an accumulated increase to per capita GDP of $4314, an average benefit of about 7.5 % per year, in the period 2008-2013. The results are robust to placebo tests and sensitivity checks. While the sign and size of the per capita GDP effects of the war are surprising, the context in which the war took place provides explanation. The causal channels are not identified by this paper, but it is believed that the costs of the war are mitigated by benefits derived from Georgia's close westward relations. The results relate to previous research on the relationship between economy and war as they testify to the importance of seeing a conflict in its wider political and economic context. This case study contributes to the understanding of the war itself while adding to the literature on the relationship between economy and conflict.

This section is followed by section two, which provides an account of the Russo-Georgian war. Section three contains a review of the literature on war and economy. Section four deals with the theoretical framework supporting the research hypothesis and question. Section five provides a description of the origins of SCM, discusses the data used, and describes how the method and data are operationalised to answer the research question. The fifth section also deals with questions on inference and robustness. Section six presents the results from my synthetic control model. The model is tested by running robustness checks before the results are discussed in section seven. Section eight concludes. I use the words 'treatment', 'intervention', and 'the war/conflict' interchangeably. Also, the word 'unit' is used interchangeably with 'country'. 'Unit of interest' is the same as 'treated unit', while the terms 'untreated units' and 'donor units' are synonymous.

2 BACKGROUND

2.1 Historical context and an account of the war

The dissolution of the Soviet Union (USSR) led to the emancipation of Georgia and other Soviet republics. The 1990s were turbulent for many of these newly formed, sovereign states. Georgia belongs to those in the post-Soviet sphere that were most severely affected by political and economic turmoil. In the years immediately following independence, the Georgian economy contracted by 65 %, an exceptionally large drop even when compared to its peers. Georgia's situation stabilised in 1995, when economic reforms and a new currency were introduced (World Bank, 2018).

In the 1990s, Georgia not only experienced economic problems, but also border disputes and internal conflicts that had gained momentum in the latter days of the USSR. The regions of South Ossetia and Abkhazia challenged Georgian supremacy, causing civil war between secessionist groups and the Georgian government. These internal conflicts ended in 1993 with an agreement that positioned Russia as a guarantor of security in the two breakaway regions. Hundreds of Russian so-called 'peacekeeping' troops were stationed in South Ossetia and up to 3,000 in Abkhazia. Through its trip-wire presence in the secessionist regions, Russia cemented its influence over these regions (Kofman, 2018).

While the Georgian government retained control over some villages and areas of South Ossetia and Abkhazia, the internal conflicts led to the displacement of people within Georgia as well as the de-facto establishment of self-governing entities within Georgia's internationally recognised borders. The conflicts froze and remained frozen during Eduard Shevardnadze's presidency, but rekindled after the ousting of Shevardnadze and the election of Mikheil Saakashvili following Georgia's bloodless 'Rose Revolution' of 2003.

Saakashvili launched vast economic reforms and anti-corruption campaigns (Reuters, 2009; World Bank, 2018). From the beginning of his presidency, Saakashvili had ambitious plans to deepen Georgia's Western integration politically and economically (Khidasheli, 2011). With the policies implemented following the Rose Revolution, Georgia became a star reformer in the post-Soviet space. Starting in 2006, Georgia 'unilaterally and radically' liberalised its foreign trade policies and exposed itself to international competition (Emerson and Kovziridze, 2016). In addition to economic reforms, Saakashvili promised the re-establishment of territorial control within Georgia's borders. These efforts included a reform to modernise and expand the Georgian armed forces. Georgia's defence spending soared from roughly 1 % of GDP in 2003 to more than 9 % by 2007 (Kofman, 2018).

Meanwhile, the Kremlin saw the Rose Revolution as a vehicle, and Saakashvili as an agent, of Western attempts at encroaching on Russian interests in its perceived sphere of influence. Tensions rose between Tbilisi and Moscow and led to the Kremlin dealing a blow to Tbilisi by imposing an embargo which restricted Georgia's ability to trade with a major trading partner (Reuters, 2009; World Bank, 2018). Georgia continued its military reforms and made attempts to retake control over strategic parts of the breakaway areas. In early 2008, tensions increased as Russia shot down Georgian drones and reinforced its peacekeeping forces on Georgian territory with units of special forces (Kofman, 2018).

On 6 August, Georgian forces and South Ossetian militias engaged in intense fighting. Georgia started to amass forces on the de-facto border to South Ossetia and later that same day, Georgian artillery shelled South Ossetian positions in the city of Tskhinvali. Soon thereafter, Russian troops started to move across the border of Georgia proper. On 8 August, Georgian units clashed with Russian regulars. While Russia's 58th army was

moving into South Ossetia, Russia dropped airborne troops deep into Georgia and opened a second front in Abkhazia. Georgian forces were soon overwhelmed (Kofman, 2018).

The European Union's (EU) fact-finding mission to Georgia found that while Georgia's shelling of Tskhinvali sparked the large-scale armed conflict, it was merely the 'culminating point' of a longer period of rising tension and provocations. Russia had by then already deployed non-peacekeeping troops to the disputed regions and had also granted citizenship *en masse* to people in Abkhazia and South Ossetia (Independent International Fact-Finding Mission on the Conflict in Georgia, 2009). Kofman (2018) writes that Russia actively sought the war and was attempting at luring Saakashvili into making a first move that would justify a Russian full-scale military operation.

The war lasted for five days and resulted in 412 Georgian servicemen and civilians dead or missing, according to Georgian authorities. Russia in turn reports the loss of 64 servicemen and 162 South Ossetian civilians (Reuters, 2009). Another 192,000 people became displaced after having fled their homes due to the conflict (UN News, 2008). After having crossed the Georgian border with close to 20,000 troops, Russia was able to advance far towards the Georgian capital of Tbilisi before a ceasefire was negotiated under heavy international pressure from the EU and USA (World Bank, 2018; Vendil Pallin and Westerlund, 2009). Appendix A contains a map detailing Russia's offensive.

2.2 The war's aftermath

There is no doubt that the strategic consequences of the war were severe for Georgia: a fifth of Georgia's territory ended up under Russian occupation, Russian forces came to be based merely 30 km from the Georgian capital, and infrastructure was destroyed (Hamilton, 2018). What makes the case of the Russo-Georgian war curious is that its

economic consequences are not as clear. There are several contradictory and ambiguous forces in play that likely affected the economy.

First, the war did cause destruction to Georgian infrastructure and to the environment, but its destructiveness was not large (United Nations and World Bank, 2010). The war had economic implications also through lives lost, destruction of military materiel, and more. However, considering the limited extent and duration of the war, the effects of these factors on the macro level would not seem to be very large. What would seem to matter more is the war adding (in the short term) another 127,000 internally displaced people (IDP) to Georgia's 223,000 pre-existing IDPs (UN, 2008).

The direct consequences of the war can be believed to have adversely affected the economy, but there are other, opposing forces. Indirectly, the war resulted in vast international assistance to Georgia in its post-war recovery. In October 2008, the EU together with the World Bank arranged a donors conference in support of Georgia (Brunnstrom, 2008; Gogolashvili, 2017). The conference led to donors pledging $4.5bn, including $2.5bn in concessional loans and $2bn in grants. Thus, more than a third of Georgian GDP in 2008 was promised in support of Georgia's post-war recovery. Additionally, the International Monetary Fund (IMF) made available $750m in a stand-by arrangement (International Monetary Fund, 2008; Otarashvili, 2013; World Bank, no date). The case has been made that Georgia was in fact helped by the war having taken place right before the financial crisis of 2008 (Otarashvili, 2013). In 2008, Georgia had a small external debt, but its debt was issued at junk bond rates, making market financed loans expensive (Trading Economics, no date). With financial support from international donors, the Georgian government had the means to present large and inexpensively financed fiscal stimulus packages to mitigate the economic damage of the two crises of

2008 (United Nations and World Bank, 2010). The Georgian government increased its spending by about one third between 2007 and 2010 and by another quarter between 2010 and 2012. While foreign investments dropped after 2008, in 2010 they started growing again in parallel with tourism and public investments that together contributed to economic growth (Otarashvili, 2013).

Politically in Georgia, the effects were not very dramatic. It is often the case that conflict and turmoil result in increased government repression in the name of national security. The opposite happened in Georgia after the war. Before the invasion, Georgia's democracy weakened due to President Saakashvili's increasingly repressive policies, but Georgia's democratic corrosion was not exacerbated by the war of 2008 (Hamilton, 2018). On the contrary, Saakashvili *did not* act to prevent the formation and growth of a strong opposition in the post-war years. The president and his party were eventually outvoted in the free and fair 2012 and 2013 elections (Hamilton, 2018). Hence, the war does not at first glance seem to have dramatically affected Georgia's economy through changes to its political institutions.

Conversely, the foreign policy pursued post-war may have had economic implications for Georgia. Following the war, EU-Georgian relations entered a new phase as Georgia shifted its focus from trans-Atlantic relations and North Atlantic Treaty Organization (NATO) membership to integration with the EU (Khidasheli, 2011). Due to the war, a future NATO membership was rendered unlikely, making the EU Georgia's best alternative (ibid). Georgia's pivot to the EU was met by a reciprocal European interest. The war spurred on the development of the EU's Eastern Partnership (EaP), which was launched in 2009 and works towards the European integration of former Soviet republics. Less than a month after the war, the European Council decided on beginning negotiations of a free

trade area and a visa-free regime with Georgia (Gogolashvili, 2017). The war's consequences for Georgia's foreign policy alignment may thus have put Georgia in a better position to grow in the post-war years.

Another factor affecting the Georgian economy in the post-war period is the financial crisis, which may confound attempts at singling out the impact of the war. Russia's invasion of Georgia took place in the context of Georgia having undergone far-reaching economic reforms ignited by the Rose Revolution. It has been argued that Georgia's reforms aimed at diversifying its economy and exposing itself to international competition helped Georgia overcome the perils of the conflict with Russia as well as the financial crisis (Emerson and Kovziridze, 2016).

In all, there is much uncertainty as to what the aggregate consequences of the Russo-Georgian war were. There are several possible, opposing channels through which the Georgian economy might have been affected. Knowledge of the net economic impact of this specific conflict would help raise our general understanding of the relationship between war and economy and of the role of politics before and after a war.

3 LITERATURE REVIEW

It is difficult to estimate the costs of war. Overall, one does not always know where a war's costs arise and how these are to be quantified. It is obvious that a war or conflict can have socially, psychologically, and physically destructive effects on the individuals involved. There is, however, no consensus as to what is to be generally expected from a war in terms of aggregate economic outcomes. Effects of conflict may be both direct and indirect and separation into the two categories is not always easy. To estimate the costs of rebuilding a destroyed bridge is not difficult, but it is hard to measure the costs of trade foregone due

to that bridge's initial destruction. Similarly, the effects of war on a directly involved state are easier to estimate than the costs caused by spill-overs on neighbouring states.

In the literature, I identify three main channels through which wars affect economies: physical destruction in terms of infrastructure and lives, changes to policy and political institutions, and diverted investments. In addition, there may be spill-over effects on neighbouring states. It has been found that a conflict's indirect effect on the economic growth of neighbouring states is negative, but it depends on the definition of what constitutes a 'neighbouring' state (see Murdoch and Sandler, 2004; De Groot, 2010). The literature review below focuses on how a conflict affects a directly involved country.

3.1 Destruction of capital

The most evident direct consequence of conflict is physical destruction, eg the destruction of critical infrastructure such as energy grids, airports, harbours, and bridges. The consequences for an economy are clear, as disruptions to energy supplies or to the flow of goods and services severely limit an economy's ability to function. Likewise, the deaths of citizens incur human capital-related costs to an economy.

In the literature, different tools are employed to estimate these direct costs of war. One method is cost accounting, which measures and adds the costs of resources destroyed (see, for instance, Nordhaus (2002), an ex-ante analysis of the costs of the invasion of Iraq). The problem with the cost accounting method is that results are highly dependent on the difficult choice of what costs to include. When using cost accounting methods, there is also a risk of double counting biasing results (Brück, De Groot and Bozzoli, 2012). An alternative to cost accounting is counterfactual analysis, ie to compare a country having

undergone conflict to its counterfactual (Brück, De Groot and Bozzoli, 2012). More on this later.

Organski and Kugler (1977) relate post-war performance to the destruction caused to a country's industry. They condition the scale of a war's impact on whether a country was on the winning or losing side. A loser is assumed to have suffered from more destruction and is hence prone to grow faster postbellum. While defining 'power' as a country's gross national product, Organski and Kugler use time-series analysis and find that major wars affect the power distribution between states. States on the losing side of a conflict suffer large losses in the short term, while the power levels of winners and neutral states are only marginally affected. Along the lines of Solow's neoclassical growth theory, Organski and Kugler find that the effects of war dissipate with time as economies rebuild and catch up. When the losers of the war recover from the conflict – for instance by adopting new technologies in the wake of the war – the system's power balance converges to the level that would have been expected had the conflict never occurred. The authors refer to this effect as 'the phoenix factor' (Organski and Kugler, 1977). While the existence of such a phoenix factor may be reassuring, others have noted that this idea of convergence does not consider that before trends converge, an economy is still adversely affected by performing subpar for a number of time periods (Brück, De Groot and Bozzoli, 2012). Naturally, the convergence-related theories on increased returns to physical capital is conditional on there being human capital. In other words, this will not be the case if the war caused too many casualties and vast migration.

3.2 Policy and political institutions

Regarding the economic effects of war that arise not from physical destruction but from political change, there are interesting contributions to the literature. Like Organski and

Kugler (1977), Olson (1982) finds that war relates positively to economic performance. However, unlike Organski and Kugler, Olson conditions a war's impact on economic performance on how the war affects a country's political environment. Olson argues that an economy can end up benefitting from a conflict in the longer term if the war causes creative political destruction. If politicians, policies, and vested interests that have previously hampered economic development are replaced because of the war, then an economy can benefit from a better economic and political development (Olson, 1982).

Another plausible effect of looming conflict or direct hostilities is political instability. It has been found that political instability, measured in terms of variation to political freedoms, has detrimental effects on physical capital accumulation. Hence, if conflict is related to political instability, an economy might be affected by the war through its effect on investors' willingness to invest (Feng, 2001).

In the literature, the dynamic effects of conflicts have been found to be conditional on its length and severity. Koubi (2005) employs cross-sectional data and a regression model to measure the effect on economic performance of intra-state wars. Koubi (2005) acknowledges that the war affects the economy through a number of channels and refer to the destruction channel discussed by Organski and Kugler (1977) as well as to the policy channel discussed by researchers like Olson (1982) and Feng (2001). Koubi's main contribution is on the temporal aspects to the costs of war. Koubi finds evidence in support of a causal 'peace dividend' effect: 'the longer or more severe the war, the higher the subsequent long-term rate of economic growth' (2005, 67).

There are, however, contributions to the literature suggesting that intertemporal effects are more ambiguous. In the post-war environment, with its lack of revenue and need to restore

infrastructure, Collier and Hoeffler (2004) note that aid is especially productive. The authors find that extraordinary growth postbellum is not naturally occurring, but rather caused by an interaction between aid and policies and institutions. A typical post-conflict economy experiences an additional growth of two percentage points per annum. The size of the growth spurt is conditional on institutions and the policies pursued after the end of a conflict and is due to aid (Collier and Dollar, 2002; Collier and Hoeffler, 2004). The returns to aid are, however, diminishing (Collier and Dollar, 2002). Collier and Hoeffler (2004) find that the absorptive capacity of aid of an economy, and hence the contribution of aid to super-normal growth, is greatest a few years after peace is restored.

3.3 (Re)allocation of resources

How a conflict affects the economy by changes to investment patterns is another ambiguous issue. Defence industry investments may give rise to new, dual-use technologies that in the long-term spur economic growth. In addition, an increase to government spending through higher military expenditure increases an economy's overall output. However, as military expenditure and war related investments go up, these investments may have a crowding-out effect that is detrimental to the economy (see papers by Benoit, 1978; Dunne and Vougas, 1999; and Cappelen, Gleditsch and Bjerkholt, 1984). Hence, resources that could have contributed to economic performance elsewhere by increasing productivity and promoting growth is, due to war or the threat of war invested in less productive ways (Koubi, 2005).

Investments may also be affected by changes to the perceived risk environment. In his study of civil wars, Collier (1999) finds that war causes a gradual decline of the capital stock due to lower savings, reallocation of investments to foreign markets, and destruction. Collier suggests that even after peace has been restored, the risk characteristics of the post-

war economy are not the same as those antebellum (ibid). Like Koubi (2005), Collier finds that the dynamic effects of war are conditional on the war's characteristics. In case of a short war, the slow-changing capital stock is likely to still be above the desired level post-war, which causes a continuing economic decline. A sufficiently long conflict, however, will result in a post-conflict economy's capital stock being below the desired level postbellum and lead to growth. Collier's empirical findings suggest that such a peace dividend of a long war is large (1999).

3.4 Synthetic control method literature

With the synthetic control method, treatment effects are estimated by constructing a synthetic comparison unit to be used as counterfactual. By design, the synthetic control method estimates an aggregate effect of conflict. Hence, the channels through which the war affects Georgia are not detected by an SCM model (Cavallo et al, 2013).

Since its introduction, the SCM method has been used in comparative case studies to evaluate the costs of conflict. Abadie and Gardeazabal (2003) first apply the method to estimate the impact of terrorism on the Basque Country of Spain. They find that the intensity of terrorist activities had a negative causal effect on Basque per capita GDP. The robustness of their results is tested by running a synthetic control model on a similar Spanish region. Similarly, Echevarría and García-Enríquez (2020) find that the Arab spring and the political turmoil and conflict associated with it led to welfare losses in Egypt. The Basque conflict and the revolution in Egypt both are internal conflicts characterised by killings and repression. Judging by the findings of the authors, the start of these two conflicts seem to have caused an immediate dip in economic welfare, leaving the counterfactual units and the observed units at different levels but with about parallel growth trajectories post-treatment (Abadie and Gardeazabal, 2003; Echevarría and García-

Enríquez, 2020). The results relate to those of Feng (2001) and Collier (1999), as a deterioration of the political environment plausibly scared off investors.

Unlike the studies mentioned above, Bluszcz and Valente (2020) research a large-scale inter-state conflict. The authors find a causal negative effect of Russia's aggression against eastern Ukraine on Ukrainian per capita GDP. Similarly, Echevarría and García-Enríquez (2019) consider the Libyan revolution, which resulted in a prolonged and bloody civil war. Their findings indicate a sharp and sudden decline in Libyan GDP per capita compared to its counterfactual. Thereafter, the Libyan economy has been highly volatile. Any clear sign of continued divergence or convergence is not seen. Both Ukraine and Libya are still undergoing treatment, which makes the findings difficult to relate to the literature.

4 THEORETICAL FRAMEWORK

The three identified channels through which war affects an economy can be applied to the Georgian case. Like the literature, the relationship between Russia's invasion and Georgia's economic performance is ambiguous.

First, considering a war's destruction of capital, the research of Organski and Kugler (1977) suggests that the Russo-Georgian war caused Georgia to lose GNP in the short term to then converge to a pre-war state. Koubi (2005) in turn suggest that the war's severity matters for its effect on Georgia's economy. The Russo-Georgian war was intense but very short and thus limited in its destructiveness. By the findings of Koubi (2005), this suggests that the Russo-Georgian war's effect on Georgian post-war economic performance was limited.

Second, how the war affected the economy through policy. The Russo-Georgian war's economic impact is neither by Olson's (1982) nor Feng's (2001) research obvious, because there are no dramatic consequences of the war on vested political interests or political repression in Georgia. The reforms launched by Saakashvili in the mid-2000s continued after the war. Saakashvili did not increase repression and his party were peacefully voted out of office in 2012 and 2013. Still, Georgia moved westwards politically after the war. Domestically, Georgia made slight improvements to its Freedom House index score and raised its polity score, which implies a more democratic development (Hamilton, 2018). In addition, by the findings of Collier and Hoeffler (2004), the soundness (and continued strengthening) of Georgian institutions coupled with effective policies and international aid together point toward a positive economic post-war development.

Third, findings by Collier (1999) would suggest that this short war had lingering, negative effects on the Georgian economy due to a deteriorated risk environment causing the desired level of capital to be lower than the pre-war levels. In the case of Georgia, FDIs indeed went down after the war. They soon started growing again in 2010, but then from a lower level than pre-war.

As discussed above, a war can be believed to influence economic performance through several channels. When measuring the aggregate effect of the conflict using SCM, these channels and their effects cannot be disentangled from one another. This is a limitation to this paper. The consequences of war will therefore include those resulting directly from the destruction of infrastructure and deaths as well as those caused by the reallocation of resources, the economic costs of displaced people, the potentially mitigating effects of international support given to Georgia, and by political changes. Other costs of war, such as social costs, are not measured in this paper, unless they are reflected in GDP levels.

Rather than considering the indirect consequences of the war as confounders of the 'true' (direct) impact of the war, a general equilibrium outcome of the war is to be found. The consequences of the conflict are thus seen against the backdrop of Georgia's wider (geo)political and economic context.

4.1 Research question and hypothesis

This paper aims at answering the question:

How did Russia's 2008 invasion of Georgia affect Georgian per capita GDP?

This research question connects to the paper's purpose of measuring the economic consequences of the conflict and thereby contribute to the literature on economy and war. Based on the theoretical framework, the first and third channels identified in section three suggest a negative but limited economic impact of the war while the second channel would indicate a positive effect post 2008. Weighing these together, this paper follows the hypothesis that the Russo-Georgian war's overall effect on the Georgian economy was negative but negligible. While an immediate decline is thought to have occurred due to the war's direct consequences (eg IDPs and physical destruction), Georgia is believed to then have converged to its pre-treatment trajectory helped, eg, by international aid. Georgia is thus believed to perform subpar for some periods before there is convergence, making the overall effect in the studied period negative. The potentially detrimental and lingering effects of the war on investments are believed to have been mitigated by foreign aid and Georgia's economic reform pace.

5 EMPIRICAL SPECIFICATION

5.1 Methodological approach

The research question of this paper is answered by applying SCM to the Russo-Georgian war. SCM has been described as '[a]rguably the most important innovation in the policy

evaluation literature in the last 15 years' (Athey and Imbens 2017, 9). SCM was first introduced by Abadie and Gardeazabal (2003). Since then, it has been further developed by Abadie, Diamond and Hainmueller (2010) who develop inference techniques for the SCM method in their study on Californian tobacco taxes. In addition to papers already mentioned, SCM has been applied to measure the effects of Germany's reunification (Abadie, Diamond and Hainmueller, 2015), natural disasters (Cavallo et al, 2013), economic reform (Adhikari et al, 2018; Billmeier and Nannicini, 2013), Georgia's Rose Revolution political and market reforms (Lawson, Grier and Absher, 2019), and the effect of government formation deadlocks on economic growth (Albalate and Bel, 2019), to mention a few.

When conducting studies using population data, the uncertainties relating to the true population values are removed. However, there remains an uncertainty as to whether the aggregate unit chosen for comparison offers a valid counterfactual. Before SCM, in comparative case studies at the aggregate level, control units could be arbitrarily chosen based on geographic proximity to the treated unit or for other subjective reasons. However, to choose a control unit based on subjective reasons is a bad method of selection if countries vary greatly in political or economic characteristics (McClelland and Gault, 2017). For instance, Norway's oil wealth makes it a very different from Sweden, regardless of the two countries' proximity. Traditional ways of finding a comparison unit are thus suffering from omitted variable biases. These problems can be mitigated by using SCM.

The SCM method follows from the potential outcome approach proposed by Rubin (1974) and provides a data-driven method of selecting comparisons units in comparative studies (Abadie, Diamond and Hainmueller, 2010; Adhikari et al, 2018). By constructing a counterfactual through the *combination* of control units, the characteristics of the unit of

interest are easier to reproduce (Abadie, Diamond and Hainmueller, 2015). It is a method well-suited for case studies on a single country (Gardeazabal, 2012). The synthetic unit is chosen so that it matches the pre-treatment trajectory and economic characteristics of the treated unit. If the synthetic unit tracks the observed unit well, any divergence occurring after the intervention presumably reflects a treatment effect (McClelland and Gault, 2017).

There are several additional benefits to using SCM. First, with SCM, a counterfactual unit is created as a linear combination of available control units that are similar to the treated unit (the 'donor pool'). In this weighted average of donors, weights are restricted to being positive and sum to one. This way, data is not extrapolated to create a potentially extreme counterfactual not supported by the available data. When extrapolating, one relies heavily on the model used. Since a model is seldom perfect, extrapolating beyond the convex hull of available control units may lead to biases (King and Zeng, 2006). Second, SCM is applicable to small samples while still allowing for inference, because of its data-driven approach and because of how the test statistic is defined (see section 5.1.2). Third, the SCM method shows graphically how well the counterfactual replicates the unit of interest while also presenting the constitution of the counterfactual unit. Unlike regression analysis, where all control units contribute to the regression estimates, SCM provides the weights attributed to donor units and thus allows for qualitative and quantitative analyses of the origins and direction of potential biases (Adhikari et al, 2018; Abadie, 2020). Lastly, SCM does not require as stark identifying assumptions as difference-in-difference or panel data methods. When applying these methods, there are often concerns relating to non-parallel trends and omitted variable biases. It is shown in section 5.1.1. that SCM allows for time-varying unobserved factors when the pre-treatment trajectory of the treated unit is closely tracked by the synthetic control unit and there are enough pre-treatment periods (Abadie

and Gardeazabal, 2003; Abadie, Diamond and Hainmueller, 2010; McClelland and Gault, 2017; Abadie, 2020).

To obtain credible estimates of the treatment effect, SCM requires the following assumptions:[1]

- no other unit in the donor pool can have undergone a similar treatment in the sampling period;
- there cannot be spill-overs of the treatment affecting the donor pool units (ie, no violation of the stable unit treatment assumption, SUTVA);
- the predictor values of the treated unit must be replicable by a linear combination of those of the donor units;
- the treatment should not have an effect prior to implementation (if there is a risk of anticipation effects, then the treatment period could be moved); and
- the predictor variables must have an approximate linear relationship with the outcome (Abadie, Diamond and Hainmueller, 2010; McClelland and Gault, 2017).

5.1.1 *The synthetic control model*[2]

Suppose a sample of $J + 1$ units (for instance, aggregate entities such as countries) indexed by j. Without loss of generality, $j = 1$ can be defined as the treated unit. The remaining units $j = 2$ to unit $j = J + 1$ are possible control units that together form the donor pool.

Further, suppose that there is a balanced panel data set including both the treated unit and control units for the time periods $t = 1, \ldots, T$. Among these T time periods, there are T_0

[1] For a detailed discussion on these identifying assumption, I recommended reading Abadie (2020)
[2] This section is to a large extent based on Abadie, Diamond and Hainmueller (2010), Abadie (2020) and Echevarría and García-Enríquez (2020)

pre-treatment periods ($t = 1, \dots, T_0$) and T_1 post-treatment periods ($t = T_0 + 1, \dots, T$). Hence, $T = T_0 + T_1$. The first period in which the unit is subjected to the treatment is $T_0 + 1$. Before then, the treatment is assumed to have no effect on the treated unit. From time $T_0 + 1$, the treatment is assumed to be affecting the treated unit until $t = T$.

Suppose that Y_{jt}^N is the outcome of unit j at time t if treatment did not occur, while Y_{jt}^T is the outcome value for unit j at time t if the treatment did occur. Let treatment be D_{jt}:

$$D_{jt} = \begin{cases} 1, \text{ if } j = 1 \text{ and } t > T_0 \\ 0, \text{ otherwise} \end{cases}$$

Since the treatment does not affect any unit prior to the treatment period $t = T_0 + 1$, we have $Y_{jt}^T = Y_{jt}^N$ for all units before $t = T_0 + 1$. The observed outcome for unit j at time t is:

$$Y_{jt} = Y_{jt}^N + \alpha_{jt} D_{jt}$$

Hence, the effect of treatment can be defined as: $\alpha_{jt} \equiv Y_{jt}^T - Y_{jt}^N$. This paper attempts at finding the treatment effect for unit 1, Georgia. In other words, to estimate $\alpha_{1T_0+1}, \alpha_{1T_0+2}, \dots, \alpha_{1T}$.

What is needed to find the treatment effect for unit 1 is Y_{1t}^N, which is the same as the counterfactual outcome of the time periods following the treatment period. Suppose that the counterfactual outcome can be given by the model:

$$Y_{jt}^N = \delta_t + \theta_t Z_j + \lambda_t \mu_j + \epsilon_{jt}$$

Where δ_t is a time trend. θ_t is a $(1 \times r)$ vector of unknown coefficients. λ_t is a $(1 \times F)$ vector of unobserved coefficients. Z_j is a $(r \times 1)$ vector of observed predictors that are

unaffected by the treatment. μ_j is an $(F \times 1)$ vector of unknown predictors. Lastly, ϵ_{jt} are zero mean individual transitory shocks.

The synthetic control unit is a weighted average of the donors. The synthetic unit can be defined as a $(J \times 1)$ vector of donor weights $W = (w_2, \ldots, w_{J+1})'$ with $0 \leq w_j \leq 1$ for $j = 2, \ldots, J+1$ and $w_2 + \ldots + w_{J+1} = 1$. Choosing W is to choose a synthetic control. How W is determined is shown and discussed in appendix B. The outcome of each synthetic control available from the donor pool is then:

$$\sum_{j=2}^{J+1} w_j Y_{jt} = \delta_t + \theta_t \sum_{j=2}^{J+1} w_j Z_j + \lambda_t \sum_{j=2}^{J+1} w_j \mu_j + \sum_{j=2}^{J+1} w_j \epsilon_{jt}$$

Consider that there are $(w_2^*, w_3^*, \ldots, w_{J+1}^*)$ optimal weights that reproduce both the outcome in the pre-treatment period and the observed covariates of the treated unit:

$$\sum_{j=2}^{J+1} w_j^* Y_{jt} = Y_{1t} \quad \text{and} \quad \sum_{j=2}^{J+1} w_j^* Z_j = Z_1,$$

Then we have an estimate for the treatment effect for the post-treatment periods $t = T_0 + 1, \ldots, T$:

$$\hat{\alpha}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt}$$

In most cases, the outcome values and observed variables of the treated unit will not be perfectly reproduced by the weights obtained. In theory, it can hold exactly. However, given that donor weights are non-negative and add up to one, this is only possible if the outcome variables and covariates of the treated unit belong to the convex hull of the donor units' outcomes and covariates. To approximate the economic characteristics of the treated unit, one at least needs the predictors of the treated unit to fall close to the convex hull of

the donor pool. It is therefore important that the treated unit is not an extreme in relation to the donor units in order to avoid biases (Abadie and Gardeazabal, 2003; Abadie, Diamond and Hainmueller, 2010; Abadie, 2020). If the treated unit is an extreme, it will not be well replicated by the synthetic unit. With SCM, one sees how great the discrepancies are and can make an informed decision on whether the synthetic control is good enough to be used. If the match is bad, another method than SCM should be considered (Abadie, Diamond and Hainmueller, 2010).

The linear model above displays one of the advantages of SCM over traditional difference-in-difference or fixed-effect panel models. Using such models, unobserved variables are allowed if they can be assumed to be constant over time, ie assuming $\lambda_t = \lambda$. SCM does not require this assumption (Abadie, 2020). Problems with the linear function model may arise when it is not suited for all the units in the donor pool, ie when the outcome in fact is determined by a non-linear model. Then, a seemingly good match may still give biased results. This concern can plausibly be overcome by only including donor units that have characteristics like those of the treated. By only sampling similar units, the risk of having a good fit result from an 'averaging away' of large differences in predictors is reduced (ibid).

SCM provides an unbiased estimator of the counterfactual (and consequently, the treatment effect) if

$$\sum_{j=2}^{J+1} w_j^* Z_j = Z_1$$

and

$$\sum_{j=2}^{J+1} w_j^* \mu_j = \mu_1.$$

One does not observe μ_1. But, intuitively, a unit's outcome variable is matched well by its synthetic comparison only if it is similar in both observed and unobserved characteristics. A good replication of Z_1 and $Y_{11}, \ldots, Y_{1T_0}$ can be obtained either when μ_1 is well reproduced or when the individual transitory shocks, ϵ_{jt}, compensate for a bad match of μ_1. The latter is unlikely if ϵ_{jt} are small or if T_0 is large. When T_0 is small and there is enough variation in ϵ_{jt}, then a close fit may be obtained even though μ_1 is ill-replicated. This is called overfitting and may cause bias. A credible counterfactual must thus be a good fit over many pre-treatment periods (Abadie, Diamond and Hainmueller, 2010, 2015; ibid).

5.1.2 *Inference and robustness checks with SCM*

In comparative case studies, the traditional methods of conducting inference are difficult to apply. Sample sizes are usually small, they are not randomised, and the units are not chosen probabilistically (Abadie, Diamond and Hainmueller, 2015). However, with SCM inference can be done by running in-space placebo tests. In-space placebo tests are conducted by assigning treatment status to units that were in fact not treated. This allows one to see whether the estimated effect for the treated unit is abnormally large relative to any effect estimated for a randomly chosen, non-treated donor unit. The distribution of the test statistic is obtained by calculating the probability of estimating among the sampled units an effect which is greater than or equal to the estimated effect for the treated unit. An unusually large estimated effect for the treated speaks in favour of a causal effect of treatment (Abadie, Diamond and Hainmueller, 2010).

In this paper, I employ the inference procedures used by Abadie, Diamond and Hainmueller (2010, 2015). For one, these inference methods are easily performed using the statistical software packages available for synthetic control studies, `synth` (Abadie,

Diamond and Hainmueller, 2010) and `synth_runner` (Galiani and Quistorff, 2017). For another, these inference methods are the standard means of inference in the literature. The criticism levelled against these inference methods is discussed in appendix D.

The first inference method calculates the probability of estimating an effect of the same or greater magnitude for *each* post-treatment period by running the described in-space placebo tests and obtaining a test statistic distribution.

The second inference method builds upon the first. Abadie, Diamond and Hainmueller (2015) note that treatment effects may be large for some periods, but not for all. Then, relying on the p-values of the first method may be too conservative. Instead, the authors recommend SCM practitioners to consider the empirical distribution of the ratio between post- and pre-treatment prediction errors. Intuitively, for the treated unit, a large *post-treatment* root mean squared prediction error (RMSPE) – ie a bad fit – is expected if the treatment is assumed to have had an effect. This because a treatment effect implies a divergence between actual and synthetic Georgia post-treatment. With regards to the SCM requirement that the pre-treatment prediction error is low, the treated unit should have a relatively large ratio between post- and pre-treatment RMSPEs. A more detailed description on ways of calculating RMSPEs and p-values is found in appendices B and C.

Robustness of results can be checked by conducting in-time placebo tests. Intuitively, a treatment happening in year $t = T_0 + 1$ should not have an effect already at, eg, $t = T_0 - 5$. If a treatment effect is estimated in a pre-treatment year, then the credibility of an estimated effect is undermined, as it would imply that the synthetic control lacks predictive power (Abadie, Diamond and Hainmueller, 2015).

One can also perform leave-one-out (L-O-O) tests. L-O-O test checks the results' sensitivity to changes to country weights. By re-running the test while excluding one of the positively weighted donors one at a time, Abadie, Diamond and Hainmueller (2015) see how dependent their findings are on any one donor. Goodness fit is thus traded off for a sensitivity analysis.

5.2 Data collection

5.2.1 *Donor pool*

Based on section 5.1, panel data for European and Central Asian post-communist countries is collected. The baseline donor pool is seen in table 1.

Table 1: baseline donor pool

Former eastern-bloc states	
Albania	Lithuania
Armenia	Moldova
Azerbaijan	Montenegro
Belarus	North Macedonia
Bosnia	Poland
Bulgaria	Romania
Croatia	Russia
Czech Republic	Serbia
Estonia	Slovakia
Georgia (treated unit)	Slovenia
Hungary	Tajikistan
Kazakhstan	Turkmenistan
Kosovo	Ukraine
Kyrgyzstan	Uzbekistan
Latvia	

The literature recommends that only units that are similar to the unit of interest are included in the donor pool (Abadie, Diamond and Hainmueller 2010, 2015). A baseline donor pool is selected based on their historical economic and cultural similarity to Georgia. The units first chosen to construct the donor pool are the post-communist states previously belonging to the USSR, the Warsaw Pact, or to Yugoslavia. These units all have in common a history of planned economies, pervasive corruption and, in many cases, political and economic turmoil in the 1990s. Support for my choice of only including former eastern-bloc countries is found in Bluszcz and Valente (2020) on Ukraine as well

as in Lawson, Grier and Absher (2019) on Georgia. The former only uses post-communist states in their donor pool. The latter has a larger donor pool but generates a synthetic control in which only former eastern-bloc republics are given non-zero weights.

5.2.2 *Selection of variables*

A set of variables is selected based on the literature and the peculiarities of the post-communist context (see table 2). These are considered when developing my model.

Table 2: variables employed

Variable name	Description
pcGDP	Gross domestic product per capita at purchasing power parity (PPP) in constant 2017 international US dollars.
GCFshare	Gross capital formation (% of GDP)
FDIshare	Foreign direct investment, net inflows (% of GDP)
Exports	Exports of goods and services (% of GDP)
Inflation	Inflation, GDP deflator (annual %)
Urban	Urban population (% of total population)
MilExp	Military expenditure (% of GDP)
LifeExp	Life expectancy at birth, total (years)
HDI	Human Development Index, high = better
TIC	EBRD Transition Indicators (mean), high = better
EUMS	EU membership (dummy)
Freedom	Freedom House Index (mean), low = better
REMshare	Personal remittances, received (% of GDP)
ODAshare	Net official development assistance (% of GDP)

Data is provided by the World Development Indicators database (World Bank, no date), the European Bank of Reconstruction and Development's Transition Indicators data set (European Bank of Reconstruction and Development, no date), and by Freedom House (Freedom House, no date). I calculate net official development assistance (ODA) as share of GDP using separate ODA and GDP data from the World Bank. EBRD data are used to generate a variable which is the yearly average over six transition indicators for each sampled country. An EU membership dummy is generated based on EU enlargement data (De Munter, 2020). The Freedom House variable is generated from Freedom House data and is the yearly average of each country's political rights and civil liberties scores.

5.2.3 *Time period*

Data are collected to cover the period 1995-2013. While it is preferable to have as many pre-treatment periods as possible, the year 1995 is chosen as the first pre-treatment year due to the economic turbulence of the early 1990s (Lawson, Grier and Absher, 2019).

Considering that literature on economy and conflict and that capital is fixed in the short-term, a war likely has both immediate and lingering effect on the economy. With regards to the post-treatment period, observations for several years are needed if the effect of treatment is likely to appear gradually. At the same time, the further you move from the time of intervention the more likely are other events to confound estimated treatment effects. The year 2013 is chosen as the final sampled year because any estimated effect beyond 2013 may be confounded by spill-overs caused by Russia's invasion of Ukraine in 2014.

5.3 Implementation

Based on the previous sections, this section shows how SCM is operationalised to answer the research question.

5.3.1 *Model specification*

The dependent variable is per capita GDP in constant, PPP-adjusted US dollars. It is a variable used in seminal papers by Abadie and Gardeazabal (2003) and Abadie, Diamond and Hainmueller (2015). As discussed above, SCM requires that the outcome can be modelled by a linear function. Considering the frequent use of GDP per capita as outcome variable in influential SCM papers, it is assumed that this condition holds. GDP per capita *levels* are most frequently used in the literature and are therefore also used in this paper. Furthermore, a requirement for SCM is that the outcome variable is available for all units

in all time periods (Echevarría and García-Enríquez, 2020). For this reason, Montenegro and Kosovo are dropped from the baseline donor pool.

With regards to the independent variables, there is a vast literature on predictors of per capita GDP. My model specification is based on that of Lawson, Grier and Absher (2019).[3] Based on the literature on determinants of growth, Lawson and co-authors use FDI share, GCF share, exports share, and inflation rate. They also add socio-demographic variables: urban population share, life expectancy, and military expenditure (Lawson, Grier and Absher, 2019). Like Lawson and co-authors, I would prefer to include an education variable. However, lack of data makes this impossible. Serbia lacks data for some of the listed variables. Therefore, Serbia is dropped.

In addition to the variables used by Lawson and co-authors, I attempt at improving the synthetic control's fit by including combinations of other variables listed above, for instance `TIC`, `EUMS`, `Freedom`, `REMshare` and `ODAshare`. These variables are not found to increase the model's goodness of fit. They are therefore not included in my model.

As discussed in section 5.1, SCM requires that the treated unit is close to the convex hull of the donor pool. Preferably, the values of the unit of interest for each variable are centred around the middle of the distributions. Figure 1 shows how Georgia relates to the donors. Georgia does not appear to be an extreme in terms of most variables. Georgia's predictor values will therefore likely be well-replicable by the donor pool without extrapolation. In terms of military expenditure Georgia is close to being an extreme. This is a possible cause for bias. Therefore, I do not use the military expenditure variable.

[3] Results from running their model is found in appendix G.

With regards to lagged outcome predictors, the SCM literature has expended much digital ink on the appropriate way to use these variables in SCM models. The benefit of using lagged outcome variables is a good fit for the synthetic control. They are also important in reproducing unobservable predictors of the outcome variable (Abadie, 2020). However, researchers have highlighted the risk of using too many lagged outcomes (Abadie and Gardeazabal, 2003; Kaul et al, 2018). According to Kaul et al (2018), the use of lagged outcomes for all years renders all other predictors irrelevant for predicting post-treatment outcomes. To use too many lagged outcomes may thus cause biased results if influential predictors are rendered meaningless, ie assigned zero weights. (Kaul et al, 2018).

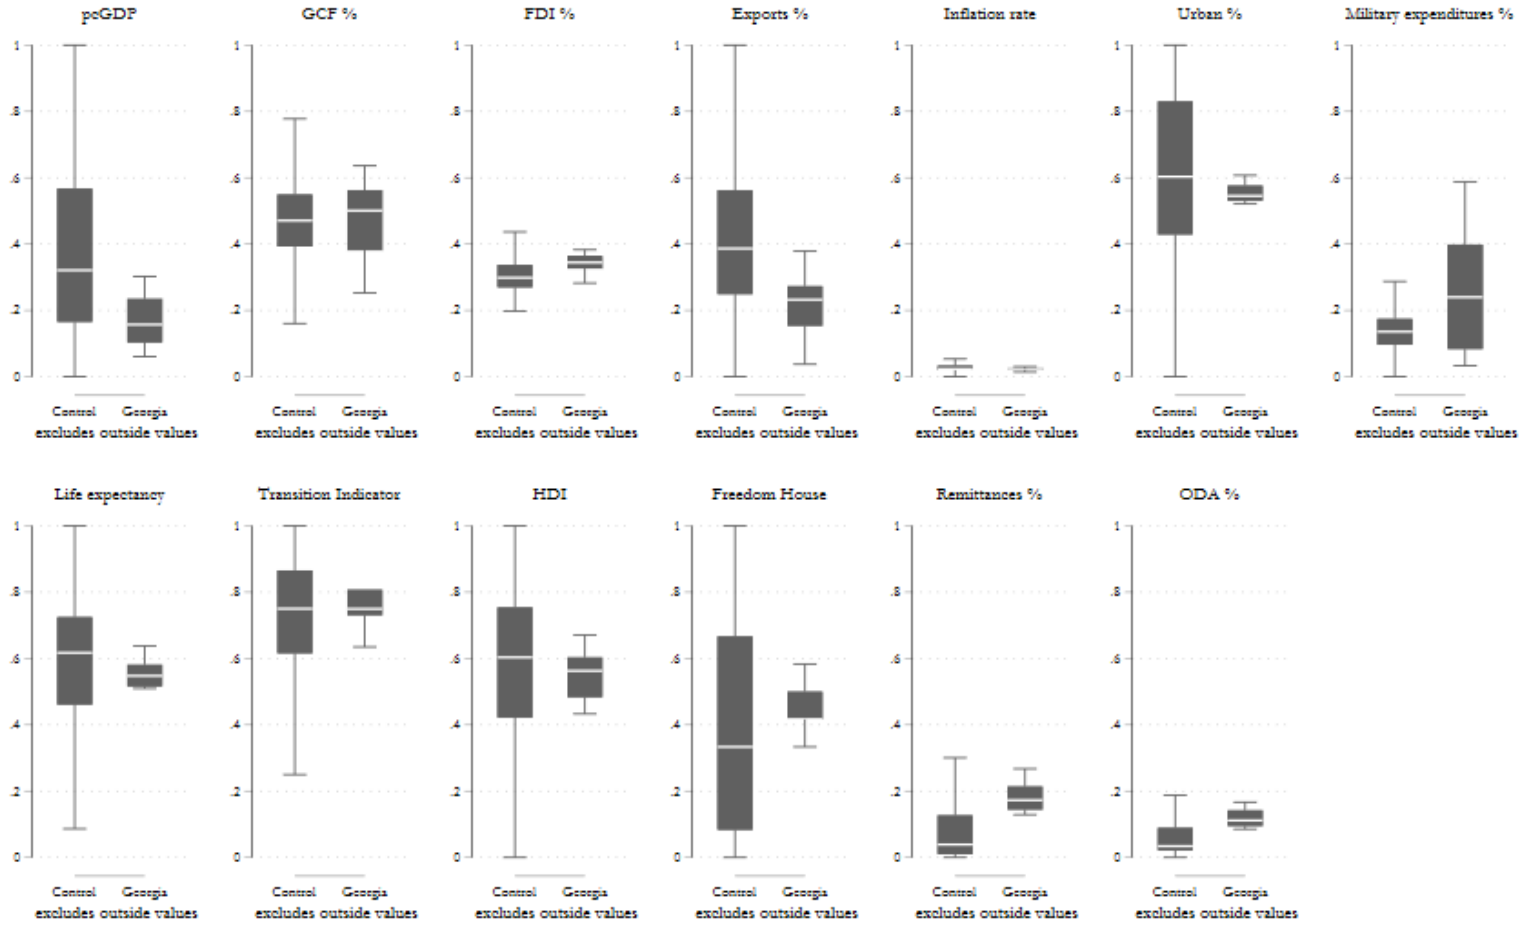


Figure 1: Georgia's place in the convex hull of donor units. Shares are measured as shared of GDP. Data is normalised. Source: my rendering of data from the World Bank (2021), EBRD (2021), Freedom House (2021)

Authors have chosen different ways to incorporate lagged outcomes in their models. Abadie (2020) writes that if a good fit is obtained, researchers are free to decide on which

lagged outcomes to use as predictors. To my knowledge, the lowest RMSPE is obtained when three lagged outcome predictors (2007, 2005, and 2000) are used in my model.

In sum, the model that provides the best counterfactual for Georgia is specified as such:

Model specification	
Lagged outcome predictors	`pcGDP(2007), pcGDP(2005), pcGDP(2000)`
Predictors	`FDIshare, Urban, GCFshare, LifeExp, Inflation, Exports`
Treatment period	2008
Results period	1995-2013

5.3.2 *Donor pool*

The baseline donor pool is constituted by former eastern-bloc republics. It is, however, important to note that even when restricting the donor pool to countries that are culturally or historically similar, biases may still be present if states within this group differ widely in terms of economic characteristics. For instance, the economies of Azerbaijan and Russia are highly dependent on oil. Therefore, Azerbaijan and Russia are excluded.

By the identifying assumptions, units should also be excluded from the donor pool if they are likely to be affected by spill-overs of the treatment. While there are many states likely affected by either Russia or Georgia (or both) being struck by conflict, the Russo-Georgian war is likely to have had relatively little effect on states other than those directly involved. Unlike Russia's war against Ukraine (which led to Western sanctions), few actions were taken in response to Russia's violation of international law in 2008. Russo-Western relations were soon back to normal following the war (Hamilton, 2018). Hence, there is not enough reason to exclude units due to them being affected by spill-overs caused by sanctions. However, Russia, as a direct participant in the conflict, should be excluded.

A country that is likely to have been severely affected by spill-overs from the conflict is neighbouring Armenia. As a result of the war, Armenia lost access to important transportation routes through Georgia. Unlike Georgia, Armenia did not receive vast financial support following the crisis (Avagyan, 2009; Remler, 2020). If Armenia is given a non-negative weight as a donor to synthetic Georgia, a negative spill-over originating from the war is likely to cause a downward bias to any estimated treatment effects in the period immediately following the war. However, Armenia will not be excluded, because the bias is assumed to be short-term and clear in its (downward) direction while Armenia is needed to acquire a good fit for synthetic Georgia. The trade-off of including Armenia is discussed in the sensitivity analysis.

Others that might have been affected by the war are exporters to Georgia. The main exporters to Georgia in 2007 were Turkey, Ukraine, Azerbaijan, Germany and Russia (World Bank, 2021). Among these, only Ukraine, Azerbaijan and Russia belong to the baseline donor pool. While there is a possibility that they were affected by decreased Georgian demand, it is unlikely that this effect would be substantial considering the relative size of the Georgian economy and the negligible importance that exports to Georgia have to these economies. For instance, Ukrainian exports to Georgia constituted merely 0.85 % of Ukrainian exports in 2008 (The Observatory of Economic Complexity, no date). Azerbaijan and Russia are already excluded due to aforementioned reasons. Ukraine will, however, not be left out due to this presumed negligible spill-over effect.

Other reasons to exclude units from the donor pool are if they experienced idiosyncratic shocks to the outcome variable or if they experienced similar treatment in the sampling period (Abadie, Diamond and Hainmueller, 2015; Adhikari et al, 2018). Abadie, Diamond and Hainmueller (2010) exclude US states that imposed similar tobacco legislation as

California in the sampling period. Abadie, Diamond and Hainmueller (2015) exclude Sweden, Finland, and Ireland due to idiosyncratic shocks in the 1990s. Like Georgia, both Bosnia and Tajikistan experienced war in the sampling period. Bosnia suffered conflict until the end of 1995 and Tajikistan saw civil war until 1997. If these conflicts cause these countries to differ greatly from pre-treatment Georgia, they are unlikely to be given positive weights as donors. If these units contribute to synthetic Georgia while the effects of these conflicts linger past the treatment year, the results may be biased. By looking at these two units' growth trajectories, it is seen that both countries follow a stable growth path from 1997 onwards (see appendix E). Bearing this in mind and noting that SCM averages pre-treatment characteristics over the entire pre-treatment period, I believe that any effects of these Bosnian and Tajiki conflicts constitute a minor problem to this paper's model**.** Bosnia and Tajikistan will therefore be kept in the donor pool. The results' sensitivity this choice is tested in the sensitivity analysis.

With regards to idiosyncratic shocks, many countries suffered from the financial crisis of 2008. This shock was, however, not idiosyncratic but global. Most of the sampled units are likely affected. If the synthetic control is the pre-treatment period is tracking actual Georgia closely while predictors are matched well, the effect of the financial crisis should be replicated by the synthetic control. The financial crisis as a possible confounder is discussed at length in section seven.

In sum, the sampled units used in the model are listed in table 3.

Table 3: Baseline donor pool and measures taken

Former eastern-bloc states 1(2)	Dropped?	Former eastern-bloc states 2(2)	Dropped?
Albania		Lithuania	
Armenia		Moldova	
Azerbaijan	Yes	Montenegro	Yes
Belarus		North Macedonia	
Bosnia		Poland	
Bulgaria		Romania	

Croatia		Russia	Yes
Czech Republic		Serbia	Yes
Estonia		Slovakia	
Georgia (treated unit)		Slovenia	
Hungary		Tajikistan	
Kazakhstan		Turkmenistan	
Kosovo	Yes	Ukraine	
Kyrgyzstan		Uzbekistan	
Latvia			

6 RESULTS

In this section, the effect of the war is estimated. I construct a synthetic control unit with the purpose of obtaining an unbiased counterfactual for Georgia. Then, the robustness of results is tested by using (two-sided) in-space placebo tests and in-time placebo tests along with a leave-one-out test and a test of the results' sensitivity to changes to variable weights.

6.1 The causal effect of the Russo-Georgian war

By running the synthetic control model, synthetic Georgia is constructed as a weighted average of Armenia, Bosnia, Kyrgyzstan, Macedonia, Slovakia, and Turkmenistan. A detailed explanation of how the model selects donors is found in appendix B. As seen in table 4, Armenia is the main donor to the synthetic control and is given more than 50 % weight.

Table 4: synthetic control donor weights

Albania	0	Lithuania	0
Armenia	.515	Macedonia	.008
Belarus	0	Moldova	0
Bosnia	.164	Poland	0
Bulgaria	0	Romania	0
Croatia	0	Slovakia	.042
Czechia	0	Slovenia	0
Estonia	0	Tajikistan	0
Hungary	0	Turkmenistan	.049
Kazakhstan	0	Ukraine	0
Kyrgyzstan	.222	Uzbekistan	0
Latvia	0		
			Sum: 1.000

The predictor averages of synthetic Georgia and actual Georgia are seen is table 5. Notice that the lagged outcome variables account for almost all predictive power in the model. It

is important to note that even though most explanatory power is given to one variable, `pcGDP(2005)`, the model's fit is dependent on the inclusion of all these variables. While it is common in the literature that lagged outcomes are given large weights, this weight is unusually high and a potential problem. The results' sensitivity to changes to variable weights is discussed in section 6.1.2.3.

Table 5: treated and synthetic unit predictor averages in pre-treatment period

		Georgia		
Variable	Weights	Treated unit	Synthetic control unit	Donor average
pcGDP(2007)	.00021333	9109.898	9044.059	17345.662
pcGDP(2005)	.00036302	7315.198	7369.354	15003.816
pcGDP(2000)	.99834707	4919.232	4921.027	11224.387
FDIshare	6.902e-08	8.038	4.913	5.021
Urban	.0009498	53.254	53.259	56.136
GCFshare	8.943e-09	25.931	25.763	25.256
LifeExp	3.504e-12	70.100	70.764	70.557
Inflation	.00012654	21.742	21.728	33.518
Exports	1.636e-07	25.059	30.314	42.779
			RMSPE 58.644	

Table 5 highlights the benefit of SCM compared to using donor pool averages as an alternative method to construct a counterfactual in case studies at the aggregate level. It is seen that the synthetic control method does a better job in constructing a counterfactual for Georgia than donor averages listed in the fifth column.

According to Abadie, Diamond and Hainmueller (2015), the credibility of a synthetic control is dependent on its ability to follow the treated unit's outcome variable trajectory pre-treatment. The outcome variable trend of both synthetic and treated Georgia is displayed in figure 2. Like Abadie, Diamond and Hainmueller (2010), a trendline in the graphs is set to year 2007. Since the war took place in August 2008, divergence is expected to occur after 2007. The graph shows that synthetic Georgia is closely tracking observed Georgia up until treatment. In other words, the combination of donor units reproduces a credible counterfactual for Georgia. Even though there was increasing tension before 2008, there does not seem to have been anticipation effects causing divergence before the

treatment year. The model has a good fit and an RMSPE that amounts to about 1 % of the average per capita GDP of Georgia in the pre-treatment period.

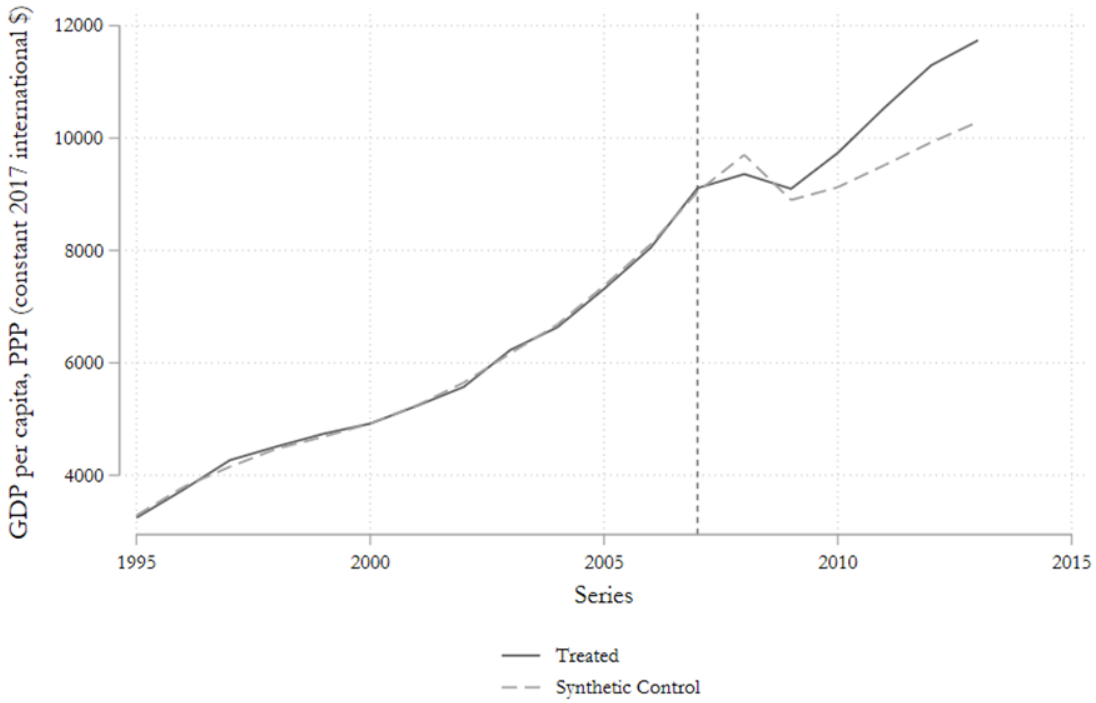


Figure 2: synthetic Georgia and observed Georgia

The estimated effect is the difference between synthetic Georgia's and observed Georgia's lines. The graph shows that in 2008, treated Georgia performs worse than synthetic Georgia, implying that there was an initial loss of GDP per capita resulting from the war. Table 6 shows that the estimated effect in 2008 amounts to a per capita GDP loss of approximately $340 (a loss of about 3.5 %). Thereafter, the estimated treatment effect is positive and increasing in size for each consecutive year. Interestingly, treated Georgia surpasses synthetic Georgia in terms of GDP per capita already in 2009.

Table 6: effects and standardised p-values

	Treated Georgia	Estimated effects	Standardised p-values
2008	9359	-343.1849	0.286
2009	9097	199.6734	0.571
2010	9737	611.3734	0.143
2011	10541	1021.679	0.143
2012	11296	1370.824	0
2013	11740	1452.536	0

6.1.1 *Inference*

Using the first inference method, the likelihood of estimating a similar treatment effect for each individual year is measured. These p-values are obtained when limiting the units used for calculating p-values to those with an RMSPE no greater than five times that of Georgia. Then, seven donors remain for calculating the p-values.[4] The effect in 2008 is found to be negative but insignificant at any conventional level (see table 6).[5] Thereafter, the effect is insignificant up until 2012. In 2012 and 2013, the war is found to have had a positive, causal effect significant at the 0 % level. The difference to per capita GDP amounts to about $1400 in each of the two periods, making the per capita GDP of treated Georgia about 14 % higher than that of counterfactual Georgia.

Moving on to the second method of inference. By comparing RMSPE-ratios, it is seen in figure 3 that Georgia's RMSPE-ratio is substantially larger than most other units. An exception is Belarus. In all, the post-war years have an accumulated net per capita GDP effect of $4314, more than $700 (on average about 7.5 %) per year. The probability of estimating an effect as high as (or higher than) Georgia's when treatment in fact did not occur is thus 1/21, or, 4.76 %.

Because Georgia is the only unit undergoing treatment, it is expected to also be the unit with the single highest RMSPE-ratio. Belarus's high ratio needs explanation. By setting

[4] These results are sensitive to the cut-off multiple chosen. When setting the cut-off to a multiple of 20, almost every donor is included in p-value calculations, which makes the results significant at the 10 % level in 2008 and at the 5 % level in every other year except for 2009. When setting the cut-off to a multiple of 10, then seventeen donors are used and results are significant at the 12 % level in 2008, at the 10 % level in 2010 and 2011, and at the 5 % level in 2012 and 2013. However, the inclusion of units with too poorly fitted synthetics puts one in the risk of mistaking divergence caused by an ill-fitting model for treatment effects. I have therefore chosen to use the p-values obtained from setting a multiple of 5. The (more significant) results from setting the cut-off to 20 or 10 are – unlike Abadie, Diamond and Hainmueller (2010) – in this paper not presented in further detail for the sake of preserving space.

[5] It is significant at the 5 % level using a one-sided test, but this paper relies on two-sided tests.

Belarus as treated unit and running the model, a synthetic Belarus is obtained. Figure 4 shows that it produces a good fit in the pre-treatment period. The model finds that Belarus's positively weighted donors are by order of the size of their weights Tajikistan (0.358), Latvia (0.232), Estonia (0.207) and Turkmenistan (0.204).

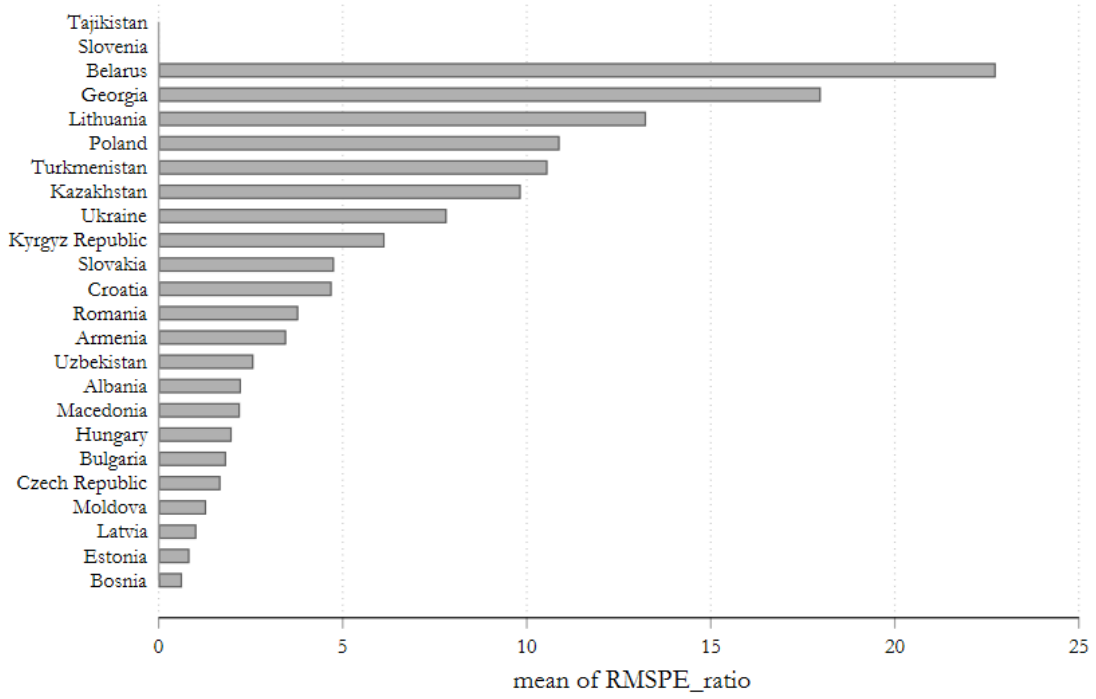


Figure 3: RMSPE-ratios of the sampled units. Values missing for Tajikistan and Slovenia due to software failure

If Belarus, or any of its synthetic control unit's donors, experienced a severe idiosyncratic shock in the post-treatment period, Belarus would be difficult to replicate for its donors and high predictive errors (and consequently a high RMSPE-ratio) would result. It is seen in appendix E that Estonia and Latvia were severely affected by the financial crisis of 2008. The two units' per capita GDP losses were much greater than that of Belarus. Since these two Baltic states together constitute more than 40 % of synthetic Belarus, the difficulty in replicating Belarus in the post-2008 years is explained. Idiosyncratic shocks among the donors cause a divergence between observed and synthetic Belarus which in turn causes a high ratio between post- and pre-treatment RMSPEs. When re-running the Byelorussian model while excluding the Baltic states from its donor pool, a worse fit is obtained, but

the large gap seen in figure 4 is much diminished (results not shown here). Hence, there is reason why Georgia is not given the highest RMSPE-ratio.

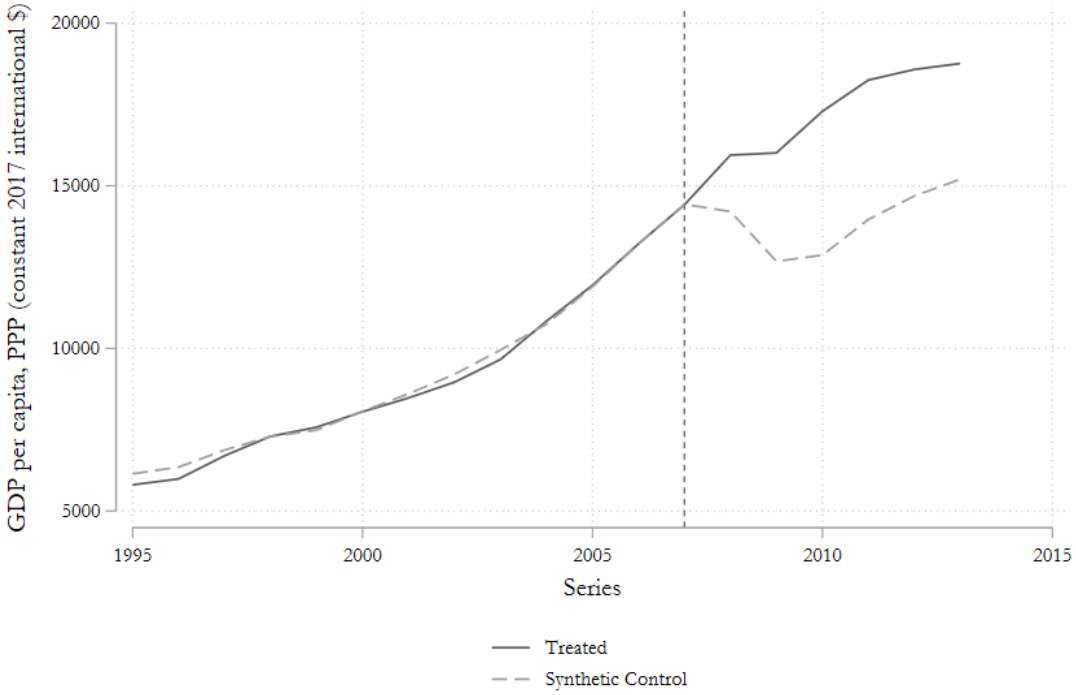


Figure 4: synthetic Belarus and observed Belarus

6.1.2 *Robustness checks*

6.1.2.1 In-time placebo tests

As a confoundedness check, the effect of the Rose Revolution on Georgian per capita GDP is estimated. The same model is used, but with lagged outcome variables adjusted accordingly in time. There is a small estimated divergence in 2003-2008, suggesting that the Rose Revolution had an effect. However, the estimated effects are small and highly insignificant. A major divergence does not occur until after 2008, which is in line with my earlier findings (see appendix F). This increases the credibility of my synthetic, as there are no major confounding effects estimated prior to 2008 (Abadie, 2020). The synthetic unit is still (decently) reproducing the outcome variable's trajectory up until the time of the actual treatment. Note that due to the scarcity of available pre-treatment years, results obtained from setting 2003 as treatment year are not as reliable as those reflected in figure 2. The results from setting treatment to 2003 differ from those of Lawson and co-authors

(2019), who estimate a positive and significant effect of the Rose Revolution on Georgian per capita GDP in 2007 and in 2009 and onwards. This may be due to differences to the models and data sets used in each respective paper. It may also be due to Lawson and co-authors having the potential problem of scarcity in pre-treatment periods and/or due to them having their results confounded by the Russo-Georgian war. In contrast to Lawson, Grier and Absher (2019), I am more conservative in selecting the donor pool units.

Another confoundedness check is performed by setting treatment to 2006. As discussed in section two, tensions between Tbilisi and Moscow began to rise with the election of Mikheil Saakashvili. In 2006, Russia imposed an embargo on Georgia that severely limited Georgian exports to Russia and cut Russian energy and gas supplies to Georgia (Transparency International, 2020). When moving treatment to 2006 and running a synthetic control, no apparent negative effect of the sanctions is identified (see appendix F). This test together with the preceding confoundedness test indicate that the causal effect estimated in the main model is not confounded by the most likely confounding events in the pre-treatment period, the Rose Revolution and the 2006 embargo.

6.1.2.2 Leave-one-out test

By iteratively removing one donor country at a time, the results' sensitivity to the inclusion of the positively weighted donor units is tested. Leave-one-out tests sacrifice fit for the benefit of seeing whether any one donor unit is driving the results. These results are not as robust as those in figure 2 and should be evaluated accordingly. By running an L-O-O test, a kind of confidence interval for the treatment effect is obtained (Echevarría and García-Enríquez, 2020; Abadie, Diamond and Hainmueller, 2015).

Table 7: iterate leave-one-out tests

	Excluded units from the model

Year	Georgia pcGDP*	Original effect	Armenia	Bosnia	Kyrgyz Republic	Macedonia	Slovakia	Turkmenistan
2008	9359	-343	**456**	*-387*	-313	-353	-220	-250
2009	9097	200	**1071**	*113*	194	188	363	247
2010	9737	611	**1818**	*517*	549	595	808	736
2011	10541	1022	**2022**	*863*	930	1001	1186	1197
2012	11296	1371	**2359**	*1164*	1232	1347	1540	1616
2013	11740	1453	**2442**	*1273*	1304	1429	1608	1720

*Observed Georgian pc GDP. **Bold** (*italics*) numbers represent the highest (lowest) estimated effect.

In figure 5, it is seen that the estimated negative effect on Georgia in 2008 is sensitive to the donors included. In year 2008, treated Georgia is under-performing compared with synthetic Georgia in all L-O-O tests, except for when Armenia is excluded from the model. When Armenia is left out, actual Georgia is over-performing its synthetic counterfactual in every post-treatment period. As seen in figure 5, the gap between treated Georgia and synthetic Georgia is relatively large when excluding Armenia. The change to results is quite large even though the pre-treatment fit looks similar. This may result from other interventions affecting new donors now given positive weight due to the exclusion of Armenia (Abadie, 2020). When Armenia is left out, Latvia (a country severely struck by the financial crisis) constitutes approximately 20 % of synthetic Georgia. Table 7 displays the estimated treatment effects obtained when iteratively excluding one donor at a time.

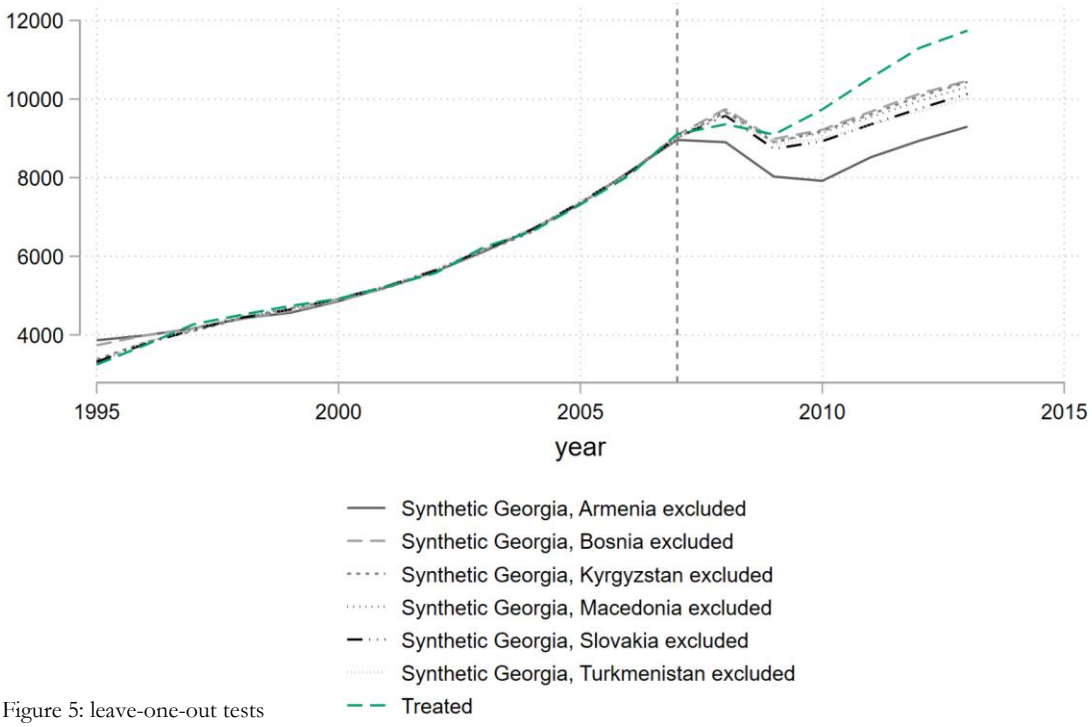


Figure 5: leave-one-out tests

The L-O-O test gives a minimum estimated average per capita GDP effect of 590 and a maximum estimated average effect of 1695 (5.7-16.5 %) post-treatment. The L-O-O test shows that the sign of the estimated effect in 2008 as well as the magnitude of the estimated overall treatment effects are sensitive to the set of donor units included. However, while the magnitudes differ, the L-O-O tests point in the same direction as my results above: the war's effect seems to have been overall positive in the post-war period.

6.1.2.3 Sensitivity to changes in variable weights

As a final robustness test, I attempt at identifying possible bias caused by the lagged outcome variables being given 99.99 % weight by my model. It is common in the literature that lagged outcome predictors are given weights of more than 90 % (see for instance Echevarría and García-Enríquez, 2019, 2020; Bluszcz and Valente, 2020). However, 99.99 % is an unusually high weight and might render in the kind of biases described by Kaul et al (2018), referenced above.

This potential bias is controlled for by running the model again, but with per capita GDP (log) instead of per capita GDP (levels) as outcome variable. By using logs, the lagged outcome variables are less heavily weighted and given about 94 % predictive power instead of almost one hundred percent. A good pre-treatment fit is obtained by the model. Now that the other predictors are given higher weights, the results are slightly different but remain substantially the same. The main donors are still Armenia, Bosnia and Kyrgyzstan and they are together given about the same weight as in the model using levels. With the first inference method, looking at individual years, the positive effect in 2011 is more significant than in the levels model and the effect in 2013 is now almost significant (~5.3 %) at the 5 % level. With the second inference method, the overall effect is still positive

and close to significant (~5.3 %) at the 5 % level. The similarity to results obtained by the level and logarithmic outcome variable models suggest that the kind of bias described by Kaul et al (2018) does not seem to be affecting my overall conclusions.

Table 8: synthetic control donor weights (log model)

Albania	0	Lithuania	0
Armenia	.487	Macedonia	0
Belarus	0	Moldova	0
Bosnia	.182	Poland	0
Bulgaria	.04	Romania	0
Croatia	0	Slovakia	0
Czechia	0	Slovenia	0
Estonia	0	Tajikistan	0
Hungary	.067	Turkmenistan	.027
Kazakhstan	0	Ukraine	0
Kyrgyzstan	.197	Uzbekistan	0
Latvia	0		
			Sum: 1.000

Table 9: treated and synthetic unit averages in pre-treatment period (log model)

		Georgia	
Variable	Weights	Treated unit	Synthetic control unit
logpcGDP(2007)	.06836143	9.117	9.090
logpcGDP(2005)	.5985284	8.898	8.907
logpcGDP(2000)	.27203121	8.501	8.502
FDIshare	.00852292	8.038	5.617
Urban	.02214868	53.254	54.656
GCFshare	00393935	25.931	25.463
LifeExp	.00008489	70.100	71.005
Inflation	.01661572	21.742	21.710
Exports	.00976741	25.059	31.010
			RMSPE 0.016

Table 10: estimated effects and standardised p-values (log model)

	Estimated effects	Standardised p-values
2008	-.013	.579
2009	.035	.632
2010	.087	.158
2011	.125	.053
2012	.162	0
2013	.159	.053

7 DISCUSSION

This paper attempts at answering the question: how did Russia's invasion of Georgia affect Georgia's per capita GDP? While the literature on economy and war is ambiguous, I

hypothesised that the war caused an immediate decline followed by a convergence to Georgia's pre-treatment trajectory. Overall, a limited effect on Georgia's economy was expected. The results are in part confirming and in part rejecting the research hypothesis. In accordance with the hypothesis, an immediate dip is estimated. It is, however, statistically insignificant using a two-sided test. Contrary to the hypothesis, I obtain robust results that show that the war had a positive, large, and significant effect on Georgian per capita GDP in 2008-2013. The effect seen is different from the 'peace dividend' effect discussed in the literature, as Georgia not only converges to the counterfactual level, but even surpasses it.

Moving from the 'how' to the 'why', Cavallo et al (2013) highlight that SCM only gives the reduced form aggregate effect of treatment. The channels of transmission through which the war affected the outcome variable are not identified by the model. Neither does it tell whether the estimated effect is driven by the studied treatment or by other events. The 'why' is thus given by qualitative reasoning. Below, I propose that the positive effect estimated could be derived from the war's effect on Georgia's resource allocation but is more clearly connected to the 'policy and political institutions' channel. The latter relates to Georgia's westward relations in both the pre- and post-treatment periods.

First, Georgia's high levels of military expenditure in the early 2000s may have diverted resources away from productivity-increasing investments. In the literature, such crowding out effects of military spending are discussed, eg, by Cappelen, Gleditsch and Bjerkholt (1984). After having failed in 2008 to restore its territorial integrity, Georgian military spending declined dramatically. It went from more than 8 % of GDP in 2008 to 4 % in 2010. Meanwhile, the spending of its main synthetic donor Armenia remained relatively stable around 4 %. Georgia's reallocation of resources likely affected its GDP in the post-

treatment years. However, whether it translates into a *positive* economic effect is conditional on the assumption that Georgia's military expenditure indeed had a crowding out effect that hampered the economy. The literature is ambiguous in this regard.

Second, considering the policy and political institutions channel, the war resulted in a Georgian pivot to the EU (Emerson and Kovziridze, 2016). As Georgia became more integrated with Western Europe, repressive policies pursued by President Saakashvili prior to the war could not reasonably be continued. Georgia's commitment to liberal reforms would thus have been reinforced by the Russo-Georgian war. This relates to Olson (1982) and his notion of creative political destruction, but also to the Feng (2001). Feng (2001) connects post-war outcomes to the war's effect on political instability and freedoms. In the post-war period, Georgia's democracy recovered from a temporary increase to political repression starting in 2007. Data shows that the average Freedom House score in 2008-2013 is lower than that in the pre-war period, but only slightly. Moreover, postbellum, Georgia's state fragility score went down while its polity score increased steadily from 2008 onwards (Hamilton, 2018). These changes are indicative of societal improvement. Hence, in line with the research by Olson (1982) and Feng (2001), Georgia's post-war economic development may have been helped by the war necessitating a deeper integration into a liberal Western community which promoted policies conducive of growth. However, the causal relation between this pivot to the West and the war is not clear. While scholars connect Georgia's westward push to the war, it is not certain that they have a causal relationship. These societal improvements may have occurred regardless. A factor that is clearly related to the war is discussed next.

Last, and most importantly, Georgia's westward relations helped it secure funds amounting to $4.5bn (more than a third of Georgia's 2008 GDP) in support of its post-war recovery.

Prior to the conflict, both Georgia and its near peer Armenia had relatively low levels of external debt to GDP. Still, both countries issued debt rated below investment grade (Trading Economics, no date). However, following the war, Georgia benefitted from inexpensive, concessional loans and grants while Armenia had to turn to the financial markets to finance fiscal stimuli (Avagyan, 2009; Marutyan et al, 2010). While Armenia also received concessional loans and aid in the studied period, the support received by Armenia was not of the same magnitude as that of Georgia (International Monetary Fund and World Bank, 2011). The sampled data show that Georgia's net official development assistance average doubled in the post-war period compared to the pre-war period. Armenia, meanwhile, only saw an increase to its average of about 5 %. In addition to Georgia receiving more funds, Georgia's post-war situation (its reformed economy and policies) may, by the findings of Collier and Hoeffler (2004), also have put Georgia in a better position than Armenia to absorb aid. Accordingly, in the war's aftermath, Georgia's government spending rose sharply while those of Armenia followed a much flatter trend (Trading Economics, no date). The United Nations and the World Bank conclude that the donor intervention in Georgia was successful in preventing a collapse of Georgian demand, supported employment, and bridged the period of low private sector investments (United Nations and World Bank, 2010). Hence, the large influx of aid is likely a causal reason why war-ridden Georgia outperforms its counterfactual.

In sum, SCM does not identify a treatment's causal channels. These can, however, be qualitatively assessed. It is possible that a Georgian reallocation of resources benefitted the Georgian economy in 2008-2013. However, the literature is ambiguous in its conclusions as to whether the net economic effect of military spending is positive or negative. Instead, I believe that policy and aid are the obvious determinants of the large, positive treatment effect estimated. The results relate to the findings of Olson (1982), Feng (2001), and Collier

and Hoeffler (2004), as an interaction between Georgia's post-war situation and the aid received could translate into the economic growth seen post 2008. Unlike its near peers, Georgia was a darling of the West that had positioned itself to both secure the funds needed in the post-war years and make the best out of the funds received thanks to its institutions and policies. It entered the global recession with billions of dollars provided because of Russia's 2008 invasion and was thus able to recover from the crises of 2008 and soon thereafter surpass its counterfactual.

7.1 Internal validity

Among the identifying assumptions discussed in section five, most are addressed in section 5.3. There are, however, outstanding concerns relating to anticipation effects, violations of SUTVA, and the method's restrictions on including donors that have undergone similar treatments. Anticipation effects do not seem to be a problem, given the in-time placebo tests and the goodness of fit up until treatment in 2008. The remaining challenges to internal validity are dealt with below.

Concerning possible violations of SUTVA, there is reason to believe that the inclusion of Armenia in the synthetic control unit causes a downward bias to the synthetic unit's outcome variable due to spill-over effects. As previously discussed, the spill-over affecting Armenia was mostly linked to the closure of trade routes running through breakaway regions. While the war thus caused immediate problems for Armenia, the country would have soon been able to reroute its trade and thereby mitigate losses. I believe that an upward bias to the estimated negative effect in 2008 is more of an issue than biases to the estimates post-2009, but an upward bias to the estimated effect is possible.

Another potential bias related to violations of SUTVA may result from the inclusion of Georgian trade partners in the synthetic control unit. Neither in the pre- nor post-war years is Georgia one of Armenia's major export destinations. Only about 6 % of Armenian exports headed to Georgia in the post-war period, on average (The Observatory of Economic Complexity, no date). In 2008, a decrease in Georgian demand would imply a downward bias to the synthetic unit in that year, meaning that the estimated negative effect is underestimated. After 2008, higher Georgian demand supported by increased fiscal spending would cause the synthetic unit's curve to be biased upward. The estimated positive effect would thus be underestimated, but the bias may be negligible.

Regarding the synthetic control method's restrictions on including units that have undergone similar treatments, the inclusion of Bosnia may bias results. However, the L-O-O test shows that the overall conclusions are not sensitive to the inclusion of Bosnia.

Furthermore, it should be mentioned that the effects of the 2008 financial crisis on Georgian GDP cannot be disentangled from the effects of the war. This is a concern when making causal interpretations of the estimated effects. It is possible that the model produces a good fit between unobserved time varying and unvarying characteristics in the pre-treatment years while failing to model a potential heterogenous impact of the global financial crisis. It has been argued that Georgia's economic reforms helped the country withstand the twin crises of 2008 (Emerson and Kovziridze, 2016). If reforms put Georgia's economy in a better position than its synthetic control unit to recover from the crises of 2008, irrespective of the donor pledges, then the estimated effect may not be solely attributable to events relating to the war. This would imply biases to estimated effects, because part of the effect would not result (directly or indirectly) from the war but from Georgia's idiosyncratic response to the financial crisis.

However, the model matches predictors well and produces a pre-treatment synthetic trajectory that closely tracks that of actual Georgia. The curves track each other closely in all pre-treatment years, also in those following the reform wave sparked by the Rose Revolution. As stated by Abadie, Diamond and Hainmueller (2010, 2015) a close fit over an extended period before treatment will help account for unobserved variables and heterogenous effects of both observed and unobserved factors. Whatever constitutes a 'extended period of time' is not clearly defined. But my model's close match in 1995-2007 favours the assumption that the impact of the financial crisis on Georgia is well replicated by the synthetic control unit. Under this assumption, disregarding the other possible causes of bias, the remaining differences between the synthetic control unit's outcome variable and actual Georgia's per capita GDP are the causal effects of the Russo-Georgian war.

In sum, the possibly present biases do not seem to override the overall conclusions drawn.

7.2 External validity

I argue in this paper that the results are conditional on Georgia's unique (geo)political situation in 2008. The model used and the results obtained in this case study do not immediately translate to other conflicts. My findings are, nevertheless, in part confirming those made by Collier and Dollar (2002) and Collier and Hoeffler (2004) in their general studies of the relationship between war and economy. Their results relating to foreign aid effectiveness in post-war situations are, according to the authors, best treated with caution. Still, they are similar to mine. In accordance with their quantitative analysis, the qualitative reasoning of this paper suggests that Georgia's successful post-war recovery relates to an interaction between policies, aid, and institutions. My case study thus provides support for more generally applicable findings on the relationship between war and economy.

7.3 Recommendations

There are outstanding questions relating to the results' sensitivity to the choice of baseline donor pool. The L-O-O tests shows that Armenia is an important driver of the estimated effect. When Armenia is removed, Latvia is instead given a large weight, which in turn gives results that support my overall conclusion of a positive treatment effect. However, the inclusion of Latvia in the L-O-O synthetic control unit may not produce reliable results, considering Latvia's severe economic downturn during the financial crisis. Further research on Russo-Georgian war could attempt at using a different donor pool that includes economically similar countries outside the former Eastern Bloc. This would give a better understanding of the results' sensitivity to the design of the baseline donor pool.

Furthermore, future research on the Russo-Georgian war could attempt at quantitatively estimating the causal impact of the individual channels possibly causing the effects identified in this paper. For instance, the implications of Georgia's varying levels of military spending and of changes to Georgia's political environment. This would give a more detailed picture of the causal channels driving Georgia's post-war recovery.

8 CONCLUSION

A gap in the literature on the economic consequences of the Russo-Georgian war was filled by measuring the war's aggregate impact. I asked: how did Russia's 2008 invasion of Georgia affect Georgian per capita GDP? Following the hypothesis that the effect was negative but negligible, the question was answered using the synthetic control method. Using a data set covering a range of economic predictors on a group of former eastern-bloc republics, Georgia's outcome in absence of the war was replicated. The chosen model gave a synthetic Georgia that closely resembled actual Georgia in both outcomes and

outcome predictors. The results show that Russia's invasion seems to have had a causal effect on Georgian per capita GDP. The divergence between actual and synthetic outcomes indicates that the war had an immediate negative effect followed by an economic development which made Georgia surpass its counterfactual. The magnitude and sign of the effect was unexpected. When considering individual years, a significant positive effect of the war is found in years 2012-2013. When considering the overall effect post-treatment, it is found that the war had a robust and positive causal effect in the post-war period.

The curious case of the Russo-Georgian war suggests that a war's effect on a country is conditional on its geopolitical context and the policies pursued post-bellum. My results reflect aggregate changes to the outcome variable, but causal channels cannot be disentangled. The estimated positive effect of the war may be explained by Georgia's westward policy alignments. International assistance likely helped Georgia overcome the costs of war. Another explanation not as clearly related to the war is that Georgia's post-war pivot towards the EU also helped. This reasoning is supported by previous studies on the relationship between aid, policies, and economic development in post-war economies. In all, this paper finds that the post-war years saw an accumulated net per capita GDP effect of $4314, an average benefit of about 7.5 % per year, which is significant at the 5 % level. Robustness checks support the results. Still, the findings of this paper are not free from internal validity concerns. Biases do not seem to override the overall conclusions drawn from my results, but there are outstanding questions relating to the channels driving the effect and to the results' sensitivity to the choice of baseline donor pool.

9 BIBLIOGRAPHY

Abadie, A. (2020) 'Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects'.

Abadie, A., Diamond, A. and Hainmueller, J. (2010) 'Synthetic control methods for comparative case studies: estimating the effect of California's tobacco control program', *Journal of the American Statistical Association*, 105(490), pp. 493–505.

Abadie, A., Diamond, A. and Hainmueller, J. (2015) 'Comparative politics and the cynthetic control method', *American Journal of Political Science*, 59(2), pp. 495–510.

Abadie, A. and Gardeazabal, J. (2003) 'The economic costs of conflict: a case study of the Basque Country', *The American Economic Review*, 93(1), pp. 113–132.

Adhikari, B. *et al.* (2018) 'Can reform waves turn the tide? Some case studies using the synthetic control method', *Open Economies Review*, 29(4), pp. 879–910. doi: 10.1007/s11079-018-9490-3.

Albalate, D. and Bel, G. (2020) 'Do government formation deadlocks really damage economic growth? Evidence from history's longest period of government formation impasse', *Governance*, 33(1), pp. 155–171. doi: https://doi.org/10.1111/gove.12410.

Arkhangelsky, D. *et al.* (2019) *Synthetic difference in differences*. w25532. National Bureau of Economic Research. doi: 10.3386/w25532.

Athey, S. and Imbens, G. W. (2017) 'The state of applied econometrics: causality and policy evaluation', *Journal of Economic Perspectives*, 31(2), pp. 3–32. doi: 10.1257/jep.31.2.3.

Avagyan, S. (2009) 'The global financial crisis: impact on Armenia', *Turkish Policy Quarterly*, 8(2), pp. 97–100.

Benoit, E. (1978) 'Growth and defense in developing countries', *Economic Development and Cultural Change*, 26(2), pp. 271–280.

Billmeier, A. and Nannicini, T. (2013) 'Assessing economic liberalization episodes: a synthetic control approach', *The Review of Economics and Statistics*, 95(3), pp. 983–1001.

Bluszcz, J. and Valente, M. (2020) 'The economic costs of hybrid wars: the case of Ukraine', *Defence and Peace Economics*, pp. 1–25. doi: 10.1080/10242694.2020.1791616.

Brück, T., De Groot, O. J. and Bozzoli, C. (2012) *How many bucks in a bang: on the estimation of the economic costs of conflict*. Oxford University Press. doi: 10.1093/oxfordhb/9780195392777.013.0012.

Brunnstrom, D. (2008) 'Donors pledge $4.5 billion for Georgia recovery', *Reuters*, 22 October. Available at: https://www.reuters.com/article/us-georgia-donors-idUSTRE49L2IB20081022 (Accessed: 24 March 2021).

Cappelen, Å., Gleditsch, N. P. and Bjerkholt, O. (1984) 'Military spending and economic growth in the OECD countries', *Journal of Peace Research*, 21(4), pp. 361–373.

Cavallo, E. *et al.* (2013) 'Catastrophic natural disasters and economic growth', *The Review of Economics and Statistics*, 95(5), pp. 1549–1561.

Collier, P. (1999) 'On the economic consequences of civil war', *Oxford Economic Papers*, 51(1), pp. 168–183.

Collier, P. and Dollar, D. (2002) 'Aid allocation and poverty reduction', *European Economic Review*, 46, pp. 1475–1500.

Collier, P. and Hoeffler, A. (2004) 'Aid, policy and growth in post-conflict societies', *European Economic Review*, 48(5), pp. 1125–1145. doi: 10.1016/j.euroecorev.2003.11.005.

De Groot, O. J. (2010) 'The spillover effects of conflict on economic growth in neighbouring countries in Africa', *Defence and Peace Economics*, 21(2), pp. 149–164. doi: 10.1080/10242690903570575.

Dunne, P. and Vougas, D. (1999) 'Military spending and economic growth in South Africa: a causal analysis', *The Journal of Conflict Resolution*, 43(4), pp. 521–537.

Echevarría, C. A. and García-Enríquez, J. (2019) 'The economic consequences of the Libyan spring: a synthetic control analysis', *Defence and Peace Economics*, 30(5), pp. 592–608. doi: 10.1080/10242694.2018.1446241.

Echevarría, C. A. and García-Enríquez, J. (2020) 'The economic cost of the Arab Spring: the case of the Egyptian revolution', *Empirical Economics*, 59(3), pp. 1453–1477. doi: 10.1007/s00181-019-01684-7.

Emerson, M. and Kovziridze, T. (2016) *Georgia and Europe: a short guide*. Brussels: Centre for European Policy Studies (CEPS). Available at: http://aei.pitt.edu/80164/1/Georgia_pop_edition_FINAL.pdf.

European Bank of Reconstruction and Development (no date) *How the EBRD assesses transition qualities*, *European Bank of Reconstruction and Development*. Available at: //www.ebrd.com/economic-research-and-data/transition-qualities-asses.html (Accessed: 7 March 2021).

Ferman, B. and Pinto, C. (2017) 'Placebo tests for synthetic controls', *MPRA Paper*, (78079). Available at: https://mpra.ub.uni-muenchen.de/78079/.

Firpo, S. and Possebom, V. (2018) 'Synthetic control method: inference, sensitivity analysis and confidence sets', *Journal of Causal Inference*, 6(2). doi: 10.1515/jci-2016-0026.

Freedom House (no date) *Countries and Territories*, *Freedom House*. Available at: https://freedomhouse.org/countries/freedom-world/scores (Accessed: 7 March 2021).

Galiani, S. and Quistorff, B. (2017) 'The Synth_runner package: utilities to automate synthetic control estimation using Synth', *The Stata Journal*, 17(4), pp. 834–849. doi: 10.1177/1536867X1801700404.

Gardeazabal, J. (2012) *Methods for measuring aggregate costs of conflict*. Oxford University Press. doi: 10.1093/oxfordhb/9780195392777.013.0011.

Gogolashvili, K. (2017) *EU-Georgia relations future perspectives*. Policy paper. Tbilisi. Available at: https://www.osgf.ge/files/2017/Publications/EU-Georgia_Relations_and_Future_Perspectives.pdf.

Hahn, J. and Shi, R. (2017) 'Synthetic control and inference', *Econometrics*, 5(4), p. 52. doi: http://dx.doi.org.ez.hhs.se/10.3390/econometrics5040052.

Hamilton, R. E. (2018) 'A ten-year retrospective on the Russia-Georgia war', *Foreign Policy Research Institute - Black Sea Strategy Papers*, p. 38.

International Monetary Fund (2008) *Press release: IMF executive board approves US$750 million stand-by arrangement for Georgia*, *International Monetary Fund*. Available at: https://www.imf.org/en/News/Articles/2015/09/14/01/49/pr08208 (Accessed: 7 April 2021).

International Monetary Fund and World Bank (2011) *Joint IMF/World Bank debt sustainability analysis*. Available at: https://www.imf.org/external/pubs/ft/dsa/pdf/dsacr11366.pdf (Accessed: 7 April 2021).

Kaul, A. *et al.* (2018) 'Synthetic control methods: never use all pre-intervention outcomes together with covariates'.

Khidasheli, T. (2011) 'Georgia's European way', *IPG*, 3, pp. 95–108.

King, G. and Zeng, L. (2006) 'The dangers of extreme counterfactuals', *Political Analysis*, 14(2), pp. 131–159.

Kofman, M. (2018) *The August war, ten years on: a retrospective on the Russo-Georgian war*, *War on the Rocks*. Available at: https://warontherocks.com/2018/08/the-august-war-ten-years-on-a-retrospective-on-the-russo-georgian-war/ (Accessed: 24 March 2021).

Koubi, V. (2005) 'War and economic performance', *Journal of Peace Research*, 42(1), pp. 67–82.

Lawson, R., Grier, K. and Absher, S. (2019) 'You say you want a (Rose) Revolution? The effects of Georgia's 2004 market reforms', *Economics of Transition and Institutional Change*, 27(1), pp. 301–323. doi: https://doi.org/10.1111/ecot.12205.

Marutyan, A. *et al.* (2010) *The 2009 annual report of the public debt of the Republic of Armenia*. Available at: https://www.minfin.am/website/images/files/2009.pdf.

McClelland, R. and Gault, S. (2017) 'The synthetic control method as a tool to understand state policy'.

Murdoch, J. C. and Sandler, T. (2004) 'Civil wars and economic growth: spatial dispersion', *American Journal of Political Science*, 48(1), pp. 138–151. doi: 10.2307/1519902.

Nordhaus, W. D. (2002) 'The economic consequences of a war with Iraq', *NBER working paper series*, (9361).

Olson, M. (1982) *The rise and decline of nations: economic growth, stagflation, and social rigidities*. New Haven: Yale University Press. Available at:

http://ebookcentral.proquest.com/lib/sciences-po/detail.action?docID=3420903 (Accessed: 28 February 2021).

Organski, A. F. K. and Kugler, J. (1977) 'The costs of major wars: the phoenix factor', *The American Political Science Review*, 71(4), pp. 1347–1366. doi: 10.2307/1961484.

Otarashvili, M. (2013) *Georgia and the global economic crisis*, *FPRI*. Available at: https://www.fpri.org/research/eurasia/recent-findings/georgia-global-econ-crisis/ (Accessed: 28 February 2021).

Remler, P. (2020) *Russia's stony path in the South Caucasus*, *Carnegie Endowment for International Peace*. Available at: https://carnegieendowment.org/2020/10/20/russia-s-stony-path-in-south-caucasus-pub-82993 (Accessed: 28 February 2021).

Reuters (2009) 'Factbox: facts about the 2008 war in Georgia', *Reuters*, 4 August. Available at: https://www.reuters.com/article/us-georgia-war-conflict-sb-idUSTRE5732TH20090804 (Accessed: 28 February 2021).

Rubin, D. B. (1974) 'Estimating causal effects of treatments in randomized and nonrandomized studies.', *Journal of Educational Psychology*, 66(5), pp. 688–701. doi: 10.1037/h0037350.

The Observatory of Economic Complexity (no date) *Georgia (GEO) exports, imports, and trade partners*, *The Observatory of Economic Complexity*. Available at: https://oec.world/en/profile/country/geo/?yearSelector2=importGrowthYear16 (Accessed: 4 March 2021).

Trading Economics (no date) *Armenia - credit rating*, *Trading Economics*. Available at: https://tradingeconomics.com/armenia/rating (Accessed: 8 April 2021).

Trading Economics (no date) *Armenia government spending*, *Trading Economics*. Available at: https://tradingeconomics.com/armenia/government-spending (Accessed: 7 April 2021).

Trading Economics (no date) *Georgia - credit rating*, *Trading Economics*. Available at: https://tradingeconomics.com/georgia/rating (Accessed: 8 April 2021).

Transparency International (2020) *Georgia's economic dependence on Russia: trends and threats*, *საერთაშორისო გამჭვირვალობა - საქართველო*. Available at: https://transparency.ge/en/blog/georgias-economic-dependence-russia-trends-and-threats (Accessed: 28 February 2021).

United Nations (2008) *Revised UN estimates show 192,000 uprooted during Georgia conflict*, *UN News*. Available at: https://news.un.org/en/story/2008/09/272702-revised-un-estimates-show-192000-uprooted-during-georgia-conflict (Accessed: 28 February 2021).

United Nations and World Bank (2010) *Georgia: donor funding in support of post-conflict recovery and reconstruction a second progress report*. Available at: https://www.mof.ge/common/get_doc.aspx?doc_id=8992.

Vendil Pallin, C. and Westerlund, F. (2009) 'Russia's war in Georgia: lessons and consequences', *Small Wars & Insurgencies*, 20(2), pp. 400–424. doi: 10.1080/09592310902975539.

World Bank (2018) *Georgia: from reformer to performer*. The World Bank Group.

World Bank (no date) *Georgia trade summary 2008*, *World Integrated Trade Solution*. Available at: https://wits.worldbank.org/CountryProfile/en/Country/GEO/Year/2008/Summarytext (Accessed: 28 February 2021).

World Bank (no date) *World Development Indicators*, *The World Bank*. Available at: https://databank.worldbank.org/source/world-development-indicators (Accessed: 7 March 2021).

'World Bank (2018)' (no date). Available at: https://openknowledge.worldbank.org/bitstream/handle/10986/29790/GEO-SCD-04-24-04272018.pdf?sequence=1&isAllowed=y (Accessed: 28 February 2021).

APPENDIX

Appendix A Map of Georgia covering Russia's offensive in 2008



Photo courtesy of Andrei Nacu

Appendix B Choosing country and variable weights

The synthetic control should be chosen so that it well replicates the treated unit. X_1 is a $(k \times 1)$ vector of the pre-treatment characteristics (ie predictors of the outcome variable) of the treated unit. X_1 contains the r observed covariates described in section 5.1 along with an M number of pre-intervention (ie lagged) outcomes used as predictors. Hence, $k = r + M$. X_0 is a $(k \times J)$ matrix containing the same predictors of the outcome variable, but for all the J countries of the donor pool.

The vector $X_1 - X_0W$ is the difference between the pre-treatment characteristics of the observed treated unit and the synthetic control. In order to obtain a good counterfactual, W^* is selected to minimise the Euclidean distance between X_1 and X_0W in the pre-treatment period:

$$\min \| X_1 - X_0W \|$$

subject to $0 \leq w_j \leq 1$ for $j = 2, \ldots, J+1$ and $\sum_{j=2}^{J+1} w_j = 1$.

Let V be a symmetric, diagonal matrix with k elements that are non-negative and sum to one ($v_1 + , \ldots , + v_k = 1$). The values of of v_1 through v_k represent the relative importance of the k variables in X_1. W^* is then selected to minimise:

$$\min \| X_1 - X_0W \|_V = \sqrt{(X_1 - X_0W)'V(X_1 - X_0W)}$$

Subject to

$$0 \leq w_j \leq 1 \text{ for } j = 2, \ldots, J+1$$

$$w_2 + \ldots + w_{J+1} = 1.$$

The choice of W^* is thus conditional on the variable weights in V. In the literature, this minimisation problem has been solved in different ways (see, for instance, Abadie and Gardeazabal, 2003; Abadie, Diamond and Hainmueller, 2015; Firpo and Possebom, 2018; Ferman and Pinto, 2017). One way is to subjectively choose the variable weights of V and then choose a W^* that minimises the problem above. However, this would require extensive prior knowledge of the determinants of the outcome variable. Instead, this paper follows the data-driven, nested double-minimisation approach proposed by Abadie and Gardeazabal (2003). The nested double-minimisation approach has proven itself better

than the alternatives in terms of computational time and accuracy (Echevarria & Garcia-Enriquez, 2020). In this paper, the variable weights of matrix V is obtained together with the country weight vector W so that both the problem above and the root mean squared prediction error (RMPSE) are jointly minimised (ibid). The root mean square prediction error is:

$$Pretreatment\ RMSPE \equiv \sqrt{\sum_{t=1}^{T_0}\left(Y_{1t} - \sum_{j=2}^{J+1} Y_{jt}\, w_j(V)\right)^2 / T_0}$$

In practice, the optimal weights for W and V are obtained in STATA by using the `synth` and `synth_runner` packages supplied by Abadie, Diamond and Hainmueller (2010) and Galiani and Quistorff (2017) respectively. If pre-intervention predictors of the treated unit have been closely replicated by the obtained weights W^* and V^*, then the counterfactual outcome Y_{1t}^N in the post-treatment period should also be well replicated and the treatment effect is obtained by:

$$\hat{\alpha}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} w_j^*\, Y_{jt}, \text{ for } t = T_0 + 1, \ldots, T.$$

Appendix C Calculating p-values[6]

If an estimated effect for the unit of interest at a given post-treatment time period is $\hat{\alpha}_{1t}$, then the distribution of in-space placebo effects is $\hat{\alpha}_{jt}^{PL} = \{\hat{\alpha}_{jt} : j \neq 1\}$. Using this paper's first inference method, two-sided p-values are calculated as being:

$$p_{1,two-sided} \coloneqq Pr(|\hat{\alpha}_{1t}^{PL}| \geq |\hat{\alpha}_{1t}|) = \frac{\sum_{j \neq 1} 1\left(\left|\hat{\alpha}_{jt}^{PL}\right| \geq |\hat{\alpha}_{1t}|\right)}{J}.$$

[6] This section is to a large extent based on Galiani and Guistorff (2017).

For one-sided p-values, the p-value for estimated positive effects is:

$$p_{1,one-sided} \coloneqq Pr(\hat{\alpha}_{1t}^{PL} \geq \hat{\alpha}_{1t})$$

For negative effects the inequality sign is simply reversed (Galiani and Quistorff, 2017). Galiani and Quistorff (2017) note that estimated placebo effects could be large simply due to the fit being bad. In other words, the effects estimated in the post-treatment period may be due to the synthetic units' inability to closely track the observed unit rather than due to treatment. In such cases, the p-values obtained may be too conservative. Their software therefore allows researchers to control for the pre-treatment goodness of fit. This can be done by calculating standardised effects by dividing estimated effects by the pre-treatment goodness of fit. This renders in the standardised p-values that are used throughout this paper. One may then also want to restrict the comparison set of units to those that have an RMSPE (goodness of fit) at most *m* times that of the unit of interest. In this paper, the benchmark multiple *m* is set to five.

The second inference method is one used by Abadie, Diamond and Hainmueller (2015). The authors note that effects may be large for some post-treatment periods, but not for all. Then, rejection of the null hypothesis of no effect is not always clear-cut. Instead, the authors recommend SCM practitioners to consider the empirical distribution of the ratio between post- and pre-treatment RMSPE-values. A p-value is calculated as the number of units having a greater than or equal ratio between post- and pre-treatment RMSPE divided by the number of control units:

$$\text{RMSPE_ratio}_j \coloneqq \frac{\sum_{T_0+1}^{T}(Y_{jt} - \hat{Y_{jt}^N})^2/(T - T_0)}{\sum_{t=1}^{T_0}(Y_{jt} - \hat{Y_{jt}^N})^2/T_0}$$

$$p_2 := \frac{\sum_{j\neq 1}^{J+1} \mathbb{I}\left[RMSPE_ratio_j \geq RMSPE_ratio_1\right]}{J},$$

where $\mathbb{I}[\ldots]$ is an indicator function of the described event (Abadie, Diamond and Hainmueller, 2015; Firpo & Possebom, 2018). Intuitively, the credibility of an estimated causal effect is low when relatively many untreated units exhibit a similar effect when assigned treatment status (Abadie, Diamond and Hainmueller, 2015).

Appendix D Criticism of inference methods

The inference methods based on placebo tests proposed by Abadie, Diamond and Hainmueller (2010, 2015) have faced scrutiny and criticism by other researchers. Arkhangelsky et al write: '[placebo tests are] closely connected to permutation tests in randomization inference. However, in many applications of synthetic control, the exposed unit was not chosen at random, in which case placebo tests do not have the formal properties of randomization tests, and so may need to be interpreted in a more qualitative lens' (2019, 22).

Hahn and Shi (2017) acknowledge that the use of such permutation tests is intuitive, but underline that in placebo runs, pre-treatment trends are often tracked poorly. While the applicability of the SCM method to the treated unit requires that the synthetic control unit has a good fit, this condition seems to be relaxed when it comes to iteration of the test on donor pool units (Hahn and Shi, 2017). However, by excluding placebo runs with particularly bad pre-treatment fits, the inference test discussed above can still provide some useful information in evaluating the likelihood of a treatment effect being non-zero. However, the points made by Hahn and Shi (2017) are important to be aware of.

Using a Monte Carlo experiment, Firpo and Possebom (2018) find that the original RMSPE-ratio inference procedure employed by Abadie, Diamond and Hainmueller (2015) performs well compared with other inference methods when there is only one treated unit. Firpo & Possebom (2018) also extend the method introduced to SCM by Abadie, Diamond, and Hainmueller by using parametric weights in calculations of p-values and by developing a method for testing any sharp null hypothesis (ie not only the traditional zero-effect null of Abadie, Diamond, and Hainmueller).

Appendix E Per capita GDP graphs

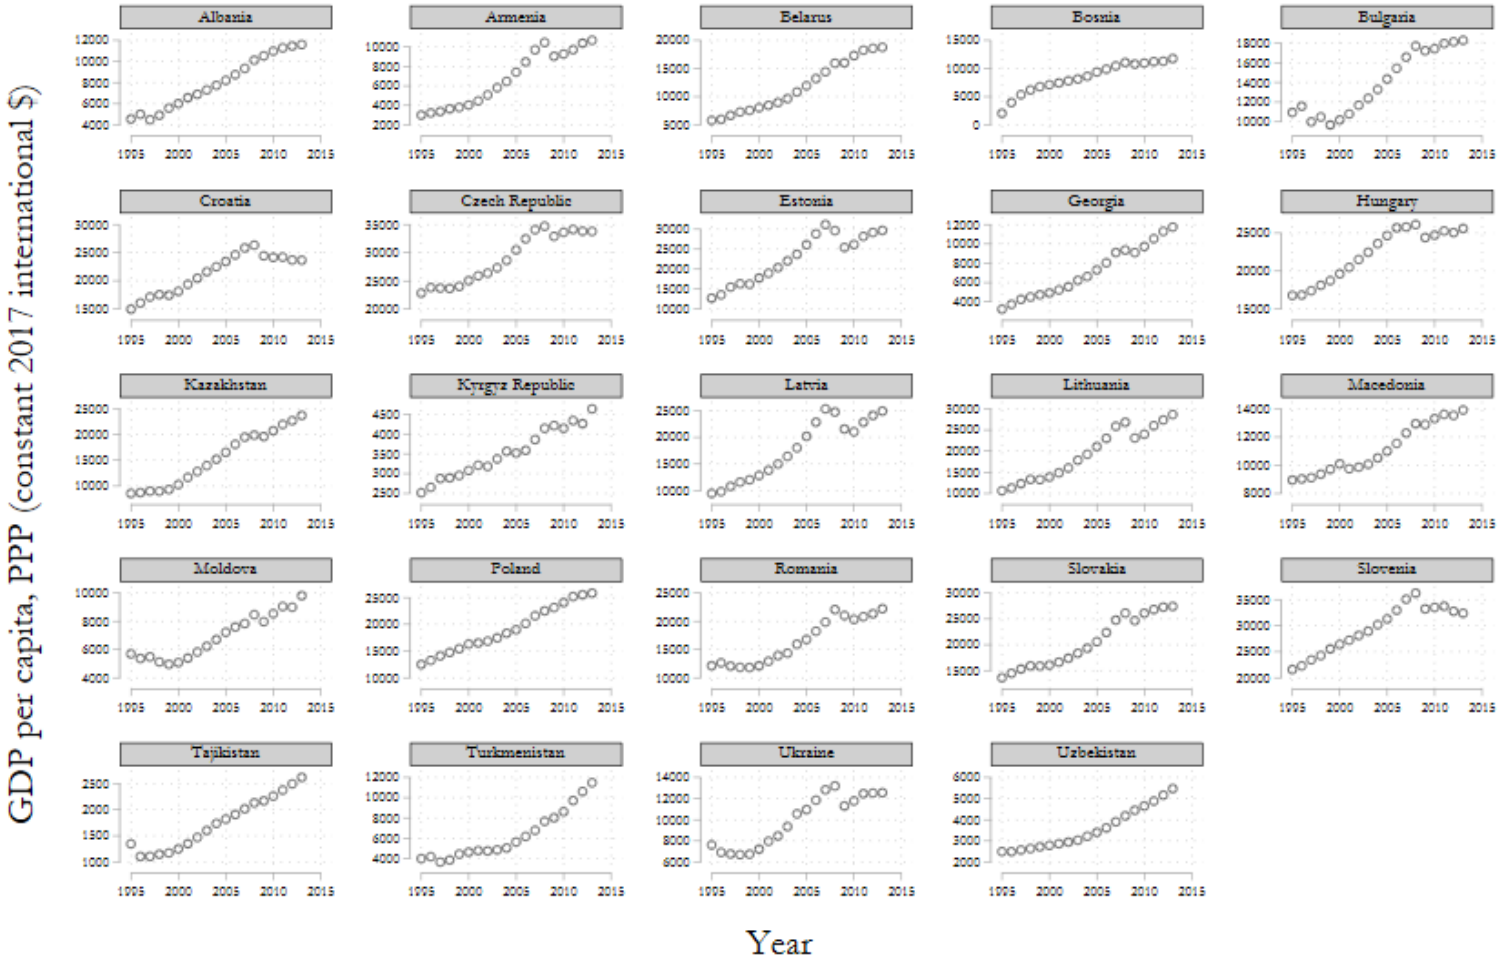


Figure E: graphs over per capita GDP levels in the sampled period among the sampled units. Source: my rendering of World Bank data (no date)

Appendix F In-time placebo tests

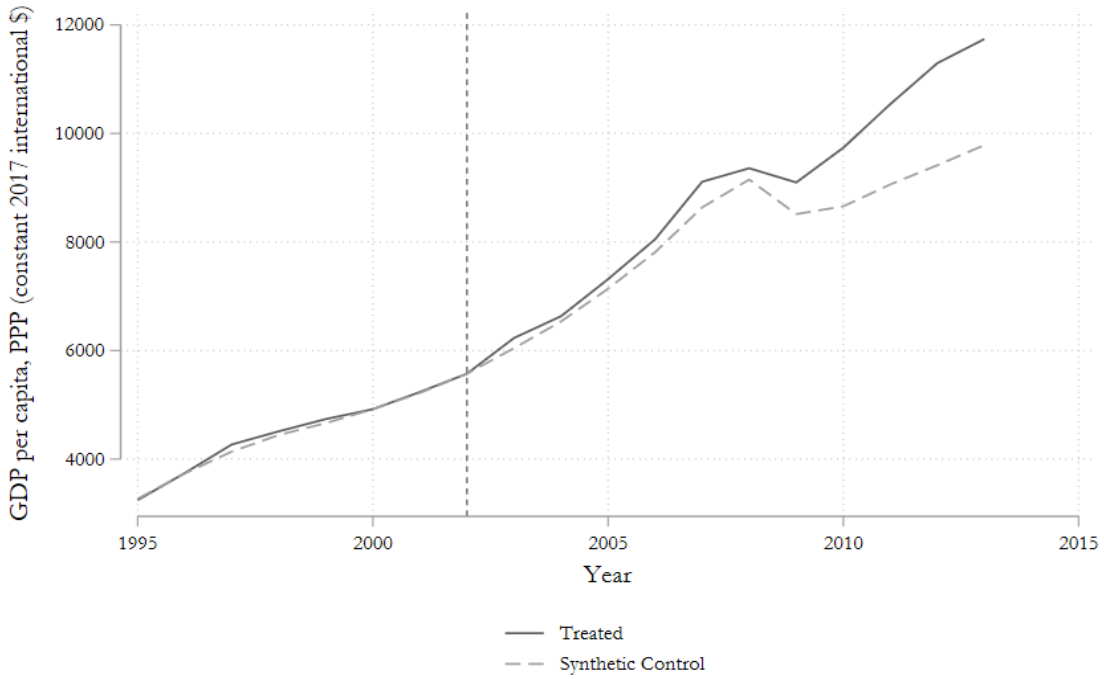


Figure F1: in-time placebo test (moving treatment to 2003)

Table F1: effects and standardised p-values

	Estimated effects	Standardised p-values
2003	187.374	.438
2004	96.801	.813
2005	175.687	.75
2006	238.852	.813
2007	472.406	.688
2008	207.967	1
2009	583.936	.875
2010	1080.006	.563
2011	1480.807	.5
2012	1881.689	.5
2013	1954.555	.563

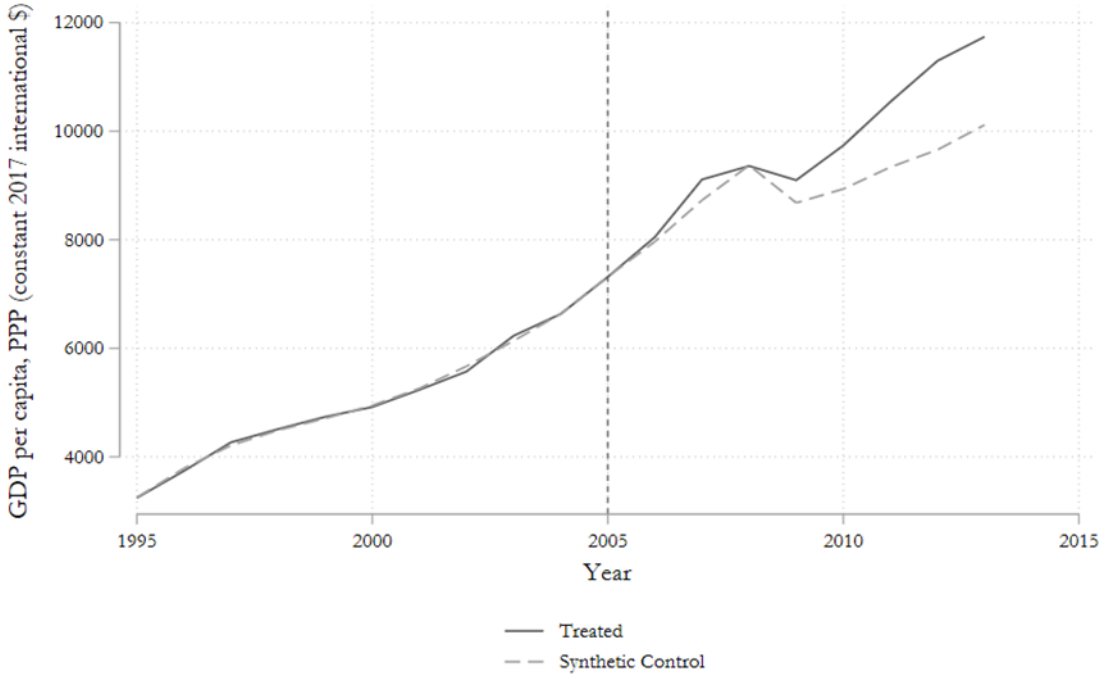


Figure F2: in-time placebo test (moving treatment to 2006)

Table F2: effects and standardised p-values

	Estimated effects	Standardised p-values
2006	82.81627	.6
2007	381.9899	.2
2008	-12.14278	1
2009	418.0049	.4
2010	800.68	.4
2011	1204.124	.2
2012	1638.419	0
2013	1623.25	0

Appendix G Benchmark model output

Model specification

Lagged variables	`pcGDP(2005), pcGDP(2000)`
Covariates	`FDIshare, Urban, GCFshare, LifeExp, MilExp, Inflation, Exports`
Treatment period	2008
Results period	1995-2013

Table G1: treated and synthetic unit averages in pre-treatment period

Variable	Weight	Treated unit	Synthetic control unit
pcGDP(2005)	.00069567	7315.198	7404.301
pcGDP(2000)	.98911292	4919.232	4917.480
FDIshare	1.737e-06	8.038	5.235
Urban	.00761916	53.254	53.237

GCFshare	.00256516	25.931	25.927
LifeExp	1.260e-07	70.100	70.794
MilExp	9.901e-07	2.337	2.543
Inflation	2.251e-06	21.7416	26.146
Exports	1.986e-06	25.059	30.276
			RMSPE 79.550

Table G2: synthetic control donor weights

Albania	0	Lithuania	0
Armenia	0.469	Macedonia	0
Belarus	0	Moldova	0.05
Bosnia	0.222	Poland	0
Bulgaria	0.041	Romania	0
Croatia	0	Slovakia	0
Czech Republic	0	Slovenia	0
Estonia	0	Tajikistan	0
Hungary	0	Turkmenistan	0.06
Kazakhstan	0	Ukraine	0
Kyrgyzstan	0.158	Uzbekistan	0
Latvia	0		

Table G3: estimated effects and standardised p-values

	Estimated effects	Standardised p-values
2008	-270.516	0.333
2009	207.134	0.778
2010	633.7694	0.444
2011	1036.187	0.111
2012	1411.501	0
2013	1455.851	0.111

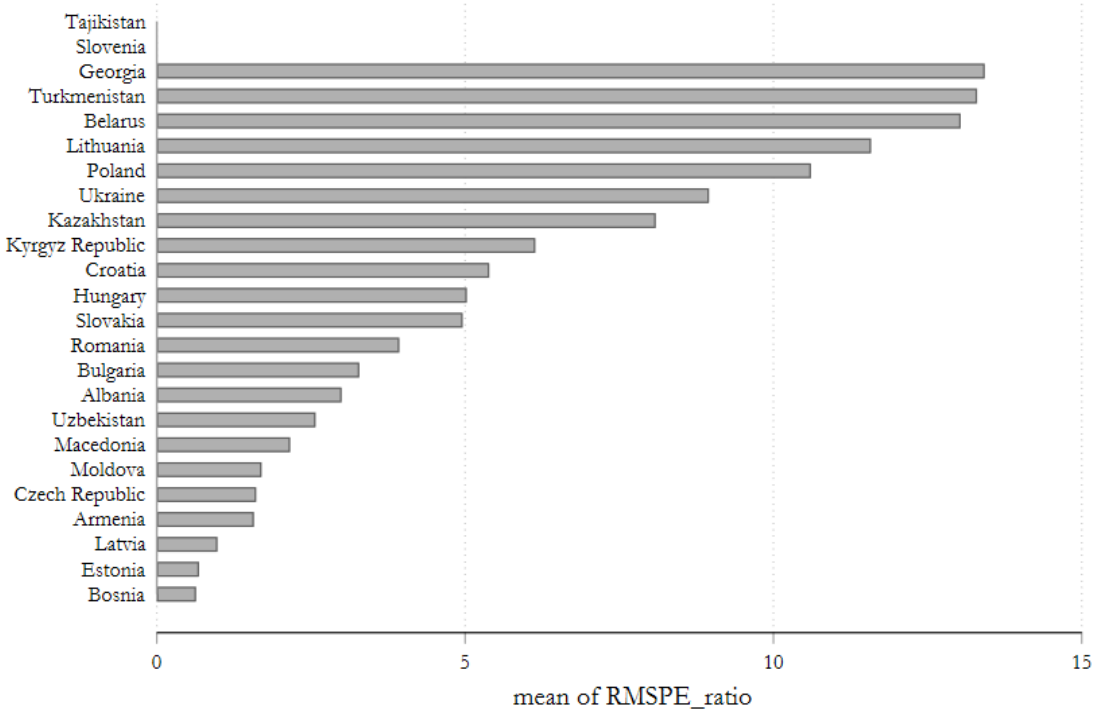


Figure G1: RMSPE-ratio distribution

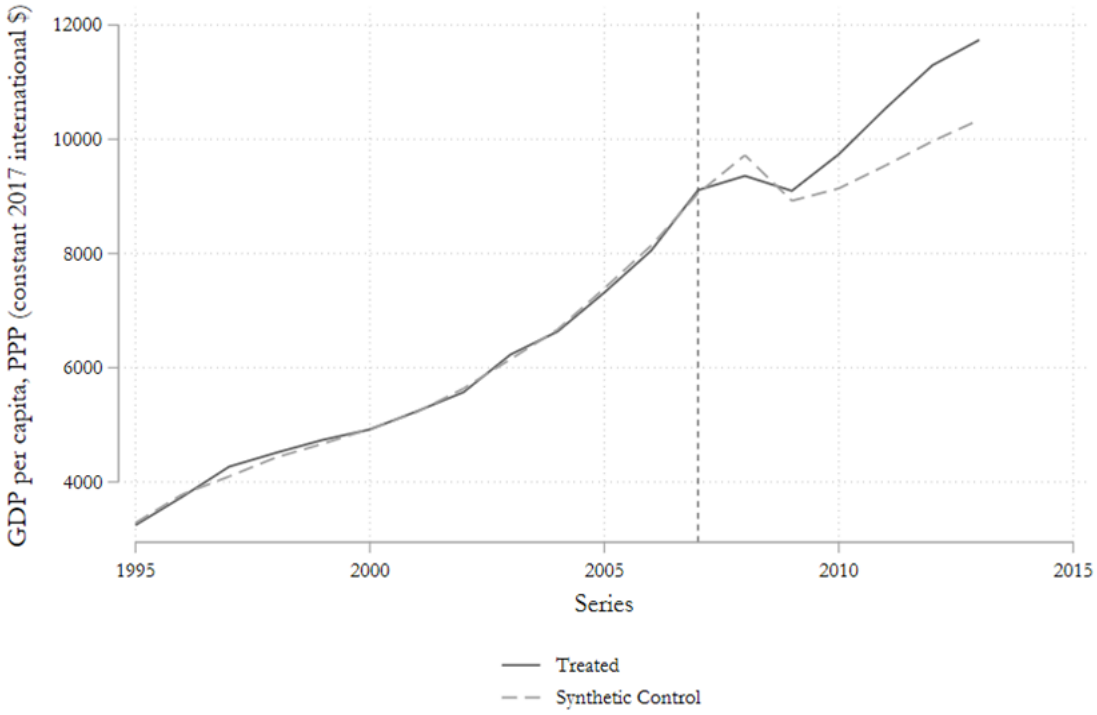


Figure G2: treated and synthetic Georgia