

Aidt, Toke; Asatryan, Zareh; Badalyan, Lusine

Working Paper

Political consequences of (consumer) debt relief

ZEW Discussion Papers, No. 24-030

Provided in Cooperation with:

ZEW - Leibniz Centre for European Economic Research

Suggested Citation: Aidt, Toke; Asatryan, Zareh; Badalyan, Lusine (2024) : Political consequences of (consumer) debt relief, ZEW Discussion Papers, No. 24-030, ZEW - Leibniz-Zentrum für Europäische Wirtschaftsforschung, Mannheim

This Version is available at:

<https://hdl.handle.net/10419/300009>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

// NO.24-030 | 05/2024

DISCUSSION PAPER

// TOKE AIDT, ZAREH ASATRYAN,
AND LUSINE BADALYAN

Political Consequences of (Consumer) Debt Relief

Political Consequences of (Consumer) Debt Relief

TOKE AIDT* ZAREH ASATRYAN† LUSINE BADALYAN‡

First version: September 28, 2022

This version: March 6, 2024

Many governments operate consumer debt relief programs, often timed to match the election cycle, but their political effects are not well understood. We ask if debt relief can influence elections in democracies. Our motivating exercise is the Biden administration's promise to relieve student debt. We utilize quasi-experimental variation generated by another very large debt relief program enacted in the Republic of Georgia that, similar to USA, affected every sixth voter. We estimate that the program helped the incumbent candidate win that election, and that its effects persisted. Overall, we show how economic power can translate into political power in democracies.

Keywords: Consumer debt relief, US student debt, distributive politics, vote buying, elections.

JEL codes: G51, D72.

*University of Cambridge & CESifo; tsa23@econ.cam.ac.uk

†ZEW & CESifo; asatryan@zew.de

‡Frankfurt School of Finance & Management; l.badalyan@fs.de

We thank Salome Baslandze, Carlo Birkholz, Ruben Enikolopov, Friedrich Heinemann, David Gomtsyan, Tommy Krieger, Jean Lacroix, and Vardges Levonyan for comments. We thank Peter Buchmann, Felix Köhler and David Westerheide for outstanding research assistance, as well as Heghine Manasyan and Koba Turmanidze for help in using the CRRC survey data.

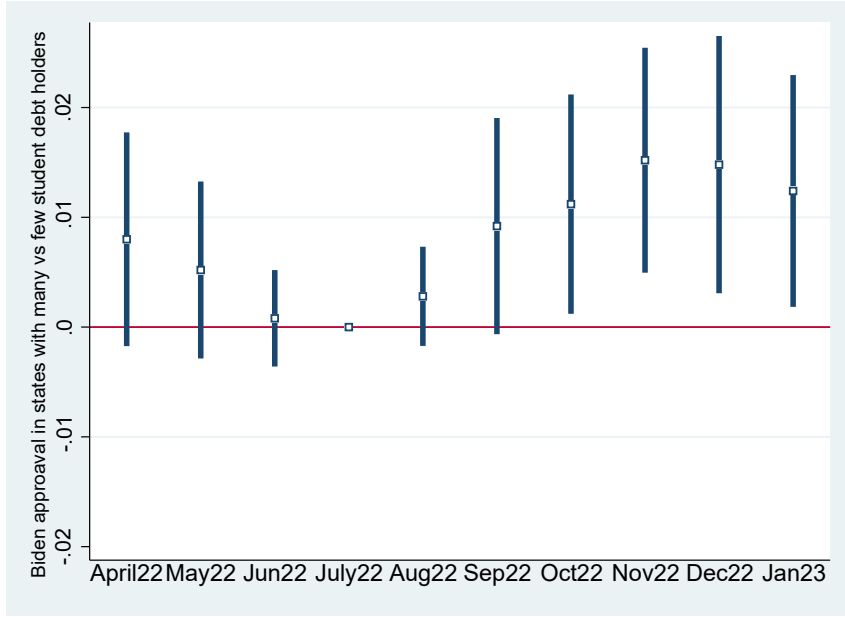
1 Introduction

Many governments operate debt relief programs that give consumers bankruptcy protection or write off consumer debt in times of crisis. The economic trade-offs associated with such programs are well-understood: debt relief helps individuals in financial trouble and improves the overall financial health of the economy, but at the cost of creating moral hazard and diverting public funds from other purposes or adding to public debt (e.g., Dobbie et al. 2017; Karlan et al. 2019; Exler and Tertilt 2020). The political consequences, in contrast, are not well-understood. Yet, debt relief is highly political and shares many features with the traditional tools of distributive politics routinely used by political parties to sway or mobilize voters to support them: debt relief bestows a (private) benefit to those who get it and the introduction of new or reformed debt relief programs can be timed to match the election cycle (e.g., Faraz and Rockmore 2020; Delatte et al. 2023). This paper studies the political consequences of debt relief. We ask whether debt relief can influence the outcome of high-stake elections, and, if so, what does it cost to win?

To answer these general questions, we consider two particular cases. The primary case is a large-scale debt relief policy enacted recently in the former Soviet Republic of Georgia, where we can use unique data on the beneficiaries of debt relief to estimate the causal effect on election outcomes. The second case provides motivating evidence that debt relief can have political effects from the well-known problem of student debt overhang in the USA.

The cumulative value of US student debt is about \$1.6 trillion or about 7% of US GDP and the question of relieving (some) federal student debt became salient during the 2020 presidential campaign. Several democratic candidates endorsed generous relief programs during the primaries, in the end leading Joe Biden to also propose forgiving \$10,000 of debt per student. The policy was announced officially by Biden in August 2022, a few months before the mid-term elections. It was estimated that it, if implemented, would

Figure 1: US student debt relief policy and Biden’s approval ratings



Notes: The figure shows the estimated differences in President Biden’s approval ratings between states that have above-median number of student debt holders and states that have below-median student debt holders (y-axis) in the months before and after announcing the student debt relief policy in August, 2022 (x-axis). The difference-in-differences estimation equation is shown in Appendix A as equation (3). We control for state-level GDP, the unemployment rate, average household income, and the education rate. The reported point estimates are relative to July. The bars represent 95% confidence intervals. The details are discussed in Appendix A.

have benefited 43 million Americans or about every sixth voter.¹ While existing research have studied the economic effects of student debt and debt relief, such as the effects on occupational choice, labor market outcomes, or financial decisions (Daniels Jr and Smythe 2019; Luo and Mongey 2019; Rothstein and Rouse 2011; Field 2009; Folch and Mazzone 2022), the electoral consequences have not been studied yet. To see if the proposed debt relief policy improved Biden’s popularity, we compare his approval rating (as measured monthly in online polls) before and after the announcement in states that had above-median number of student debt holders (“treated states”) and states that had below-median student debt holders (“control states”). Figure 1 reports the difference-in-difference estimates for the six months after the announcement (along with pre-trends). The policy announcement appears to have benefited Biden politically. After the announcement of the debt relief plan, his approval ratings increase by about 1.5

¹ See White House briefing (accessed September 2022).

percentage points in states that had many debt holders relative to states with few.² This political effect is all the more remarkable when it is taken into account that the plan was only a prospective promise of relief, and a non-credible one, as it turned out later, when the US Supreme Court in June 2023 ruled that the plan was unconstitutional. Yet, it is also clear that we cannot claim that this is a causal effect and even if the estimates in Figure 1 were causal, showing that the policy effects the presidential approval rating is not the same as showing that it has a material effect in a high-stakes election. To do that we need to turn to our primary case: the 2018 presidential election in the Republic of Georgia.

Between the two rounds of the 2018 presidential election in Georgia, a major debt relief program was announced, which got implemented the following year. Neither of the two top candidates, supported by the incumbent Georgian Dream party and the opposition United National Movement, obtained an absolute majority in the first round, leading to a second round race between them to be held after a month. In between the two rounds, the incumbent party announced that a deal had been negotiated with Georgian banks and non-bank lenders. The deal, as announced, would write off bad loans, including non-performing ones, smaller than \$770 for about 600,000 voters which, similar to the case of US student debt relief, would benefit about one-sixth of the electorate. The gross book value of the program was announced to be \$578 million, which is about 3.6% of Georgia's GDP. The debt relief program was funded by a private foundation linked to the incumbent (Georgian Dream) party. Subsequently, the names and the amount of debt that was written off for each of the beneficiaries were published on the internet. We collected and geo-located a sample of about 20,000 beneficiaries. This enables us to quantify the scale of the debt relief program across election districts.

We use a difference-in-differences strategy to estimate the impact of the program on the outcome of the 2018 presidential election. We find that a 10% increase in debt relief leads to a 7% increase in the vote for the Georgian Dream supported candidate in the

² We discuss the US case in more detail in Appendix A.

second round relative to the first round. The program did not increase aggregate turnout. The effects were stronger in poorer districts, and they persisted into the parliamentary election that took place two years later. We use individual level survey data to show that voters who likely benefited from the program report a positive attitude to it and are more likely to support the Georgian Dream party and to consider that the election was conducted fairly. Based on our estimates, a counter-factual exercise suggests that \$13.6 million would have been sufficient to swing the election. The cost of a single vote was \$81.6.

Our paper is related to four strands of literature. First, we contribute to the literature on consumer debt relief. In addition to the literature related to US student debt relief discussed above, this literature shows how bankruptcy protection regulations in mature economies have positive effects on financial and labour market outcomes and reduce foreclosure (Dobbie and Song 2015; Dobbie et al. 2017; Agarwal et al. 2017). The evidence from emerging economies, such as the field experiments conducted in India and the Philippines waiving small amounts of debt, highlights that debt relief programs create moral hazard (Kanz 2016; Mukherjee et al. 2018; Karlan et al. 2019), and Giné and Kanz (2018) finds no offsetting positive effects on productivity, wages, or consumption. The focus in this literature is on the economic effects of debt relief, and our contribution is to show that the debt relief programs can have electoral effects and that these effects can be pivotal in determining high-stakes election outcomes. Debt relief constitutes a direct benefit bestowed by politicians to those who have their debt forgiven, and, thus, discretionary debt relief programs have the potential to become a tool of distributive politics.

Second, our paper is related to a growing literature on political credit cycles. The literature has demonstrated robust evidence of election cycles in the supply of credit on both national and local levels (e.g., Kern and Amri 2021). These cycles can be achieved through the relaxation of macroprudential regulations in the run-up to elections (Müller 2023) or through the influence of politicians on state-owned or politically-connected banks

to expand credit supply during elections and re-allocate it to electorally important districts (for evidence from several countries, see Cole 2009; Carvalho 2014; Englmaier and Stowasser 2017; Faraz and Rockmore 2020; Delatte et al. 2023). The literature has often found that the share of non-performing loans increases in the aftermath of election-induced credit expansion (Cole 2009; Kern and Amri 2021). Our contribution to the literature is to show that debt relief timed around elections can help incumbents win elections. In this way, we add to the evidence showing that credit market conditions have political effects (Doerr et al. 2022; Gyöngyösi and Verner 2022; Funke et al. 2016).

Third, the paper is related to the literature on distributive politics. This literature argues that politicians promise distributive policies to sway voters to deviate from their ideological preference or to induce them to turn out to vote. Theoretically, this can be aimed at swing voters without strong ideological attachments (Lindbeck and Weibull 1987) or at core supporters (Cox and McCubbins 1986; Dixit and Londregan 1996). The evidence of this is overwhelming (Golden and Min 2013).³ Unlike this literature on distributive politics that studies retrospective rewards (e.g., disaster aid or cash transfers received prior to an election), the benefits are prospective in our case. Consequently, the debt relief is a promise of a transfer as in theoretical models of distributive politics, such as Lindbeck and Weibull (1987), which may be realized after elections (as in the case of Georgian debt relief) or not even if the incumbent is re-elected (as in the case of the US student debt relief). Our focus on prospective rewards allows us to investigate if such promises are sufficient to generate consequential voter reactions and if so, if they persist into subsequent elections. Another important difference compared to the literature on distributive politics is that the program we study in Georgia is financed and implemented by a private foundation, while the past literature, similar to the US student debt relief, is

³ A number of recent papers, including Bechtel and Hainmueller (2011) on Germany, Manacorda et al. (2011) on Uruguay, Pop-Eleches and Pop-Eleches (2012) on Romania, De La O (2013) on Mexico, Zucco Jr (2013) on Brazil, Vannutelli (2023) on Italy, report causal evidence that incumbents that implemented targeted welfare transfer programs of various sorts benefit from them electorally.

always about tax-funded government programs. This helps us understand how economic power can translate into political power through a privately funded transfer scheme.

Fourth, our study is also related to a small literature on vote buying. Debt relief shares many of the characteristics of what Stokes (2005) refers to as patronage goods. The literature on vote buying views such goods as contingent deals where a private benefit is exchanged for a vote. Vote buying is, typically, a hidden activity and not legal. A new literature on forensic economics (Zitzewitz 2012) has made progress in casually identifying instances of electoral fraud from the observable consequences of vote buying, such as anomalies in monetary cycles and household consumption patterns around elections (Mitra et al. 2021; Aidt et al. 2020). We relate to this literature by quantifying the geographical intensity of a particular type of vote buying using unique data on individual beneficiaries of the program.

Although our main results about the electoral consequences of debt relief apply to the particular cases of consumer debt relief in Georgia, we can find indications of similar effects in relation to student debt relief in USA and we note that the problem of bad debt and, consequently, the demand for debt relief is very common. According to data from the World Bank's World Development Indicators, in 2019 (before the pandemic) 6.6% of all loans were non-performing in an average country, with the percentage reaching 37% in Greece and 51% in Ukraine for example. Countries use various policies to write off consumer debt. A notable example is the consumer bankruptcy protection legislation of the USA. It is essentially a large social insurance program, larger than, for example, all state unemployment insurance programs combined (see, e.g., Dobbie and Song 2015). New debt relief programs are often enacted as a result of economic crises. For example, during the COVID-19 crisis all but 13 of the 185 countries followed by the Oxford Covid-19 Government Response Tracker introduced some form of debt relief by the end of 2021.⁴ Other recent examples include debt relief of \$14.4 billion in India following the 2008-09

⁴ See Oxford Covid-19 tracker (accessed December 2021).

crisis, or of \$10 billion in Brazil and \$2.9 billion in Thailand (Kanz 2016; Giné and Kanz 2018). The design of these programs is often highly political.

2 Background

Georgia gained independence from the Soviet Union in 1991. It has a population of about 3.7 million and is classified by the World Bank as a middle income country with a GDP per capita of \$4,289 and with a poverty head count of 20% of the population in 2019 (World Bank 2022). De jure, Georgia is a democratic republic. De facto, Freedom House classifies Georgia as a “partial free” country with scores similar to, for example, Indonesia or Mexico.

Georgia’s transition to democracy took off in 2003 following the so-called Rose Revolution. It led to the resignation of Eduard Shevardnadze – a former secretary of state of the Soviet Union and president since 1995 – and the election of Mikheil Saakashvili as president representing the United National Movement on a liberal, pro-European platform. The United National Movement’s hold on power was challenged in 2012 by the new party, Georgian Dream. The party was launched by the Russia-based billionaire businessman Bidzina Ivanishvili and it won the parliamentary election in 2012, with Ivanishvili as the prime minister, and has held power since.

Today, the prime minister is the head of the government and the president is the head of the state. This arrangement is a result of the transition from a presidential to a parliamentary system following a series of constitutional amendments passed between 2013 and 2018. These changes have curtailed the president’s executive powers in favor of the prime minister. Yet, the president still holds significant de facto power by, for example, being the commander-in-chief of the army, and because of the traditionally powerful role of the president in Georgia. The 2018 presidential elections was the last to elect a president by direct vote for a fixed six-year term. Our focus is on this crucial election. The two main candidates running for president were Grigol Vashadze backed

by the United National Movement and Salome Zurbashvili supported by the Georgian Dream party. Table D1 of the supplementary Appendix shows the results of the election. None of them obtained a majority in the first round – each getting 37-38% of the total vote – which meant that they had to compete in a run-off election. Despite entering the second round neck-to-neck, Zurbashvili almost doubled her votes and became the first female president of Georgia with a majority of 59.5%.⁵

In between the two rounds (held at the end of October and November, respectively), on November 19, prime minister Mamuka Bakhtadze from the Georgian Dream party announced that a deal had been negotiated with banks and non-banking lenders to buy a very large portfolio of bad loans.⁶ He announced that this would lead to a debt write off of loans below a threshold of GEL 2,000 (\$770) for more than 600,000 individuals. This corresponds to about one-sixth of the electorate.⁷ The reported book value of the deal was GEL 1.5 billion (\$578 million).⁸ In the announcement, Bakhtadze thanked the Cartu International Charity Foundation for funding the program. The Cartu Foundation was established in 1995 by Bidzina Ivanishvili and his family is its only donor. In December 2018, the webpage “vali.ge” (“vali” translates as debt from Georgian) opened and published information about who had their debt cancelled, the book value of that debt, the name of the lending institution, etc. We refer to this as the debt relief program.

⁵ Turnout in the second round increased by about 300,000. The other candidates running in the first round got about 432,000 votes in total and the three with the highest totals (accounting for 271,000 votes) announced that they would support Grigol Vashadze in the second round.

⁶ As reported by civil.ge available here.

⁷ The exact definition of what qualifies as a “bad” loan was not specified. We suspect that the program wrote off a range of such loans, ranging from fully non-performing loans to loans which were not in default but where payments were missed. According the World Bank’s World Development Indicators, bank non-performing loans accounted for 2.27% of the value of total gross loans in Georgia at the time.

⁸ We do not observe the program’s actual cost to the Cartu Foundation, and the foundation refused to disclose it, arguing that it was a commercial secret. Since the portfolio that the foundation bought comprised of bad loans, it is likely that it paid only a fraction of the book value. Some estimates from the Euro area (link (accessed March 2024)), suggest that portfolios of non-performing loans often trade in secondary markets at 20 to 40% of the book value.

3 Data and empirical strategy

We collected individual level data on who got debt relief and how much. The source is the program’s official website vali.ge, which was operational from early December 2018 to the end of December 2019. On this website, individuals could look up whether they would receive debt relief from the program using their unique eleven-digit social security IDs. We stratified on the eleven-digit IDs, to obtain a sample of the population of the listed individuals (IDs) that is representative at the level of election districts.⁹ The details of our sampling strategy are discussed in supplementary Appendix B.

Our sample comprises of 36,564 individual loans granted to 19,937 unique individuals. The loans in our sample total GEL 235 million, which corresponds to about 15.6% of the debt that was announced to be written off. Although the debt relief was supposed to be capped at GEL 2,000, in our data about a third of loans are larger than that threshold.¹⁰ The average (median) loan equals GEL 3380 (1090). The portfolio of loans was bought by the Cartu Foundation from 78 different Georgian banks and other financial institutions.

While we observe the book value of the loans written off for each individual in the sample, we do not observe how people voted in the two rounds. Thus, we cannot use individuals as the unit of analysis. Our analysis is, therefore, based on data aggregated to the 73 election districts which are the principal geographical units at which the Georgian Election Commission reports vote total for candidates and turnout. We count the individuals (based on the first two digits of their ID) who benefited from the program and

⁹ Scraping the full data was impossible given that the website was operational for less than a year.

¹⁰ We do not know for sure why this is the case. A plausible reason is that the published values on the program’s website included interests, penalties and other fees which were added to the principal of each loan. It is also possible that the program did not, in practice, strictly comply with the GEL 2,000 threshold. The speedy negotiations between the Cartu Foundation and the many lenders did not leave time to bargain over individual loans; instead the bargaining was over large tranches of loans which may have included loans above the threshold. Figures D1(a)-(b) in the supplementary Appendix show the size distribution of the loans and the mean values by cohort and gender.

calculate the total value of the debt that was written off in each district.¹¹ In addition to the districts, the Georgian Election Commission also reports election results on the level of 3,600 precincts (or polling stations) which on average have about 1,000 voters. To increase the sample size, we also aggregate individual level data to these units. The larger sample comes at a cost, however: we can match only 85% of individuals to precincts¹² and our sampling strategy was not designed to make the sample representative for these units.

One limitation of working with aggregated data is that they cannot, in general, tell us how an individual’s personal voting behavior changed in response to receiving debt relief (see e.g., Selvin 1958). To gain insights into the mechanisms underlying the aggregate reaction to the program, we make use of the Caucasus Research Resource Center’s (CRRC) Caucasus Barometer survey. The survey has rich data on political attitudes and its fall 2019 wave asks several questions specifically about the debt relief program. We link these individual survey data to our measures of district level debt relief intensity by utilizing CRRC’s confidential data on the place of residence of the surveyed individuals.

We adopt a difference-in-differences strategy with a continuous treatment intensity variable to estimate the effect of the debt relief program on the support for the two presidential candidates and on turnout in the two rounds of the 2018 election. The first round is the pre-treatment period, the second round is the post-treatment period, and the treatment is the announcement of the program in between the two rounds. All areas (either the 73 districts or the 3,600 precincts) were treated by the program. We, therefore, exploit variation in the program’s intensity following, e.g., Duflo (2001). We define, for each area, three continuous treatment variables: *log debt* is coded as the log monetary

¹¹ Figure D2 in the supplementary Appendix presents two heat maps showing Zurabishvili’s electoral performance and the intensity of the debt relief program across election districts.

¹² This is because the required information on their addresses was sometimes missing on the vali.ge website and, when recorded, the text data would not always match the public record, likely due to coding errors. Since we know the age and gender of the individuals in our sample, we can compare that to demographic data recorded at the precinct level. It is reassuring that we, on average, find no significant differences.

value of the total debt written off, *debt* as the monetary value of the total debt written off in billion GEL, and *log individuals* as the log number of individuals benefiting from the program.¹³ The headcount contrasts areas where many voters benefited with areas where few benefited. The monetary value of the debt write off takes into account how much debt was cancelled in each area and compares areas with high to areas with low values. The outcome variables are the percentage change in the votes for candidate i (i.e., either Zurabishvili or Vashadze) or the total votes cast for all candidates (i.e., turnout) between the two rounds in area r (denoted V_r^i and T_r , respectively). We formulate the difference-in-differences model as follows:

$$\Delta Votes_r^i = \alpha_1^i + \beta_1^i ProgramIntensity_r + C_r \gamma_1^i + \epsilon_{1r}, \quad (1)$$

where the variable *ProgramIntensity* is one of the three treatment intensity measures and the vector C_r contains pre-treatment area-specific observables. In all specifications, we include the log number of registered voters and turnout in the first round. The identifying assumption is that there are no omitted time-varying and area-specific effects correlated with the ‘treatment’ of each area. This is very likely to hold given the very short time between the first and second round of elections, but, in Section 4.2, we explicitly evaluate this common trend assumption using data from several past elections and show that it holds.

4 Results

4.1 Baseline results

Tables 1 and 2 present the baseline results at the level of election districts and precincts, respectively. Each table reports results for three outcome variables: the growth in votes

¹³ In the sample, there are many precincts without any affected individuals. These observations drop out in the logarithmic specifications, but are included in the specifications with *debt*.

Table 1: DISTRICT LEVEL RESULTS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Growth in votes between first and second round								
	Incumbent			Opposition			Turnout		
log debt	0.069** (0.029)			-0.067** (0.026)			0.004 (0.007)		
log individuals		0.038 (0.035)			-0.074** (0.031)			0.004 (0.008)	
debt (in billions)			9.807* (5.469)			-0.412 (5.026)			0.486 (1.301)
Observations	73	73	73	73	73	73	73	73	73
R-squared	0.557	0.528	0.541	0.122	0.111	0.036	0.581	0.580	0.579
F	28.86	25.73	27.16	3.201	2.858	0.860	31.86	31.76	31.69
Mean dep. var.		0.76			0.33			0.18	

Notes: The table presents estimates of equation (1) at the district level. The three dependent variables are the growth in votes between the first and second round of the election for the incumbent candidate, Zurabishvili (columns 1-3), the opposition candidate, Vashadze (columns 4-6), and for total turnout (columns 7-9). The treatment variable *log debt* is coded as the log monetary value of the total debt written off, *debt* as the monetary value of the total debt written off in billion GEL, and *log individuals* as the log number of individuals benefiting from the program. All regressions control for the log number of registered voters, log turnout in the first round of the election and include a constant term (not reported). Robust standard errors are reported in brackets. ***, ** and * indicate statistical significance at the 1%, 5% and 10% level, respectively.

between the first and second round of the election for the incumbent candidate, Zurabishvili (columns 1-3) and the opposition candidate, Vashadze (columns 4-6), and the growth rate of aggregate turnout between the two rounds (columns 7-9). The three treatment variables – *log debt*, *log individuals* and *debt* – measure the relative intensity of the debt relief program across districts or precincts.

The results reported in the two tables show that, both at the district and at the precinct level, the growth in the votes of the Georgian Dream supported candidate, Zurabishvili, between the two election rounds is positively related to the measures of debt relief intensity. The opposite is true for the opposition candidate, Vashadze. Overall, this demonstrates that the debt relief program helped the Georgian Dream get its candidate elected. Since the debt write off did not take place till after the election when the deal with the banks was finalized and implemented, the effect is driven by expectations of future benefits. Individuals could, given their knowledge of their own debt situation,

Table 2: PRECINCT LEVEL RESULTS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Growth in votes between the first and second round								
	Incumbent			Opposition			Turnout		
log debt	0.037*** (0.006)			-0.052*** (0.010)			0.003 (0.002)		
log individuals		0.029*** (0.010)			-0.051*** (0.016)			0.002 (0.003)	
debt (in billions)			151.223*** (55.984)			-55.303 (82.140)			3.289 (13.727)
Observations	2,209	2,209	3,685	2,209	2,209	3,683	2,209	2,209	3,685
R-squared	0.350	0.343	0.267	0.020	0.013	0.027	0.427	0.427	0.425
F	395.6	383.7	447.8	14.71	9.877	33.71	548.3	547.4	907.5

Notes: See the notes to Table 1. The only difference is that the units are precincts rather than districts and that, consequently, the standard errors are clustered at the district level.

estimate how likely they were to benefit from the program *if* it were to be implemented in full. On average, our estimates show that voters reacted to this by rewarding the Georgian Dream party prospectively. Moreover, the debt relief program did not increase aggregate turnout. This suggests that the swing induced by the program to Zurabishvili was driven by party switching rather than by voter mobilization. We cannot, however, rule out mobilization effects entirely because there could have been differential turnout effects that net out in the aggregate.

The program's treatment effect is substantial. The district level results reported in Table 1 show that a 10% increase in the amount of debt written off leads to a 7% increase in the votes won by the Georgian Dream candidate and a corresponding reduction in the votes for the opposition. A 10% increase in the size of the affected population leads to a 4% increase and to a 7% decrease in the votes of the incumbent and the opposition, respectively. The precinct level analysis reported in Table 2 yields more precise estimates that go in the same direction.

We discuss several robustness checks in supplementary Appendix C. These relate to whether our debt relief data are representative at the precinct level, to whether our results are sensitive to migration patterns, to controlling for all other candidates' votes obtained

in the first round, and to the possibility that the debt relief program affects the Georgian Dream party’s vote count indirectly by having effects on economic activity. The main results remain robust in all these tests.

4.2 Pre-trends and persistence

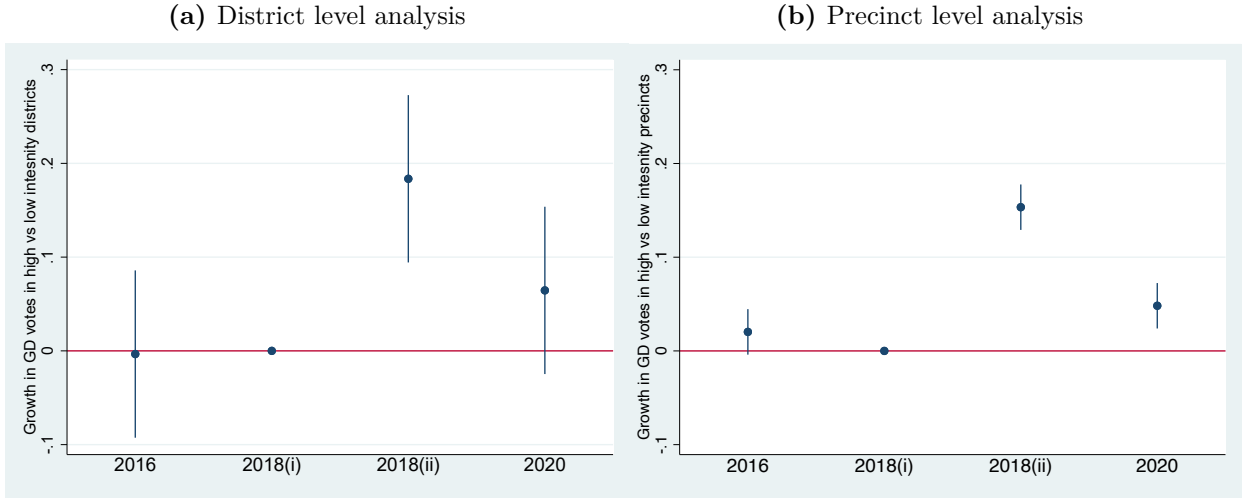
We need to address two important issues: pre-treatment trends in support for the Georgian Dream party and persistence of the effect of the debt relief program identified above. To do this, we collected data on the results of presidential and parliamentary elections in Georgia from 2013 to 2020. This corresponds to the period when the Georgian Dream was the incumbent party and the main opposition party was the United National Movement. To simplify the presentation of the results, we create a binary treatment variable. It is coded one if the total debt written off by the program in a district or a precinct is above the median, and zero otherwise.¹⁴ We interact the treatment dummy variable with indicators for the election years and estimate an event study model with these leads and lags of the treatment variable. As in the model in equation (1), the outcome variable for the event study specification is the growth of votes received by the Georgian Dream candidate. The specification is as follows:

$$\Delta Vote_{r,e} = \alpha_2 + \beta_{2e} \mathbb{1}ProgramIntensity_r \cdot \sum_{e=1}^5 Election + C_r \gamma_2^i + \epsilon_{2r}. \quad (2)$$

The difference-in-differences estimates in Tables 1 and 2 can be given a causal interpretation if the common trend assumption is valid. This assumption would be violated if the Cartu Foundation and the leaders of the Georgian Dream party, despite the unconditional nature of the program and the very short time they between the first and second rounds of the election, were able to target relief to areas with either increasing or decreasing support for the Georgian Dream party. To investigate this possibility, we consider the Georgian

¹⁴ We have experimented with other cut-offs and found qualitatively similar results. See, for example, Figure C2 in the Appendix for results comparing first and fifth quintiles. Results that replicate the baseline specifications of Tables 1 and 2 but by using a binary treatment variable as here are in Table C4.

Figure 2: THE TREATMENT EFFECT OVER TIME



Notes : The dependent variable is the growth of votes won by the Georgian Dream party in the election indicated on the x-axis compared to the previous election. We consider the presidential elections in 2013 and 2018 (in two rounds), and the parliamentary elections in 2016 and 2020. The treatment variable is a dummy coded one if the total debt written off by the program is above the median and zero otherwise interacted with an indicator for election years. The dots are the point estimates relative to the first round of the 2018 election and the bars are 95% confidence intervals. Otherwise, the specifications are similar to those of Table 1 and Table 2 for districts and precincts, respectively.

Dream party's performance in elections prior to 2018 and ask if there were systematic differences in the evolution of the votes it got in treated versus control areas. Figures 2(a)-(b) show the results across districts and precincts, respectively. They suggest that in 2016 the Georgian Dream party's election success was similar in treated and control districts and precincts. This is consistent with the common trend assumption. However, between the first and second round of the 2018 election, we observe a marked increase in the growth of votes for the Georgian Dream in treated relative to control areas. These results are very similar in the district and precinct specifications. The only difference is that, as expected, confidence intervals are larger in the district specification where we have fewer observations.

The baseline results in Tables 1 and 2 compare the first to the second round of the 2018 presidential election. Thus, they represent the short-run effect of the program and isolate the effect of prospective rewards. To investigate if the program's effect persisted into

the next election cycle, we estimate the differential effect on the growth in the Georgian Dream party’s votes in the treated districts and precincts but for the 2020 parliamentary election. Figures 2(a)-(b) show that two years after the debt relief program was enacted (and the debt had, in fact, been forgiven) support for the Georgian Dream had not returned to the pre-treatment trend. Instead, the treated areas continued to reward the Georgian Dream party. This result is consistent with retrospective voting and with existing clientelistic links between the Georgian Dream party and the recipients of the debt relief being perpetuated by the program.

4.3 Mechanisms

We consider the debt relief program to be a tool of distributive politics. Our interpretation is that the (prospective) beneficiaries of the program associated it with the Georgian Dream party and rewarded its candidate for it. The aggregate data do not allow us to infer this directly, so to substantiate this claim, we present additional evidence that speaks to the mechanism underlying the program’s average treatment effect. To do this, we leverage survey data and explore auxiliary predictions derived from the distributive politics literature.

We begin by showing that the treatment intensity variables used in the main analysis correlate with the individuals’ knowledge of the program and their attitude to and evaluation of it. To this end, we draw on individual level survey data from the 2019 wave of the Caucasus Barometer which asks questions about the debt relief program. With these data, we can regress the main district level measures of treatment intensity (*log debt* and *log individuals*) on binary indicators of individual attitudes towards the debt relief program as captured in the survey. Table 3 reports the results. The intercepts reported at the bottom of the table show that the vast majority of respondents are both aware of the program and view it in a positive light. More importantly, we see that individuals from districts that benefited more from the program are more likely to be aware of it (columns 1-2), and that they are more likely to report that they feel positively or very positively

Table 3: VALIDATION OF PROGRAM INTENSITY MEASURES WITH INDIVIDUAL SURVEY DATA

VARIABLES	(1) I have not heard about the debt relief program	(2) I feel (very) positively about the program	(3) I feel (very) positively about the program	(4) I feel (very) positively about the program	(5) Either I and/or my family benefited from the program	(6) Either I and/or my family benefited from the program
log debt	-0.0153*** (0.0031)		0.0163*** (0.0059)		0.0102** (0.0051)	
log individuals		-0.0216*** (0.0047)		0.0275*** (0.0085)		0.0082 (0.0076)
Constant	0.2352*** (0.0411)	0.1324*** (0.0300)	0.8199*** (0.0765)	0.9222*** (0.0544)	0.2704*** (0.0674)	0.3517*** (0.0491)
Observations	1,479	1,479	1,335	1,335	1,434	1,434
R-squared	0.0380	0.0363	0.0124	0.0143	0.0341	0.0322
F	11.65	11.11	3.332	3.869	10.08	9.493

Notes: The table presents linear probability regressions using individual level data from the fall 2019 wave of the Caucasus Barometer. The dependent variables ask various questions about the debt relief program and the answers are coded as dummy variables. The two treatment intensity variables are coded at the district level, and the surveyed individuals are matched to these districts using the geographical information in the survey. The regressions control for four respondent level characteristics: gender, age, years of education, and employment status. ***, ** and * indicate statistical significance at the 1%, 5% and 10% level, respectively.

about it (columns 3-4). Columns 5-6 relate the two treatment intensity measures to a self-reported measure of whether a respondent or their family benefited from the debt relief program. We see that the correlation is positive, but statistically significant only for *log debt*. The self-reported share of families benefiting from the program is, on average, several times smaller than the share that benefited according to our objective data. It is, therefore, likely that the survey respondents were reluctant to report that they had benefited from the program and that this creates measurement error that inflates the standard errors. Overall, the results reported in Table 3 validates that the treatment intensity variables capture the impact of the program.

Next, we use the individual level survey data to investigate if the self-reported awareness of and attitude to the program are related to self-reported party identification, and to trust in the president and in fair elections. Table 4 reports the results. The results in columns 1-4 show that respondents who are more aware of the program and who view it

Table 4: POLITICAL ATTITUDES AND THE DEBT RELIEF PROGRAM

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	The party closest to me is		UNM		I rather or fully		Elections were	
	GD				trust the president		conducted fairly	
I have not heard about the debt relief program	-0.1771*** (0.0488)		0.0814** (0.0336)		-0.1239*** (0.0413)		-0.1163** (0.0504)	
I feel positively or very positively about the program		0.1757*** (0.0283)		-0.0851*** (0.0189)		0.1031*** (0.0243)		0.1726*** (0.0264)
Constant	0.1761*** (0.0517)	0.0016 (0.0614)	0.2186*** (0.0356)	0.2956*** (0.0409)	0.1779*** (0.0437)	0.0623 (0.0526)	0.1030** (0.0498)	-0.0788 (0.0581)
Observations	1,950	1,781	1,950	1,781	2,164	1,955	1,873	1,724
R-squared	0.0140	0.0279	0.0124	0.0182	0.0093	0.0155	0.0135	0.0367
F	5.508	10.19	4.895	6.577	4.071	6.152	5.092	13.11

Notes: The table presents linear probability regressions using individual level data from the fall 2019 wave of the Caucasus Barometer. The dependent variables ask various questions related to the political attitudes of the respondents and are coded as dummy variables. The main independent variables code questions about the debt relief program as dummy variables. The regressions control for the same respondent level characteristics as in Table 3. GD = Georgian Dream; UNM = United National Movement. ***, ** and * indicate statistical significance at the 1%, 5% and 10% level, respectively.

in a more positive light are more likely to self-identify with the Georgian Dream party and less likely to self-identity with the United National Movement. Moreover, knowledge of the program and a positive attitude towards it are positively correlated with trust in the elected president (columns 5-6), and, somewhat ironically, with the view that the 2018 election was conducted fairly (columns 7-8). These correlations are consistent with our interpretation of the difference-in-differences estimates as evidence that voters rewarded prospectively the Georgian Dream party for the debt relief program.

Finally, the literature on distributive politics and vote buying show that it is easier to “buy” votes from poor than from rich voters (Stokes et al. 2013; Aidt and Jensen 2017). If the debt relief program caused the observed swing towards the Georgian Dream, then we would expect to observe a stronger effect in poor districts than in rich ones. To test this, we rely on night light emission data aggregated to the district levels as a proxy for average income of a district. We augment the baseline difference-in-differences model

in equation (1) with an interaction between this proxy for income and the treatment intensity variable *log debt*. Figure C3 in the supplementary Appendix shows that the debt relief program was successful in swinging votes towards Zurabishvili and away from Vashadze in poorer districts, while it was ineffective in richer districts. This is consistent with the distributive politics interpretation of the program.

4.4 What does it cost to win?

As a counter-factual exercise, we use our estimate of the treatment effect to calculate how much it would have cost the Cartu Foundation to swing the presidential election in favor of the Georgian Dream party and to calculate the cost of a single vote (see Condra et al. (2018) for a similar approach). Supplementary Appendix E presents the details behind these calculations.

We report in Table 1, column 3 the program’s treatment effect on the growth rate of votes for Zurabishvili and the Georgian Dream per billion GEL. We use this point estimate to calculate the sum that would be required for Zurabishvili to reach the votes that the opposition candidate actually received in the second round of the election. This calculation suggests that the foundation needed to buy up debt with a book value of \$67.8 million to swing the election by one vote in favor of Zurabishvili. The actual cost of buying the loans was almost surely less than the book value. If we assume that the foundation paid about 20% of the book value (as discussed in footnote 8), we get that the cost of buying victory with a minimal margin was \$13.6 million. Zurabishvili in fact won with a clear margin, and our calculation suggests that buying debt with a book value of \$217.1 million or a likely market value of \$43.4 million was sufficient to guarantee this victory. The actual size of the program was much larger than what, as we estimate, was needed to reach these victories, and stood at \$578 million in book value, or \$144.4 in market value. Thus, we conclude that the election results could have easily been the reverse absent the debt relief program.

The Georgian debt relief program was large-scale and benefited one-sixth of the electorate. In contrast to small, targeted programs, these features make it plausible to extrapolate the cost of one vote from the cost of winning the overall election. Our calculation suggests that the cost of a single vote was \$81.6 of debt relief (GDP per capita is about \$4200). This is ‘cheap’ relative to other estimates of the ‘price of a vote’ in the literature. Bechtel and Hainmueller (2011), for example, report in the context of disaster relief in Germany in 2002 that the price per vote was about EURO 63,000 (GDP per capita around EURO 25,000), while Healy and Malhorta (2009) estimate one additional vote cost \$27,000 (GDP per capita around \$45,000) in relief spending in the USA.

5 Conclusion

We argue that debt relief programs can serve as a tool of distributive politics and help politicians win elections. This is a new perspective on the role of such programs which are widely used both in established and in new democracies. Our study of the 2018 presidential election in Georgia shows that a large-scale privately funded debt relief program announced between the two rounds of the election had a substantive impact on the aggregate votes cast in favor of the main two candidates. The effect persisted into the next election, hinting at perpetuated clientelistic links between the winner of the election and those who benefited from the program and in that way undermining electoral accountability (see Leight et al. 2020).

References

- Agarwal, S., G. Amromin, I. Ben-David, S. Chomsisengphet, T. Piskorski, and A. Seru (2017). Policy intervention in debt renegotiation: Evidence from the home affordable modification program. *Journal of Political Economy* 125(3), 654–712.
- Aidt, T., Z. Asatryan, L. Badalyan, and F. Heinemann (2020). Vote buying or (political) business (cycles) as usual? *The Review of Economics and Statistics* 102(3), 409–425.
- Aidt, T. S. and P. S. Jensen (2017). From open to secret ballot: Vote buying and modernization. *Comparative Political Studies* 50, 555–93.
- Bagues, M. and B. Esteve-Volart (2016). Politicians’ luck of the draw: Evidence from the Spanish christmas lottery. *Journal of Political Economy* 124(5), 1269–1294.
- Bechtel, M. M. and J. Hainmueller (2011). How lasting is voter gratitude? An analysis of the short- and long-term electoral returns to beneficial policy. *American Journal of Political Science* 55(4), 852–868.
- Carvalho, D. (2014). The real effects of government-owned banks: Evidence from an emerging market. *The Journal of Finance* 69(2), 577–609.
- Cole, S. (2009). Fixing market failures or fixing elections? Agricultural credit in India. *American Economic Journal: Applied Economics* 1(1), 219–50.
- Condra, L. N., J. D. Long, A. C. Shaver, and A. L. Wright (2018). The logic of insurgent electoral violence. *American Economic Review* 108(11), 3199–3231.
- Cox, G. W. and M. D. McCubbins (1986). Electoral politics as a redistributive game. *Journal of Politics* 48, 370–389.
- Daniels Jr, G. E. and A. Smythe (2019). Student debt and labor market outcomes. *AEA Papers and Proceedings* 109, 171–175.

- De La O, A. L. (2013). Do conditional cash transfers affect electoral behavior? Evidence from a randomized experiment in Mexico. *American Journal of Political Science* 57(1), 1–14.
- Delatte, A.-L., A. Matray, and N. Pinardon-Touati (2023). Political quid pro quo in financial markets. *Working Paper, HAL-Id 04336193*.
- Dixit, A. and J. Londregan (1996). The determinants of success of special interests in redistributive politics. *Journal of Politics* 58(4), 1132–1155.
- Dobbie, W., P. Goldsmith-Pinkham, and C. S. Yang (2017). Consumer bankruptcy and financial health. *Review of Economics and Statistics* 99(5), 853–869.
- Dobbie, W. and J. Song (2015). Debt relief and debtor outcomes: Measuring the effects of consumer bankruptcy protection. *American economic review* 105(3), 1272–1311.
- Doerr, S., S. Gissler, J.-L. Peydró, and H.-J. Voth (2022). Financial crises and political radicalization: How failing banks paved hitler’s path to power. *The Journal of Finance* 77(6), 3339–3372.
- Duflo, E. (2001). Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment. *The American Economic Review* 91(4), 795–813.
- Englmaier, F. and T. Stowasser (2017, 04). Electoral cycles in savings bank lending. *Journal of the European Economic Association* 15(2), 296–354.
- Exler, F. and M. Tertilt (2020). Consumer debt and default: A macro perspective. *Oxford Research Encyclopedia of Economics and Finance*.
- Faraz, N. and M. Rockmore (2020). Election cycles in public credit: Credit provision and default rates in Pakistan. *Journal of Development Economics* 147, 102528.

- Field, E. (2009). Educational debt burden and career choice: Evidence from a financial aid experiment at NYU Law School. *American Economic Journal: Applied Economics* 1(1), 1–21.
- Folch, M. and L. Mazzone (2022). Go big or buy a home: The impact of student debt on career and housing choices. *Available at SSRN 3805220*.
- Funke, M., M. Schularick, and C. Trebesch (2016). Going to extremes: Politics after financial crises, 1870–2014. *European Economic Review* 88, 227–260.
- Giné, X. and M. Kanz (2018). The economic effects of a borrower bailout: Evidence from an emerging market. *The Review of Financial Studies* 31(5), 1752–1783.
- Golden, M. and B. Min (2013). Distributive politics around the world. *Annual Review of Political Science* 16(1), 73–99.
- Gyöngyösi, G. and E. Verner (2022). Financial crisis, creditor-debtor conflict, and populism. *The Journal of Finance* 77(4), 2471–2523.
- Healy, A. and N. Malhorta (2009). Myopic voters and natural disaster policy. *American Political Science* 103, 387–406.
- Kanz, M. (2016). What does debt relief do for development? Evidence from India’s bailout for rural households. *American Economic Journal: Applied Economics* 8(4), 66–99.
- Karlan, D., S. Mullainathan, and B. N. Roth (2019). Debt traps? Market vendors and moneylender debt in India and the Philippines. *American Economic Review: Insights* 1(1), 27–42.
- Kern, A. and P. Amri (2021). Political credit cycles. *Economics & Politics* 33(1), 76–108.
- Leight, J., D. Foarta, R. Pande, and L. Ralston (2020). Value for money? Vote-buying and politician accountability. *Journal of Public Economics* 190, 104227.

- Lindbeck, A. and J. W. Weibull (1987). Balanced-budget redistribution as the outcome of political competition. *Public Choice* 52, 273–297.
- Luo, M. and S. Mongey (2019). Assets and job choice: Student debt, wages and amenities. *National Bureau of Economic Research Working Paper* 25801.
- Manacorda, M., E. Miguel, and A. Vigorito (2011, July). Government transfers and political support. *American Economic Journal: Applied Economics* 3(3), 1–28.
- Mitra, A., S. Mitra, and A. Mukherji (2021). Cash for votes: Evidence from India. *University of Kent, Working Paper*.
- Mukherjee, S., K. Subramanian, and P. Tantri (2018). Borrowers’ distress and debt relief: Evidence from a natural experiment. *The Journal of Law and Economics* 61(4), 607–635.
- Müller, K. (2023). Electoral cycles in macroprudential regulation. *American Economic Journal: Economic Policy* 15(4), 295–322.
- Pop-Eleches, C. and G. Pop-Eleches (2012). Targeted government spending and political preferences. *Quarterly Journal of Political Science* 7(3), 285–320.
- Rothstein, J. and C. E. Rouse (2011). Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics* 95(1-2), 149–163.
- Selvin, H. C. (1958). Durkheim’s suicide and the problems of empirical research. *The American Journal of Sociology* 63(6), 608–619.
- Stokes, C. S. (2005). Perverse accountability: A formal model of machine politics with evidence from Argentina. *American Political Science Review* 99(3), 325–325.
- Stokes, C. S., T. Dunning, M. Nazareno, and V. Brusco (2013). *Brokers, Voters, and Clientelism: The puzzle of distributive politics*. Cambridge University Press.
- Vannutelli, S. (2023). The political economy of stimulus transfers. *Working paper*.

World Bank (2022). *World Development Indicators* (2022 ed.). Washington DC: World Bank. <http://datatopics.worldbank.org/world-development-indicators> (accessed August 2022).

Zitzewitz, E. (2012). Forensic economics. *Journal of Economic Literature* 50(3), 731–769.

Zucco Jr, C. (2013). When payouts pay off: Conditional cash transfers and voting behavior in Brazil 2002–10. *American Journal of Political Science* 57(4), 810–822.

Supplementary (Online) Appendix

Political Consequences of (Consumer) Debt Relief

by Toke Aidt, Zareh Asatryan and Lusine Badalyan

Appendix A: US student debt relief

The relief plan: On 24th of August 2022, US president Joe Biden announced a student loan debt relief plan that would write off \$10,000 of student debt for American citizens with income lower than \$125,000 (\$250,000 for married couples).¹⁵ The plan builds on promise that Biden made during the 2018 presidential election campaign, but was announced formally in August 2022, two months prior to the mid-term elections that year. The plan was supposed to provide assistance to up to 43 million people, and, according to the Congressional Budget Office¹⁶, it would cost about \$400 billion over a period of thirty years. According to the US Department of Education, about 65% of the relief would benefit individuals below the age of 40, and around 90% of the funds would go to individuals with annual incomes under \$75,000.

The formal procedure for claiming relief was released by the Biden administration on 17th of October 2022, only three weeks before the mid-term elections. However, six Republican states¹⁷ immediately contested the president's authority to cancel student debt and initiated a legal challenge. Eventually, in June 2023, the plan was ruled out as


¹⁵ With respect to debts held by the Department of Education, the American Department of Education would write off up to \$20,000 in debt for Pell Grant recipients.

¹⁶ Source: <https://www.cbo.gov/system/files/2022-09/58494-Student-Loans.pdf>.

¹⁷ Arkansas, Iowa, Kansas, Missouri, Nebraska and South Carolina.

being unconstitutional by the US Supreme Court.¹⁸ Figure A1 shows google trends data for searchers for the term “student debt relief” in the US. As expected, search results take off around the week of the announcement in August 2022 and reach their peak around the week when the formal application process opens. Searches die out after the US Supreme Court decision in June 2023.

Figure A1: Google Trends search results for “student debt relief” in USA



Notes: The figure shows the popularity index by week in 2022 and 2023 related to the search term “student debt relief” in the US based on Google Trends. The popularity index ranges from 0 (lowest) to 100 (highest).

Data: We collect data on the number of student debt holders in each US state.¹⁹ From this we construct our treatment variable which is a dummy variable equal to one if the state has more student debt holders than the median state and zero otherwise. The outcome variable is the monthly approval rate for President Biden from April 2022 to January 2023 as captured by online opinion polls.²⁰ We complement these variables with a number of economic and demographic control variables measured at the level of US states. Table A1 reports summary statistics and lists the sources of the variables.

Empirical strategy: We want to know if voters in states with above medium student debt holders (i.e., many potential relief beneficiaries) are more likely to support the Biden administration after the announcement of the relief policy in August 2022 compared to the periods before the announcement using states with below medium student debt holders

¹⁸ See US Supreme Court ruling (accessed February 2024).

¹⁹ The data are collected by the US Census Bureau statistics and is available at <https://www.studentaid.gov/> (accessed March 2024).

²⁰ Source: <https://civiqs.com/> (accessed March 2024).

Table A1: Summary statistics for variables used in the US student debt relief exercise

	Mean	SD	Min	Max	Source
Log number of student debt holders	13.11	1.06	10.92	15.20	StudentAid.gov
Approval rates, difference November to July 2022	0.08	0.02	0.03	0.12	Civiqs
Log GDP per capita Q3 2022	11.20	0.25	10.77	12.40	U.S. Bureau of Economic Analysis
Unemployment rate, %	3.52	0.72	2.20	5.20	U.S. Bureau of Labor Statistics
% population with at least high school diploma	0.91	0.03	0.80	0.95	U.S. Census Bureau
Log households income	11.07	0.17	10.75	11.42	World Population Review
Log population	15.23	1.02	13.27	17.50	U.S. Census Bureau

(few potential relief beneficiaries) as the control group. To this end, we estimate the following equation:

$$\Delta Polls_{s,m} = \alpha + \beta_m \mathbb{1}Borrowers_s \cdot \sum_{m=1}^8 Month + \gamma^i \mathbb{X}_s + \epsilon_{s,m} \quad (3)$$

where $\Delta Polls_{s,m}$ is the monthly change in Biden’s approval rating in each state, $\mathbb{1}Borrowers_s$ is a dummy variable that equals one for states that have more debt holders than the median state and zero otherwise, $\sum_{m=1}^8 Month$ are a set of dummies for each month taking the pre-reform month of July as the base month, \mathbb{X}_s are the set of control variables as defined above, and $\epsilon_{s,m}$ is the error term. Standard errors are clustered by state.

Results: Table A2 shows the estimates of equation (3), and the estimates of β_m are plotted graphically in Figure 1 in the main text. The treatment variables, indicating differences in Biden approval ratings across states that have many debt holders as opposed to states that have fewer debt holders, become positive and statistically different from zero from September on wards, that is, right after the announcement of the student debt relief policy in August, 2022. The treatment effect peaks in November, during the month of mid-term elections, and gradually declines but still indicating a significant difference in January 2023, the last month of our sample. Approval ratings for the treatment and control states are not statistically different in the months before the announcement of the debt relief program. Column 1 to 3 of Table A2 include different sets of state-level control variables. We see that the results are always similar.

Table A2: US student debt relief and opinion polls

VARIABLES	Monthly change in Biden's approval ratings		
High borrower states	0.030 (0.021)	0.018 (0.014)	0.025 (0.028)
High borrower states \times April 2022	0.008 (0.005)	0.008 (0.005)	0.008 (0.005)
High borrower states \times May 2022	0.005 (0.004)	0.005 (0.004)	0.005 (0.004)
High borrower states \times June 2022	0.001 (0.002)	0.001 (0.002)	0.001 (0.002)
High borrower states \times July 2022	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
High borrower states \times August 2022	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)
High borrower states \times September 2022	0.009* (0.005)	0.009* (0.005)	0.009* (0.005)
High borrower states \times October 2022	0.011** (0.005)	0.011** (0.005)	0.011** (0.005)
High borrower states \times November 2022	0.015*** (0.005)	0.015*** (0.005)	0.015*** (0.005)
High borrower states \times December 2022	0.015** (0.006)	0.015** (0.006)	0.015** (0.006)
High borrower states \times January 2023	0.012** (0.005)	0.012** (0.005)	0.012** (0.005)
April 2022	0.054*** (0.004)	0.054*** (0.004)	0.054*** (0.004)
May 2022	0.039*** (0.003)	0.039*** (0.003)	0.039*** (0.003)
June 2022	0.028*** (0.001)	0.028*** (0.001)	0.028*** (0.001)
August 2022	0.022*** (0.002)	0.022*** (0.002)	0.022*** (0.002)
September 2022	0.076*** (0.004)	0.076*** (0.004)	0.076*** (0.004)
October 2022	0.077*** (0.004)	0.077*** (0.004)	0.077*** (0.004)
November 2022	0.071*** (0.004)	0.071*** (0.004)	0.071*** (0.004)
December 2022	0.084*** (0.005)	0.084*** (0.005)	0.084*** (0.005)
January 2023	0.085*** (0.004)	0.085*** (0.004)	0.085*** (0.004)
Log population			-0.004 (0.020)
Log GDP Q3 2022 pro capita		-0.151* (0.078)	-0.150* (0.080)
Unemployment rate		0.032** (0.012)	0.032** (0.013)
% of population with at least high school diploma		0.189 (0.383)	0.126 (0.328)
Log households income		0.455*** (0.080)	0.457*** (0.079)
Constant	0.276*** (0.017)	-3.344*** (0.530)	-3.255*** (0.606)
Observations	500	500	500
R-squared	0.165	0.599	0.600
F	126.3	111.3	108.2

Notes: The table shows estimates of equation (3). The coefficients on the treatment dummies (β_m) are also plotted in Figure 1 in the main text. The outcomes variable measures the change in monthly approval ratings for President Biden from April 2022 to January 2023. The treatment variables are interaction terms between the dummy taking the value of one if the log number of student debt holders is above that in the median state and zero otherwise, and an indicator variable for the month of the year. The three columns of the table use different sets of control variables. Statistical significance denoted as: *** p<0.01, ** p<0.05, * p<0.1. Errors are clustered by state.

Appendix B: Sampling strategy

The website vali.ge, operational between early December 2018 and end of December 2019, allowed Georgian voters to look up if they had received debt relief using their eleven-digit social security ID. This ID system was introduced in the late 2000s. We collected the data on who benefited from the debt relief program by randomly sampling these IDs. The structure of the ID numbers are

$$YYZZZXXXXXX$$

where the two first digits (YY) represent the 73 districts used by the Georgian Election Commission to record election results,²¹ the next three digits (ZZZ) represent sub-districts within some of the larger districts, including the capital Tbilisi, and the last six digits, ($XXXXXX$), is a unique number assigned to an individual residing inside the geography defined by the first five digits. In practice, almost all IDs have the six last digits between 000000 and 170000, i.e., IDs outside this range are not being used.

Our sampling strategy was as follows:

1. We stratified on districts and sampled S IDs within each of the 73 districts based on a randomization of the six last digits (within the range from 000000 to 170000).
2. Within the (larger) districts which have sub-districts, we sampled $\frac{S}{D_r}$ ID numbers from each of the D_r sub-districts within district r . We then apply weights proportional to the number of sub-districts when aggregating the data to the level of districts.

This sampling strategy delivers an estimate of the number of beneficiaries within a district which is proportional to the true number of beneficiaries as well as an estimate of the total debt that was relieved in the district.

²¹ We exclude voters residing outside Georgia and four tiny districts with a total of 138 residents.

To see this, let C denote the 170,000 potential IDs that could be in use within each district $r \in \{1, \dots, 73\}$, let n_r be the population in district r , and let b_r be the number of beneficiaries in the district.

We sample ID numbers, not individuals. If all ID numbers had been in use in each district ($C = n_r$ for $\forall r$), then it is clear that a random sample of IDs (or individuals) would, as the sample size converges to the population, give us an unbiased estimate of the proportion of the population who benefited from the program:

$$\frac{\hat{b}_r(S)}{S} \rightarrow \frac{b_r}{n_r}$$

where $\hat{b}_r(S)$ is the number of IDs in a sample of size S that were associated with an individual who benefited.

However, not all ID numbers were in use in each district, i.e., $C > n_r$ for $\forall r$. Accordingly, when we draw a number which does not match a beneficiary, we do not know if this is because no person with that number exists or if the person with the number did not benefit from the program. In this case, $\frac{\hat{b}_r(S)}{S}$ converges to $\frac{b_r}{C}$ as the sample size converges to the total number of possible IDs, C . From this, we get

$$b_r \approx \frac{\hat{b}_r(C)}{S} C$$

Our main regression uses the cross sectional variation in the expected number of beneficiaries in each district. Using the notation from above, we can write this as

$$\Delta V_r = \alpha' + \dots = \beta_1 \ln(b_r) + \alpha' + \beta_1 \ln\left(\frac{\hat{b}_r}{C} S\right) + \dots = \alpha + \beta_1 \ln(\hat{b}_r) + \dots$$

where $\alpha = \alpha' + \beta_1(\ln(S) - \ln(C))$. From this, we observe that the regression exploits that our sampling strategy provides appropriate cross sectional variation in how many individuals benefited, i.e., of the program intensity.

A similar argument can be made with respect to the total debt relieved in a given district. If we let B_r be the average debt relieved in district r , then we can write the total debt relieved as $b_r B_r$ and for a large sample S , we got $b_r B_r \approx \frac{\hat{b}_r \hat{B}_r}{C} S$.

Appendix C: Additional tests

We implement several robustness tests related to the sampling strategy and to other potential concerns.

Sample strategy

Since we stratified the sampling at the district level, the sample may not be representative at the precinct level. Precincts in Georgia are constructed such that they include about a 1,000 voters. However, the allocation of voters within districts to precincts is not random. Our precinct level results could be biased if the allocation of individuals to precincts is correlated with the measures of program intensity. To evaluate this, we report in Table C1 results from a specification that includes district fixed effects in the precinct level model. This controls for the potential correlation between the sample measures of program intensity and the rules of allocating voters into precincts that are specific to districts. It is reassuring to observe that these estimates are similar to the baseline results.

Another potential issue with our sampling strategy relates to migration. The social security ID encodes an individual's place of residence when the ID is issued (in the late 2000s), not the residence at the time of voting (in 2018). Thus, if (many) people migrated between the time when the IDs were issued and the 2018 election, we will incorrectly inflate the program intensity measures in the out-migration areas where the IDs were issued. Unfortunately, we do not observe migration patterns. However, several factors mitigate this problem. First, the social security system was established in the late 2000s. Thus, the IDs represent the place of residence about a decade before the election in 2018 and not at a place much earlier in a voter's life which would be the case if, for example, IDs were given at birth. Second, most migration is rural to urban, and in Georgia in particular from rural areas to the capital Tbilisi where almost a third of Georgia's population lives. We can, therefore, evaluate the sensitivity of our results to migration by omitting all the

precincts of Tbilisi. Table C2 shows that the baseline estimates remain robust to this. Third, we can use the CRRC survey data to estimate the likelihood that an individual migrates as a function of individual observable characteristics. To do this, we use a question that asks if respondents are registered (got their ID) in the place where they live at the moment. Migration decisions are strongly negatively correlated with age. We can use these estimated migration probabilities by age as inverse sample weights (i.e., giving less weight to individuals with an age associated with a high migration probability) when aggregating the individual level data to the measures of district level debt relief intensity. Table C3 shows the results remain similar to our baseline estimates.

Indirect channels

The short time window between the first and the second round of the presidential election mitigates concerns about other policy changes potentially driving the results. Nevertheless, we test here two possible alternative explanations.

First, as long as the debt relief program could have led to effects on economic activity, these effects could have affected the Georgian Dream party’s vote indirectly (Bagues and Esteve-Volart 2016). In the absence of geographically disaggregated economic data in Georgia, we measure economic activity with night light emission.²² Using mean monthly growth in night light as the outcome variable, we estimate the differential effect of the program on economic activity using a specification similar to the one used to construct Figure 2. That is, we compare economic activity in districts with high versus low program intensity. The results are summarized in Figure C1. We do not find significant differences between treated and control areas, neither before and nor after the treatment. We conclude that it is unlikely that the debt relief program’s treatment effect was caused by differential economic trends.

²² The source of the night light data is the Visible and Infrared Imaging Suite (VIIRS) of the Earth Observation Group. The data is available here (accessed April 2021).

Second, the pool of candidates changes between the two elections as only the runner-up candidates make it to the second round. This can potentially pose a challenge if the spatial distribution of votes that went to the losing candidates in the first round is correlated with the program intensity. To rule out this possibility, in Table C5 we control for the share of votes obtained in the first round by all candidates who did not make it to the second round. The results remain similar to the baseline results.

Table C1: RESULTS WITH DISTRICT FIXED EFFECTS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Growth in votes compared to first round								
	Incumbent		Opposition				Turnout		
log debt	0.024*** (0.006)			-0.025** (0.010)			0.004** (0.002)		
log individuals		0.025*** (0.009)			-0.015 (0.016)			0.004* (0.002)	
debt (in billions)			33.959 (49.505)			-7.165 (76.640)			3.315 (12.488)
Observations	2,209	2,209	3,685	2,209	2,209	3,683	2,209	2,209	3,685
R-squared	0.532	0.531	0.455	0.182	0.180	0.194	0.573	0.572	0.547
F	31.93	31.70	39.60	6.245	6.168	11.39	37.62	37.53	57.38

Notes: The regressions are similar to those in Table 2 except that they include district fixed effects. See the notes to Table 2 for further details.

Table C2: RESULTS WITHOUT TBILISI

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Growth in votes compared to first round								
	Incumbent			Opposition			Turnout		
log debt	0.0352*** (0.0070)			-0.0596*** (0.0115)			0.0028 (0.0018)		
log individuals		0.0166 (0.0111)			-0.0653*** (0.0181)			0.0027 (0.0029)	
debt (in billions)			17.6832*** (5.0113)			-15.2834** (7.3969)			1.0366 (1.2340)
Observations	2,073	2,073	3,549	2,073	2,073	3,547	2,073	2,073	3,549
R-squared	0.3397	0.3324	0.2632	0.0216	0.0150	0.0278	0.4247	0.4243	0.4239
F	354.9	343.4	422.0	15.20	10.49	33.83	509.2	508.3	869.5


Notes: The regressions are similar to those in Table 2 except that they exclude precincts located in the largest city and capital of Georgia, Tbilisi. See the notes to Table 2 for further details.

Table C3: RESULTS FOR DATA AGGREGATED WITH SAMPLE WEIGHTS BASED ON MIGRATION PROBABILITIES

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Growth in votes compared to first round								
	Incumbent			Opposition			Turnout		
log debt	0.070** (0.030)			-0.081*** (0.027)			0.007 (0.007)		
log individuals		0.033 (0.036)			-0.086*** (0.031)			0.007 (0.008)	
debt (in billions)			20.332* (11.751)			-5.705 (10.761)			1.514 (2.788)
Observations	73	73	73	73	73	73	73	73	73
R-squared	0.555	0.526	0.540	0.149	0.133	0.040	0.584	0.583	0.580
F	28.65	25.51	27.01	4.040	3.514	0.955	32.28	32.16	31.81


Notes: The regressions are similar to those in Table 1 except that they use migration probabilities by age as inverse sample weights when aggregating the individual level data to the measures of district level debt relief intensity. See the notes to Table 1 for further details.

Figure C1: THE TREATMENT EFFECTS ON AVERAGE GROWTH IN MONTHLY NIGHT-LIGHT LUMINOSITY OVER TIME



Notes: The dependent variable is the growth of average district level monthly night light luminosity emission compared to the same month of the previous year. The treatment variable is defined as a dummy coded one if the total debt written off by the program is above the median and zero otherwise interacted with an indicator for time. Estimates are relative to September 2018.

Figure C2: THE TREATMENT EFFECT OVER TIME



Notes : The regressions are similar to those in Figure 2 except that the treatment variable is a dummy coded one if the total debt written off by the program is at the top quintile and zero in the bottom quintile of the distribution of program intensity.

Table C4: TREATMENT EFFECTS BY LOW-VS-HIGH DEBT INTENSITY

	(1)	(2)	(3)	(4)	(5)	(6)
		Districts			Precincts	
		Growth in votes compared to first round				
	Incumbent	Opposition	Turnout	Incumbent	Opposition	Turnout
Median of debt (in billions)	0.146** (0.066)	-0.103** (0.049)	0.030** (0.014)	0.120*** (0.021)	-0.168*** (0.026)	0.013** (0.006)
Observations	73	73	73	3,685	3,683	3,685
R-squared	0.339	0.078	0.518	0.009	0.011	0.001
F	17.92	2.942	37.67	33.83	40.50	5.318

Notes: The regressions in columns (1-3) and (4-6) are similar to those in Tables 1 and 2, respectively, except that the treatment variable follows the definition of Figure 2 and is defined to be a binary variable coded one if the total debt written off by the program is above the median and zero otherwise.

Table C5: TREATMENT EFFECTS CONTROLLING FOR OTHER CANDIDATES VOTES


	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Growth in votes compared to first round								
	Incumbent			Opposition			Turnout		
Panel A: Districts									
log debt	0.057** (0.024)			-0.075*** (0.023)			0.004 (0.007)		
log individuals		0.052* (0.028)			-0.065** (0.028)			0.005 (0.008)	
debt (in billions)			4.051 (4.659)			-4.317 (4.716)			0.222 (1.335)
Observations	73	73	73	73	73	73	73	73	73
R-squared	0.709	0.699	0.688	0.302	0.253	0.204	0.586	0.586	0.584
F	41.44	39.49	37.45	7.339	5.748	4.349	24.05	24.08	23.90
Panel B: Precincts									
log debt	0.027*** (0.006)			-0.061*** (0.010)			0.002 (0.002)		
log individuals		0.023** (0.009)			-0.057*** (0.016)			0.002 (0.003)	
debt (in billions)			57.636 (51.107)			-120.343 (80.710)			-3.388 (13.664)
Observations	2,209	2,209	3,683	2,209	2,209	3,681	2,209	2,209	3,683
R-squared	0.427	0.424	0.392	0.058	0.048	0.065	0.428	0.428	0.433
F	411.0	405.2	593.2	33.78	28.01	63.97	412.5	411.9	702.8

Notes: The regressions of Panels A and B are similar to those in Tables 1 and 2, respectively, except that they additionally control for the votes of all other candidates in the first round of the election (see vote shares in Table D1).

Table C6: THE TREATMENT EFFECT IN POOR AND RICH DISTRICTS

	(1)	(2)	(3)
	Growth in votes	Growth in votes	Growth in votes
	Incumbent	Opposition	Turnout
log debt	0.1569*** (0.0322)	-0.0883*** (0.0253)	0.0270*** (0.0098)
mean night light	0.0885** (0.0429)	-0.0574* (0.0337)	0.0158 (0.0130)
log debt x mean night light	-0.0049 (0.0030)	0.0042* (0.0024)	-0.0009 (0.0009)
Constant	-1.5100*** (0.4513)	1.5676*** (0.3545)	-0.2128 (0.1369)
Observations	73	73	73
R-squared	0.4477	0.1518	0.1771
F	18.64	4.118	4.949

Notes: This table adds an interaction term between one of the measures of treatment intensity (*log debt*) and district level average night light emission to the baseline specification of equation (1). The average marginal effects are plotted in Figure C3. See the notes to Table 1 for further details.

Figure C3: TREATMENT EFFECTS IN POOR AND RICH DISTRICTS

Notes: This figure plots the average marginal effects of the regressions shown in Table C6. The y-axis on the left shows the average marginal effects on the growth of votes for the respective candidate. The y-axis on the right corresponds to the histogram of the night light data. District level average night light emission are plotted on the x-axis.

Appendix D: Descriptive statistics

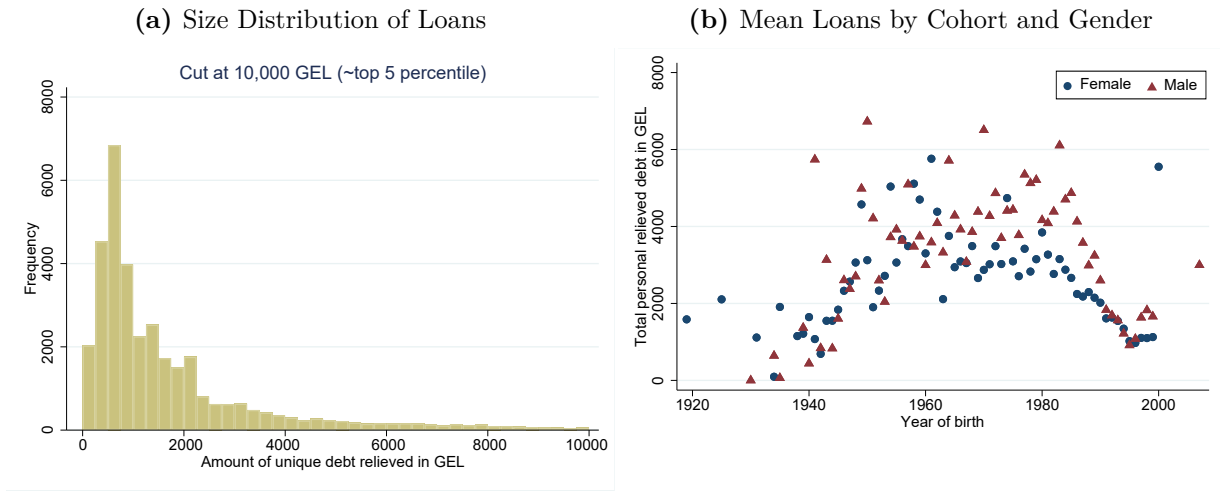
Table D1: RESULTS OF 2018 PRESIDENTIAL ELECTION

	First round: October 28th	
	Votes	%
Salome Zurabishvili	615,572	38.64
Grigol Vashadze	601,224	37.74
Davit Bakradze	174,849	10.97
Shalva Natelashvili	59,651	3.74
David Usupashvili	36,037	2.26
Zurab Japaridze	36,034	2.26
Kakha Kukava	21,186	1.33
...		
Turnout	1,647,878	46.83

	Second round: November 28th	
	Votes	%
Salome Zurabishvili	1,147,701	59.52
Grigol Vashadze	780,680	40.48
Turnout	1,988,787	56.36

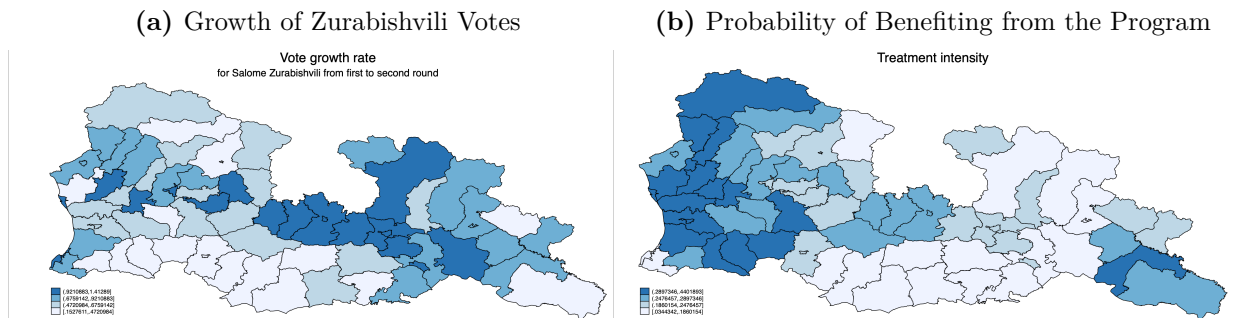
Notes: The data come from the Election Administration of Georgia. The top panel reports the first round results for candidates who received more than 1% of the votes. The bottom panel reports the result for the two candidates in the run-off in the second round.

Figure D1: INDIVIDUAL LOAN DATA



Notes: The data refer to our sample of 36,564 individual loans belonging to 19,937 unique individuals. Loans are measured in GEL, GEL 1,000 equals about \$385. The Figure on the left shows the size distribution of the loans, leaving out the top five percentile. The Figure on the right shows the average value of the loans relieved by gender and birth cohort.

Figure D2: MAP OF THE 73 GEORGIAN ELECTION DISTRICTS



Notes: The two maps show the election districts of Georgia. The heat map on the left shows the growth in the votes cast for Zurabishvili between the first and the second round of the election. The head map on the right shows the probability that a randomly selected voter benefited from the debt relief program. The regions of Abkhazia and North Ossetia are not under Georgian de facto control, and are, thus, missing from the map.

Appendix E: Estimating the costs of winning and of buying a single vote

In this appendix, we explain how we calculate the monetary cost of winning the election and the cost of buying one vote. The calculations are summarized in Table E1.

To calculate the cost of winning, we proceed in two steps. First, we use the total vote counts in Table D1 to calculate the growth in votes from the first to the second round needed for the Georgian Dream to just win the election (getting the same votes (780,680) as the opposition in the second round) or win with the votes it actual got (1,147,701). To just win the election, the total vote for the party needed to growth with 26.9% ($=100(780,680-615,572)/615,572$); to win with its actual votes, the growth in vote needed is 86.5% ($=100(1,147,701-615,572)/615,572$). Second, we use the point estimate of 9.8067 from Table 1, column 3 to calculate what it would cost in billion GEL to induce such an increase. This estimate is based on our sample of debt, not the total debt forgiven. The foundation claimed that the total debt forgiven was GEL 1,500 million and the debt recorded in our sample is GEL 234.58 which is 15.64% of the announced value. To account for this, we multiply the coefficient of 9.81 with 0.1564 to get the scale coefficient equal to 1.5338. Based on this, it would cost 1000 times $0.2699/1.534 = \text{GEL } 176$ million to tie the election. Using a GEL/\$ exchange rate of 0.385, this is equal to \$67.75 million. All this is based on the book value of the loans. If we assume that the actual cost of buying the loans is about 20% of that, then we get that the cost is \$13.55 million. To win with with 1,147,701 votes, it would cost $0.865/1.5338 = \text{GEL } 564$ million or \$217.14 million with a likely market value of \$43.43 million.

We also calculate the cost of a vote in two steps. First, the growth in votes required to gain one vote starting from the 615,572 votes that the Georgian Dream obtained in the first is $1/615,572$. Second, using the adjusted point estimate of 1.5338 from Table 1,

column 3 estimates the cost of one vote to be GEL 1,060(=1/(615,572*1.5338)) or \$408 based on book values. The likely market value is one fifth of this, i.e., \$81.63 per vote.

Table E1: THE COST OF WINNING AND OF A SINGLE VOTE

	Actual (i.e. from round i to ii)	Hypothetical victory with minimal margin (i.e. 50%+1)
Increase in votes, %	86.50%	26.99%
	In the sample	Announced value
Total debt, million GEL	234.58	1,500.00
Total debt, million USD	90.33	577.59
	Baseline beta (i.e. Table 1, col 3)	Adjusted beta (by sample weight i.e. 15.64%)
Treatment effect, per billion GEL	9.8067	1.5338
	Total costs	
	Actual (i.e. from round i to ii)	Hypothetical victory with minimal margin (i.e. 50%+1)
Book value, million GEL	564.00	175.98
Book value, million USD	217.14	67.75
Market value, million USD	43.43	13.55
	Cost of one vote	
Book value, GEL	1,060.15	
Book value, USD	408.16	
Market value, USD	81.63	

Notes: 1 USD equals 2.597 GEL.



Download ZEW Discussion Papers:

<https://www.zew.de/en/publications/zew-discussion-papers>

or see:

<https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html>

<https://ideas.repec.org/s/zbw/zewdip.html>



IMPRINT

**ZEW – Leibniz-Zentrum für Europäische
Wirtschaftsforschung GmbH Mannheim**

ZEW – Leibniz Centre for European
Economic Research

L 7,1 · 68161 Mannheim · Germany

Phone +49 621 1235-01

info@zew.de · zew.de

Discussion Papers are intended to make results of ZEW research promptly available to other economists in order to encourage discussion and suggestions for revisions. The authors are solely responsible for the contents which do not necessarily represent the opinion of the ZEW.